7705 Service Aggregation Router Gen 2

Release 25.3.R2

Classic CLI Command Reference Guide

3HE 21565 AAAA TQZZA 01
Edition: 01
April 2025

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

Table of contents

1 Getting started 5

2 Classic CLI overview 10

3 Command Trees 43

4 a Commands – Part I 259

5 a Commands – Part II 441

6 b Commands 631

7 c Commands 758

8 d Commands 963

9 e Commands 1235

10 f Commands 1455

11 g Commands 1574

12 h Commands 1649

13 i Commands 1751

14 j Commands 2045

15 k Commands 2048

16 l Commands 2075

17 m Commands – Part I 2353

18 m Commands – Part II 2535

19 n Commands....2726
20 o Commands....2817
21 p Commands – Part I....2899
22 p Commands – Part II....3058
23 p Commands – Part III....3202
24 q Commands....3363
25 r Commands – Part I....3443
26 r Commands – Part II....3601
27 s Commands – Part I....3779
28 s Commands – Part II....4020
29 s Commands – Part III....4247
30 t Commands....4335
31 u Commands....4553
32 v Commands....4620
33 w Commands....4701
34 x Commands....4725
35 y Commands....4727
36 z Commands....4728

1 Getting started

This guide contains command descriptions for the classic SR OS CLI commands that are used to manage the SR OS.

See the *7250 IXR, 7450 ESS, 7705 SAR Gen 2, 7750 SR, 7950 XRS, and VSR Acronyms Reference Guide* for expansions of acronyms used in this guide.

In this guide, the term CLI refers to the classic CLI unless otherwise specified.

This guide does not include **clear**, **monitor**, **show**, or **tools** commands. These commands are documented in the *7705 SAR Gen 2 Clear, Monitor, Show, Tools CLI Command Reference Guide*.

Note: This guide generically covers Release 25.*x*.R*x* content and may contain some content that will be released in later maintenance loads. See the *SR OS R25.x.Rx Software Release Notes*, part number 3HE 21562 000*x* TQZZA, for information about features supported in each load of the Release 25.*x*.R*x* software. For a list of features and CLI commands that are present in SR OS but not supported on the 7705 SAR Gen 2 platforms, see "SR OS Features not Supported on SAR Gen 2" in the *SR OS R25.x.Rx Software Release Notes*.

The full set of CLI commands supported by the SR OS is documented in three related guides that are listed in the following table.

Table 1: Documentation for SR OS CLI commands

Guide title	Classic CLI commands	MD-CLI commands
7705 SAR Gen 2 Clear, Monitor, Show, Tools CLI Command Reference Guide	All **clear**, **monitor**, **show**, and **tools** commands	All **clear**, **monitor**, **show**, and **tools** commands
7705 SAR Gen 2 Classic CLI Command Reference Guide	All other commands	—
7705 SAR Gen 2 MD-CLI Command Reference Guide	—	All other commands

1.1 Platforms and terminology

Note:
Unless explicitly noted otherwise, this guide uses the terminology defined in the following table to collectively designate the specified platforms.

Table 2: Platforms and terminology

Platform	Collective platform designation
7705 SAR-1	7705 SAR Gen 2

1.2 Command tree

The SR OS CLI command tree is a hierarchical inverted tree. The highest level is the root level. Below this level are other tree levels with the major command groups; for example, **configuration** commands and **admin** commands are levels below root.

In the tree, you can click a command to link directly to the command description.

Note: Commands that are listed in the tree but are not linked to an associated description are available on one or more platforms but are not currently described in the guide.

1.3 Command descriptions

Command descriptions are listed in alphabetical order by command name.

The following figure shows an example of a command description.

Figure 1: Command description example

aa-sub-study

Syntax	as-sub-study *study-type*
Context	[Tree] (config>app-assure>group>statistics aa-sub-study)
Full Context	configure application-assurance group statistics aa-sub-study
Description	This command enables the context to configure accounting and statistics collection parameters per application assurance special study subscribers.
Parameters	*study-type* — Specifies special study protocol subscriber stats.
Values	application, protocol

sw3044

The following table describes the fields that may be shown for a command. Not all fields are applicable for all commands.

Table 3: Command description fields

Field	Description
Command Name	Name of the command
Syntax	Command syntax required to execute the command. See Table 4: Command syntax symbols for information about syntax symbols.
Context	Path to the command as it is displayed in the CLI prompt. Clicking on [Tree] links to the command in the CLI tree.
Full Context	Complete contextual path to perform the command
Description	Description of the command functionality and any restrictions

Field	Description
Default	Command default value
Parameters	Descriptions of command parameters
Values	Values allowed for the parameter
Default	Parameter default value
Platforms	Hardware platforms on which the command is available. See Platforms and terminology for more information about the platforms. **Note**: Some SR OS features are platform-specific and therefore may not be available or visible on all platforms. See the *SR OS R25.x.Rx Software Release Notes*, part number 3HE 21562 000*x* TQZZA, for information about platform support.

Note: All options for enumerated types and numerical ranges are listed in the command descriptions; however, not all options or ranges are valid on all platforms.

1.4 Navigational aids

The following aids help you navigate the guide and find specific commands.

1.4.1 Context path

In the CLI tree section, the complete contextual path to the first command on the page is displayed at the top of the page, as shown in the following figure.

Figure 2: Command tree navigation

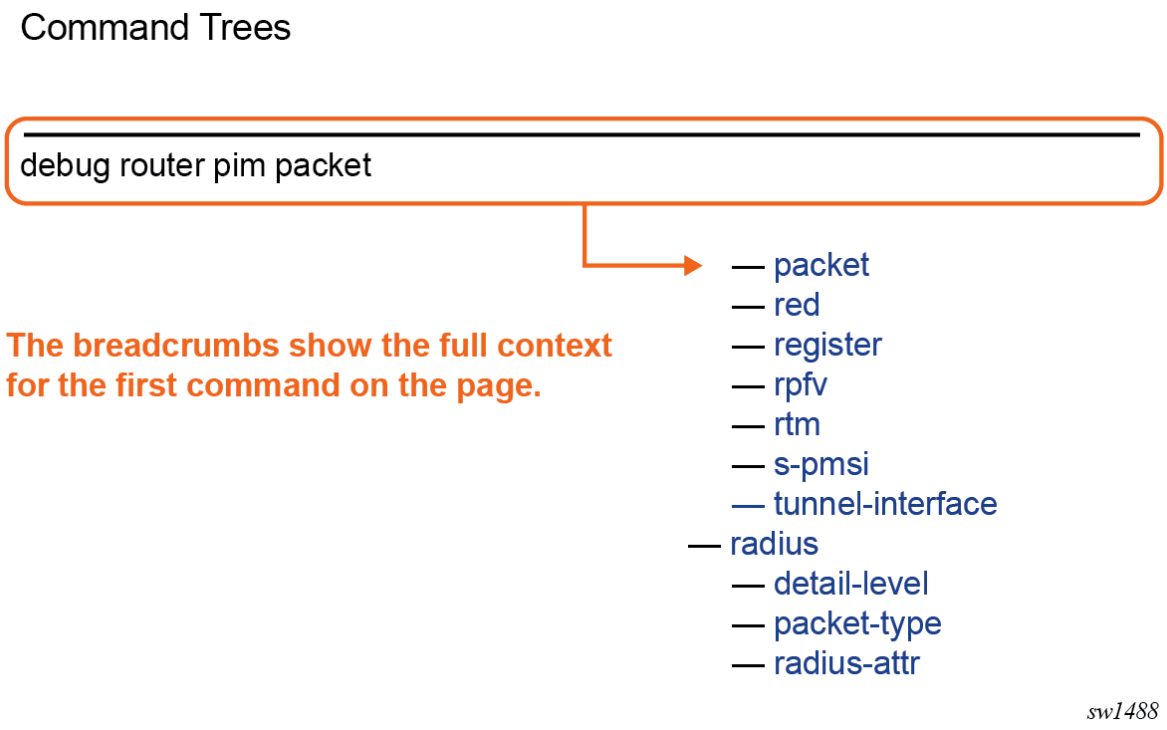

1.4.2 Searching

The Context field in each command description shows, in parentheses, the full path to the command as displayed in the CLI prompt. This form of the contextual path often abbreviates terms. For example:

(cfg>qos>qgrps>egr>qgrp>queue wred-queue)

To search this guide for a specific command using the Acrobat search function, enter the command name and append a closing parenthesis in the search window. For more efficient searching, add the previous level of the contextual path before the command name.

If you add the context and closing parenthesis, the resulting search returns only matching Context entries. It will not return instances of the same command found elsewhere in the guide. The following figure shows an example of a search.

Figure 3: Search window

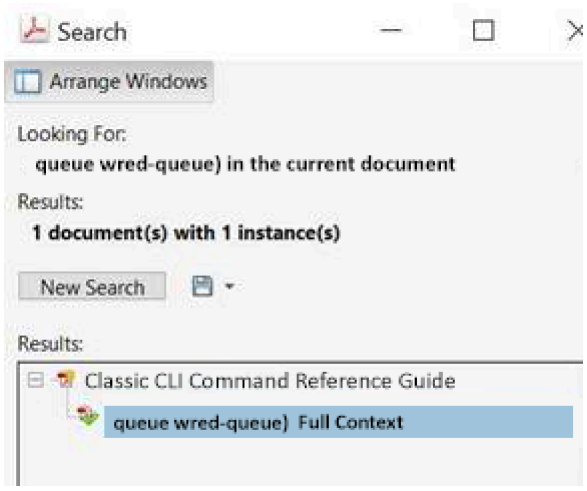

2 Classic CLI overview

In this chapter, "CLI" refers to the classic CLI unless otherwise specified.

2.1 CLI structure

The SR OS CLI is a command-driven interface accessible through the console, Telnet, or secure shell (SSH). The CLI can be used for the configuration and management of routers.

The SR OS CLI command tree is a hierarchical inverted tree. The operational root level is the highest level of this tree. When the user enters a CLI session, the user is in the operational root context. Below this level are other tree levels for the major command groups; for example, **configure** commands and **show** commands are levels below the operational root level.

The CLI is organized so that related commands with the same scope are at the same level or in the same context. Sublevels or subcontexts have related commands with a more refined scope.

Note: The CLI engine used to execute scripts is the primary CLI engine configured with **config**>**system**>**management-interface**>**cli**>**cli-engine** {[**classic-cli**] [**md-cli**]}.

2.2 Navigating in the CLI

The following table describes command syntax symbols used in this guide.

Table 4: Command syntax symbols

Symbol	Description
|	A vertical line indicates that only one of the parameters within the brackets or braces can be selected.
[ ]	Brackets indicate optional parameters.
{ }	Braces indicate that one of the parameters must be selected.
[{ }]	Braces within square brackets indicate that the parameters are optional, but if one is selected, the information within the braces is required.
Bold	**Bold** indicates commands and keywords.
Italic	*Italics* indicates that you must enter text for the parameter.

2.2.1 CLI contexts

Use the CLI to access, configure, and manage Nokia routers. CLI commands are entered at the command line prompt. Access to specific CLI commands is controlled by the permissions set by your system administrator. Entering a CLI command makes navigation possible from one command context (or level) to another.

When the user enters a CLI session, the user is in the operational root context. Navigate to another level by entering the name of successively lower contexts. For example, at the command prompt (#), enter **configure** or **config**. See Command completion for alternative syntax for the same command. The active context displays in the command prompt.

```
A:node-2# config
A:node-2>config#
```

In a CLI context, enter commands at that context level by entering the text. Press <Enter> to move to a lower context. The user can also include commands from lower context at one context level as long as the command and parameter syntax is correct.

The following example shows two methods to navigate to a service SDP ingress level.

Method 1:

```
A:node-2# configure service epipe 6 spoke-sdp 2:6 ingress
*A:node-2>config>service>epipe>spoke-sdp>ingress#
```

Method 2:

```
A:node-2>config# service
A:node-2>config>service# epipe 6
*A:node-2>config>service>epipe# spoke-sdp 2:6
*A:node-2>config>service>epipe>spoke-sdp# ingress
*A:node-2>config>service>epipe>spoke-sdp>ingress#
```

The CLI returns an error message if the syntax is incorrect. For example, if the user enters **rooter** for the **root** command, it would result in an error.

```
*A:node-2>config# rooter
Error: Bad command.
```

The command parameter is the means by which a value is passed to the command processing program. The user must enter the value according to the syntax rules and, where applicable, the defined range. In the previous example, "6" and "2:6" are the parameter values that the user must enter to execute the command. The value "6" is the value for the Epipe service identifier parameter. For the value "2:6", "2" is the value for the SDP identifier parameter and "6" is the value for the virtual circuit identifier.

2.2.2 Operational root and global commands

The commands in the following table are available at the operational root level of the CLI hierarchy. For the command descriptions, see the respective command sections in this guide. For descriptions of the **clear**, **monitor**, **show**, and **tools** commands, see the *7705 SAR Gen 2 Clear, Monitor, Show, and Tools CLI Command Reference Guide*.

Table 5: Operational root commands

Command	Description
admin	Enters the administrative context for system operations
bof	Enters the context to configure the boot options file
clear	Clears statistics or resets the operational state
configure	Enters the configuration context
[**no**] **debug**	Enters the context to enable or disable debugging and specify debug options
environment	Enters the environment configuration context
file	Enters the context for file system commands
help	Displays help in the CLI
monitor	Enters the context to monitor statistics
password	Enters the context to change the user CLI login password
show	Shows operational information
tools	Enters the tools context for troubleshooting and debugging

Global commands are commands that can be entered at any level in the CLI hierarchy. To display the list of all system global commands, enter **help globals** in the CLI.

The global commands are listed in the following table. For the command descriptions, see the respective command sections in this guide.

Table 6: Global commands

Command	Description
back	Navigates to the parent context
candidate	Enters the context to configure candidate parameters
echo	Echoes the text that is typed in. The primary use is to display messages to the screen within an exec file.
enable-admin	Enables the user to become a system administrator
exec	Executes the contents of a text file as if they were CLI commands entered at the console
exit	Returns to the previous higher context
help	Displays help in the CLI

Command	Description
history	Displays a list of the most recently entered commands
logout	Terminates the CLI session
mrinfo	Requests multicast router information
mstat	Traces a multicast path from a source to a receiver and displays multicast packet rate and loss information (IGMP-based)
mstat2	Traces a multicast path from a source to a receiver and displays multicast packet rate and loss information (UDP-based)
mtrace	Traces a multicast path from a source to a receiver (IGMP-based)
mtrace2	Traces a multicast path from a source to a receiver (UDP-based)
oam	Provides OAM test suite options. See the *7705 SAR Gen 2 OAM and Diagnostics Guide*.
ping	Verifies the reachability of a remote host
pwc	Displays the present or previous working context of the CLI session
sleep	Causes the console session to pause operation (sleep) for 1 s or for the specified number of seconds. The primary use is to introduce a pause in the execution of an exec file.
ssh	Opens a secure shell connection to a host
telnet	Connects to a host using Telnet
traceroute	Determines the route to a destination address
tree	Displays a list of all commands at the current level and all sublevels
write	Sends a console message to a specific user or to all users with active console sessions

2.2.3 CLI environment commands

The CLI **environment** commands listed in the following table are found in the **environment** context of the operational root of the CLI tree. These commands control session preferences for a single CLI session. For more information on the commands, see the respective command sections in this guide.

Table 7: CLI environment commands

Command	Description
alias	Enables the substitution of a command line by an alias
create	Enables or disables the use of a create parameter check

Command	Description
kernel	Enables or disables the kernel
more	Enables the CLI output to be displayed one screen at a time, awaiting user input to continue
reduced-prompt	Configures the maximum number of higher-level CLI context nodes to display by name in the CLI prompt for the current CLI session
saved-ind-prompt	Saves the indicator in the prompt
shell	Enables or disables the shell
suggest-internal-objects	Enables the suggestion of internally created objects while auto-completing
terminal	Configures the terminal screen length for the current CLI session
time-display	Specifies whether time should be displayed in local time or UTC
time-stamp	Specifies whether the timestamp should be displayed before the prompt

2.3 Getting help in the CLI

The **help** system commands and the **?** key display different types of help in the CLI. The following table lists the help commands.

Table 8: Online help commands

Command	Description
help	Describes the help system
help globals	Displays information about global commands
help edit	Displays information about all editing keystrokes
help special-characters	Displays information about all special characters that can be entered at the command prompt
?	Displays context-sensitive help information and displays a full list of options and commands available from the current context When entered at the root context, displays a full list of top-level contexts and global commands
command ? *command parameter* ? *command* **keyword** ?	Displays the available syntax options for the command, lists the associated parameters and keywords, and lists all commands available from the *command* context

Command	Description
string?	Lists all commands available in the current context that start with *string*
command string? *command parameter string*? *command* **keyword** *string*?	Lists all commands, parameters, or keywords available in the *command* context that start with *string*
*string***Tab**	Completes a partial command name (auto-completion) or lists available commands that match *string*

The **tree** and **tree detail** system commands are useful when searching for a command in a lower-level context.

The following example shows a partial list of the **tree** and **tree detail** command outputs on the 7705 SAR Gen 2.

Example: tree and tree detail command outputs

```
*A:node-2>config# tree
+---router
|   |
|   +---admin-tags
|   |   |
|   |   +---admin-tag
|   |   |
|   |   +---route-admin-tag-policy
|   |   |   |
|   |   |   +---exclude
|   |   |   |
|   |   |   +---include
|   |
|   +---aggregate
|   |
|   +---allow-bgp-to-igp-export
|   |
|   +---allow-icmp-redirect
|   |
|   +---allow-icmp6-redirect
|   |
|   +---autonomous-system
|   |
|   +---bfd
|   |   |
|   |   +---abort
|   |   |
|   |   +---begin
|   |   |
|   |   +---bfd-template
|   |   |   |
|   |   |   +---echo-receive
|   |   |   |
|   |   |   +---multiplier
|   |   |   |
|   |   |   +---receive-interval
|   |   |   |
|   |   |   +---transmit-interval
|   |   |   |
|   |   |   +---type
```

```
|   |   |
|   |   +---commit
|   |
|   +---bgp
|   |   |
|   |   +---add-paths
|   |   |   |
|   |   |   +---evpn
|   |   |   |
|   |   |   +---ipv4

*A:node-2>config# tree detail
...
|   router <router-instance> [create]
|   |
|   +---admin-tags
|   |   |
|   |   +---admin-tag <tag>
|   |   |   no admin-tag <tag>
|   |   |
|   |   +---route-admin-tag-policy <policy-name>
|   |   |   no route-admin-tag-policy <policy-name>
|   |   |   |
|   |   |   +---exclude <tag>
|   |   |   |   no exclude <tag>
|   |   |   |
|   |   |   +---include <tag>
|   |   |   |   no include <tag>
|   |
|   +---no aggregate <ip-prefix/ip-prefix-length>
|   |   aggregate <ip-prefix/ip-prefix-length> [summary-only] [as-set] [aggregator
 <as-number:ip-address>] [discard-component-communities] [black-hole [generate-icmp]]
 [community <comm-id1> [<comm-id2> <comm-id3> .. up to 12]] [description <description>]
 [local-preference <local-preference>] [tunnel-group <tunnel-group-id>] [policy <policy-
name>]
|   |   aggregate <ip-prefix/ip-prefix-length> [summary-only] [as-set] [aggregator <as-
number:ip-address>] [discard-component-communities] [indirect <ip-address>] [community
 <comm-id1> [<comm-id2> <comm-id3> .. up to 12]] [description <description>] [local-
preference <local-preference>] [tunnel-group <tunnel-group-id>] [policy <policy-name>]
|   |
|   +---allow-bgp-to-igp-export
|   |   no allow-bgp-to-igp-export
|   |
|   +---allow-icmp-redirect
|   |   no allow-icmp-redirect
|   |
|   +---allow-icmp6-redirect
|   |   no allow-icmp6-redirect
|   |
|   +---autonomous-system <autonomous-system>
|   |   no autonomous-system
|   |
|   +---bfd
|   |   |
|   |   +---abort
|   |   |
|   |   +---begin
|   |   |
|   |   +---bfd-template <[32 chars max]>
|   |   |   no bfd-template <[32 chars max]>
|   |   |   |
|   |   |   +---echo-receive <milli-seconds>
|   |   |   |   no echo-receive
|   |   |   |
```

```
|   |   |   +---multiplier <multiplier>
|   |   |   |   no multiplier
|   |   |   |
|   |   |   +---no receive-interval
|   |   |   |   receive-interval <milli-seconds>
|   |   |   |
|   |   |   +---no transmit-interval
|   |   |   |   transmit-interval <milli-seconds>
|   |   |   |
|   |   |   +---no type
|   |   |
|   |   +---commit
|   |
|   +---bgp
|   |   no bgp
|   |   |
|   |   +---add-paths
|   |   |   no add-paths
|   |   |   |
|   |   |   +---no evpn
|   |   |   |   evpn send <send-limit>
|   |   |   |   evpn send <send-limit> receive [none]
|   |   |   |
|   |   |   +---ipv4 send <send-limit>
|   |   |   |   ipv4 send <send-limit> receive [none]
|   |   |   |   no ipv4
```

2.4 The CLI command prompt

By default, the CLI command prompt indicates the device being accessed, the active CPM, and the current CLI context. For example, the prompt: **A:node-2>config>router>if#** indicates that the active CPM is CPM A, the user is on the device with the hostname **node-2**, and the current context **config**>**router**>**interface**. In the prompt, the separator used between contexts is the ">" symbol. The active CPM on the 7705 SAR Gen 2 is CPM A.

At the end of the prompt, there is either a pound sign (#) or a dollar sign ($). A "#" at the end of the prompt indicates the context is an existing context. A "$" at the end of the prompt indicates the context has been newly created. Contexts are newly created for logical entities when the user first navigates into the context.

Because there can be a large number of sublevels in the CLI, the **environment** command **reduced-prompt** *no of nodes in prompt* allows the user to control the number of levels displayed in the prompt.

WARNING: In CLI command configurations, allowed values in parameter strings are printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes ("). Double quotes within a string are not supported. Parameter strings input by the user must follow this format. This rule supersedes all related parameter descriptions found in this guide.

When changes are made to the configuration file, a "*" appears in the prompt string (*A:node-2), indicating that the changes have not been saved. When an **admin save** command is executed, the "*" disappears. This behavior is controlled by the **saved-ind-prompt** command in the **environment** context.

2.5 Displaying configuration contexts

The **info**, **info detail**, and **objective** commands display the configuration for the current level. The **info** command shows non-default configurations. The **info detail** command shows the entire configuration for the current level, including defaults. The **info objective** command provides an output objective that controls the configuration parameters to be displayed.

Example: info and info detail command outputs

The following example displays the output from the **info** command and the **info detail** command.

```
*A:node-2>config>router# interface system
*A:node-2>config>router>if# info
----------------------------------------------
            address 10.10.0.1/32
----------------------------------------------
*A:node-2>config>router>if#


*A:node-2>config>router>if# info detail
----------------------------------------------
            address 10.10.10.103/32 broadcast host-ones
            no description
            no arp-timeout
            no allow-directed-broadcasts
            tos-marking-state trusted
            no local-proxy-arp
            no proxy-arp
            icmp
                mask-reply
                redirects 100 10
                unreachables 100 10
                ttl-expired 100 10
            exit
            no mac
            no cflowd
            no shutdown
----------------------------------------------
*A:node-2>config>router>if#
```

2.6 exec Files

The **exec** command allows you to execute a text file of CLI commands as if it were typed at a console device.

The **exec** command and the associated exec files can be used to conveniently execute a number of commands that are always executed together in the same order. For example, an **exec** command can be used to define a set of commonly used standard command aliases.

The **echo** command can be used within an exec command file to display messages on screen while the file executes.

Arguments can be specified with the **exec** command. These arguments are passed in to be used inside the text file that includes the CLI commands. The passing of arguments with the **exec** command only works in the classic CLI. The passing of arguments with the **exec** command cannot be used in the MD-CLI.

For example, if the file `cf3/Test.txt` contains the following set of CLI commands:

```
echo $(1)
echo $(2)
echo $(3)
```

then executing the following commands:

 # exec cf3:/Test.txt -arguments var1=10 var2=20 var3=30

or

 # exec cf3:/Test.txt -arguments 10 20 30

produces the following output:

```
10
20
30
```

2.7 CLI script control

SR OS provides centralized script management for CLI scripts that are used by CRON and the Event Handling System (EHS). A set of script policies and script objects can be configured to control such things as:

- where scripts are located (local compact FTP server)
- where to store the output of the results
- how long to keep historical script result records
- how long a script may run

If the scripts are located on local compact flash devices, the user must ensure that the scripts are on the compact flash devices of both CPMs so that operation of EHS continues as expected if a CPM switchover occurs.

Only one script can execute at a time. An SNMP table (smRunTable in the DISMAN-SCRIPT-MIB) is used as both an input queue of scripts waiting to be executed and for storage of records for completed scripts. If the input queue is full, the script request is discarded.

2.8 Entering CLI commands

The following sections outline the steps to entering CLI commands.

2.8.1 Command completion

The CLI supports both command abbreviation and command completion. If the keystrokes entered are enough to match a valid command, the CLI displays the remainder of the command syntax when **Tab** or **Spacebar** is pressed. When typing a command, **Tab** or **Spacebar** invokes auto-completion. If the keystrokes entered are sufficient to identity a specific command, auto-completion completes the

command. If the letters are not sufficient to identify a specific command, pressing **Tab** or **Spacebar** displays commands matching the letters entered.

System commands are available in all CLI contexts.

2.8.2 Unordered and unnamed parameters

In a command context, the CLI accepts command parameters in any order as long as the command is formatted in the proper command keyword and parameter syntax. Command completion works as long as enough recognizable characters of the command are entered.

The following output shows the command syntax for **static-route-entry**.

```
*A:node-2>config>router# static-route-entry ?
 - no static-route-entry <ip-prefix/prefix-length> [mcast]
 - static-route-entry <ip-prefix/prefix-length> [mcast]

<ip-prefix/prefix-*> : ipv4-prefix    - a.b.c.d (host bits must be 0)
                       ipv4-prefix-le - [0..32]
                       ipv6-prefix    - x:x:x:x:x:x:x:x   (eight 16-bit pieces)
                                        x:x:x:x:x:x:d.d.d.d
                                        x - [0..FFFF]H
                                        d - [0..255]D
                       ipv6-prefix-le - [0..128]
<mcast>              : keyword - Indicates that static-route being configured
                       is used for mcast table only

 [no] backup-tag      - Create/Configure or Delete/Deconfigure backup tag for
                        static-route-entry
 [no] black-hole      + Create/Configure or Delete/Deconfigure blackhole
                        nexthop for static-route-entry
 [no] community       - Create/Configure or Delete/Deconfigure community for
                        static-route-entry
 [no] indirect        + Create/Configure or Delete/Deconfigure indirect
                        next-hop for static-route-entry
 [no] next-hop        + Create/Configure or Delete/Deconfigure next-hop for
                        static-route-entry
 [no] tag             - Create/Configure or Delete/Deconfigure tag for
                        static-route-entry
```

Some SR OS CLI commands have multiple unnamed parameters. For example, the **subrate** *csu-mode rate-step* command has both a *csu-mode* parameter and a *rate-step* parameter that do not have leading keywords. SR OS uses a best-match algorithm to select which parts of the user input are intended to be used for each unnamed parameter. This best-match algorithm depends on the specific command.

In some cases, it is not possible for the algorithm to be 100% accurate, and SR OS may assign an unintended value to a parameter when two unnamed parameters have similar constraints and syntax. For example, the **environment alias** *alias-name alias-command-name* command may reverse the *alias-name* and *alias-command-name* parameters if the first parameter entered is more than 80 characters.

2.8.3 Using editing keystrokes

When entering a command, special keystrokes allow for editing of the command. The following table lists the command editing keystrokes.

Table 9: Command editing keystrokes

Editing action	Keystrokes
Stop the current command	**Ctrl-C**
Delete current character	**Ctrl-D**
Delete text up to cursor	**Ctrl-U**
Delete text after cursor	**Ctrl-K**
Move to beginning of line	**Ctrl-A**
Move to end of line	**Ctrl-E**
Get prior command from history	**Ctrl-P**
Get next command from history	**Ctrl-N**
Move cursor left	**Ctrl-B**
Move cursor right	**Ctrl-F**
Move back one word	**Alt-B** or **Esc+B**
Move forward one word	**Alt-F** or **Esc+F**
Convert rest of word to uppercase	**Alt-C** or **Esc+C**
Convert rest of word to lowercase	**Alt-L** or **Esc+L**
Delete remainder of word	**Alt-D** or **Esc+D**
Delete word up to cursor	**Ctrl-W**
Transpose current and previous character	**Ctrl-T**
Return to operational root. If using **Ctrl-Z** after a command, return to the operational root after executing the command (equivalent to pressing **Enter** after the command and **exit all** after the command has executed).	**Ctrl-Z**
Refresh input line	**Ctrl-L**

2.8.4 Entering absolute paths

CLI commands can be executed in any context by specifying the full path from the CLI root. To execute an out-of-context command, enter a forward slash (/) or backward slash (\) at the beginning of the command line. The commands are interpreted as absolute paths. Spaces between the slash and the first command return an error.

```
*A:node-2# configure router
```

```
*A:node-2>config>router# interface system address 10.2.3.4
*A:node-2>config>router# /admin save
*A:node-2>config>router# \clear router interface
*A:node-2>config>router#
```

The "/" or "\" cannot be used as an absolute path at the beginning of the command string of the **environment alias** command. The command may change the current context depending on whether it is a leaf command. This is the same behavior the CLI performs when CLI commands are entered individually; for example:

```
*A:node-2# admin
*A:node-2>admin# save
```

or

```
*A:node-2# admin save
*A:node-2#
```

An absolute path command behaves in the same way as manually entering a series of command line instructions and parameters.

For example, beginning in an IES context service ID 4 (IES 4):

```
config>service>ies> /clear card 1
```

behaves in the way same as the following series of commands:

```
config>service>ies>exit all
clear card 1
configure service ies 4 (returns you to your starting point)
config>service>ies
```

If the command takes you to a different context, the following occurs:

```
config>service>ies>/configure service vpls 5 create
```

becomes:

```
config>service>ies>exit all
configure service vpls 5 create
config>service>vpls>
```

2.8.5 Displaying the command history

The CLI maintains a history of the most recently entered commands. The **history** command shows the most recently entered CLI commands.

```
*A:node-2# history
   1 environment terminal length 48
   2 environment no create
   3 show version
   4 configure port 1/1/1
   5 info
   6 \configure router isis
   7 \port 1/1/2
```

```
   8 con port 1/1/2
   9 \con port 1/1/2
  10 \configure router bgp
  11 info
  12 \configure system login-control
  13 info
  14 history
  15 show version
  16 history
*A:node-2# !3

A:node-2# show version
TiMOS-B-0.0.I2016 both/i386 Nokia 7450 ESS Copyright (c) 2000-2016 Nokia
All rights reserved. All use subject to applicable license agreements.
Built on Sun Oct 12 20:01:13 PDT 2008 by builder in /rel0.0/I2016/panos/main
A:node-2#
```

2.8.6 Entering numerical ranges

The SR OS CLI allows the use of a single numerical range as an argument in the command line. This range can be a set or a sequence of numbers, or a combination of both.

A set is a range of numerical values, from a minimum to a maximum, incremented by 1. For example:

```
configure service vpls [1..10] create customer 1
```

A sequence is a list of discrete integer elements, in any order. For example:

```
configure service vpls [1,2,3] no shutdown
```

A sequence can contain sets as well as integer elements. For example:

```
configure service vpls [4..6,7,8..10] no shutdown
```

For example, it is possible to shut down ports 1 through 10 on an XMA/MDA 1 in chassis slot 1. A port can be denoted by "*slot*/*mda*/*port*", where *slot* is the slot number, *mda* is the XMA/MDA number and *port* is the port number. To shut down ports 1 through 10 on an XMA/MDA 1 in slot 1, the command is entered as follows:

```
configure port 1/1/[1..10] shutdown
```

Ctrl-C can be used to abort the execution of a range command.

CLI commands can contain ranges of hexadecimal values. This allows ranges to be used when working with data normally expressed in hexadecimal instead of decimal, such as IPv6 or MAC addresses. For example:

```
#config>service>vpls>sap$ static-mac aa:bb:[0x19..0x21]:dd:ee:ff create
#config>service>vpls>sap$ info
----------------------------------------------
               static-mac aa:bb:19:dd:ee:ff create
               static-mac aa:bb:1a:dd:ee:ff create
               static-mac aa:bb:1b:dd:ee:ff create
               static-mac aa:bb:1c:dd:ee:ff create
               static-mac aa:bb:1d:dd:ee:ff create
               static-mac aa:bb:1e:dd:ee:ff create
               static-mac aa:bb:1f:dd:ee:ff create
```

```
                static-mac aa:bb:20:dd:ee:ff create
                static-mac aa:bb:21:dd:ee:ff create
-------------------------------------------
```

A range can also be a reference to a previous range in the same command. This reference takes the form [$*x*], where *x* is an integer between 0 and 5. For example:

```
configure service vprn [11..20] router-id 10.20.[$0].1
```

This gives vprn 11 the router-id "10.20.11.1", vprn 12 the router-id "10.20.12.1", and so on.

Specifying a range in the CLI does have limitations. These limitations are summarized in the following table.

Table 10: CLI range use limitations

Limitation	Description
Up to 6 ranges (including references) with 20 range elements in each range may be specified in a single command, and they may not combine to more than 1000 iterations of the command.	For example, ports on two adapter cards can be shut down in one command by using two ranges: **configure port 1/[1..2]/[1..10]**
Ranges within quotation marks are interpreted literally.	In the CLI, enclosing a string in quotation marks ("*string*") causes the string to be treated literally and as a single parameter. For example, several commands in the CLI allow the configuration of a descriptive string. If the string is more than one word and includes spaces, it must be enclosed in quotation marks. A range that is enclosed in quotes is also treated literally. For example, **configure router interface "A[1..10]" no shutdown** creates a single router interface with the name "A[1..10]". However, a command such as: **configure router interface A[1..10] no shutdown** creates 10 interfaces with names A1, A2 .. A10.
Command completion does not work when entering a range.	After entering a range in a CLI command, command and key completion, which normally occurs by pressing **Tab** or **Spacebar**, does not work. If the command line entered is correct and unambiguous, the command works properly; otherwise, an error is returned.

2.8.6.1 Using regular expressions in numerical ranges

The user can include a regular expression inside the numerical ranges of any **clear**, **config**, **show**, or **tools** CLI commands. The beginning and ending of the regular expression must be delimited with the "/" symbol.

SR OS performs the following steps:

- auto-completes the command to get all the possible names
- performs a match of the regular expression against all the names
- executes the command for the names for which the match was successful

Note: The order of execution is the same as the order in which the names are listed in the output display of the CLI **info** command or in the output display when you invoke the auto-complete function using **Tab**. If the execution of the command fails for one of the matching object names, the execution is aborted and the remaining matching object names are not processed.

For example, the following SR-TE LSP names are configured on the router:

```
*A:bkvm35# show router mpls sr-te-lsp
===============================================================================
MPLS SR-TE LSPs (Originating)
===============================================================================
LSP Name                        To              Tun      Protect     Adm    Opr
                                                Id       Path
-------------------------------------------------------------------------------
sr-te-pce                       192.0.2.198     1        N/A         Up     Dwn
RENO194_DET190_LSP1_Profile10   192.0.2.190     2        N/A         Up     Dwn
RENO194_DET190_LSP3             192.0.2.190     3        N/A         Up     Dwn
RENO194_ATL224_LSP1             192.0.2.224     4        N/A         Up     Dwn
-------------------------------------------------------------------------------
LSPs : 4
===============================================================================
```

The following command displays the subset of all SR-TE LSPs with names that include the expression "LSP":

show router mpls sr-te-lsp [/LSP/]

The SR OS expands this command into the following individual commands:

show router mpls sr-te-lsp RENO194_DET190_LSP1_Profile10

show router mpls sr-te-lsp RENO194_DET190_LSP3

show router mpls sr-te-lsp RENO194_ATL224_LSP1

The output of the three **show** commands is displayed in the following example:

```
*A:bkvm35# show router mpls sr-te-lsp [/LSP/]
===============================================================================
MPLS SR-TE LSPs (Originating)
===============================================================================
LSP Name                          To              Tun     Protect   Adm  Opr
                                                  Id      Path
-------------------------------------------------------------------------------
RENO194_DET190_LSP1_Profile10     192.0.2.190     2       N/A       Up   Dwn
-------------------------------------------------------------------------------
LSPs : 1
===============================================================================
===============================================================================
MPLS SR-TE LSPs (Originating)
===============================================================================
LSP Name                          To              Tun     Protect   Adm  Opr
                                                  Id      Path
-------------------------------------------------------------------------------
RENO194_DET190_LSP3               192.0.2.190     3       N/A       Up   Dwn
-------------------------------------------------------------------------------
LSPs : 1
```

```
===============================================================================
===============================================================================
MPLS SR-TE LSPs (Originating)
===============================================================================
LSP Name                            To              Tun     Protect   Adm  Opr
                                                    Id      Path
-------------------------------------------------------------------------------
RENO194_ATL224_LSP1                 192.0.2.224     4       N/A       Up   Dwn
-------------------------------------------------------------------------------
```

2.8.6.1.1 Regular expression symbols in a regular expression match operation

The user can use all the regular expression symbols listed in Table 12: Regular expression symbols and Table 13: Character class expressions inside the regular expression to match.

For example, the user can list all LSP names that begin with the string "RENO194_" followed by the string "ATL" as follows:

```
*A:bkvm35# show router mpls sr-te-lsp [/^RENO194_\['ATL'\]/]
===============================================================================
MPLS SR-TE LSPs (Originating)
===============================================================================
LSP Name                        To              Tun     Protect     Adm    Opr
                                                Id      Path
-------------------------------------------------------------------------------
RENO194_ATL224_LSP1             38.120.48.224   4       N/A         Up     Dwn
-------------------------------------------------------------------------------
LSPs : 1
===============================================================================
```

Note: The following conventions are used in the previous example.

- Use the character "^", which matches the start of the string, directly inside the regular expression to indicate a match at the start of the string. However, if you want to match it as a character, enter it as "\\^".
- Use the range delimiter with the escape symbol in front "\[" inside the regular expression because the range delimiter encloses the regular expression.

The following table summarizes special rules governing the use of some of the regular expression symbols inside a regular expression match operation. Any symbol from Table 12: Regular expression symbols or Table 13: Character class expressions that is not listed in Table 11: Rules governing regular expression symbols can be used directly inside a regular expression match operation.

Table 11: Rules governing regular expression symbols

String	Description
?	[/\?/] if using as a regular expression and [/\\\?/] if using to match the character ?
[ ]	[/\[\]/] if using as a regular expression and [/\\\[\\\]/] if using to match the characters []
$	[/$/] if using as a regular expression and [/\\$/] if using to match the character $

String	Description
\	[/\\\\/] if using to match the character \
/	[/\/] if using to match the character /
'	[/\'/] if using to match the character '
*	[/*/] if using to match the character *
.	[/\./] if using as a regular expression and [/\\\./] if using to match the character .
+	[/\\+/] if using to match the character +
,	[/\,/] if using to match the character ,
^	[/\\^/] if using to match the character ^
(	[/\\(/] if using to match the character (
)	[/\\)/] if using to match the character)
space	[/\ /] if using to match the character space

The SR OS does not support the combination of a partial string with a regular expression match operation.

For example, if the operator wants to display the SR-TE LSP names that begin with the string "RENO194_ATL", if part of the string is entered directly and the rest of the string is entered inside a regular expression, the command returns no match. The following example demonstrates the incorrect syntax:

```
*A:bkvm35# show router mpls sr-te-lsp RENO194_[/ATL/]
```

To obtain a match, the entire string must be inside the regular expression. The following example demonstrates the correct syntax for finding a match:

```
*A:bkvm35# show router mpls sr-te-lsp [/^RENO194_ATL/]
===============================================================================
MPLS SR-TE LSPs (Originating)
===============================================================================
LSP Name                        To              Tun      Protect     Adm    Opr
                                                Id       Path
-------------------------------------------------------------------------------
RENO194_ATL224_LSP1             38.120.48.224   4        N/A         Up     Dwn
-------------------------------------------------------------------------------
LSPs : 1
===============================================================================
```

2.8.7 Using the | match output modifier

The **| match** output modifier searches for a character string or pattern. When using the **| match** output modifier, the variables and attributes must be spelled correctly. The attributes follow the output modifier and must come before the expression or pattern. The following are examples of how to use the **| match** output modifier to complete different tasks:

- Task: Capture all the lines that include "echo" and redirect the output to a file on the compact flash:

  admin display-config | match "echo" > cf1:\test\echo_list.txt
- Task: Display all the lines that do not include "echo":

  admin display-config | match invert-match "echo"
- Task: Display the first match of "vpls" in the configuration file:

  admin display-config | match max-count 1 "vpls"
- Task: Display everything in the configuration after finding the first instance of "interface":

  admin display-config | match post-lines 999999 interface
- Task: Display a count of the total number of lines of output instead of displaying the output itself:

  admin display-config | match interface | count

Command syntax:

match *pattern* **context** {**parents** | **children** | **all**} [**ignore-case**] [**max-count** *lines-count*] [**expression**]

match *pattern* [**ignore-case**] [**invert-match**] [**pre-lines** *pre-lines*] [**post-lines** *lines-count*] [**max-count** *lines-count*] [**expression**]

where:

```
pattern         string or regular expression
context         keyword:  display context associated with the matching line
parents         keyword:  display parent context information
children        keyword:  display child context information
all             keyword:  display both parent and child context information
ignore-case     keyword
max-count       keyword:  display only a specific number of instances of matching
                          lines
lines-count     1 — 2147483647
expression      keyword:  pattern is interpreted as a regular expression
invert-match    keyword
pre-lines       keyword:  display some lines prior to the matching line
pre-lines       0 — 100
post-lines      keyword:  display some lines after the matching line
lines-count     1 — 2147483647
```

For example:

```
A:Dut-C# show log log-id 98 | match ignore-case "sdp bind"
"Status of SDP Bind 101:1002 in service 1001 (customer 1) changed to admin=up oper=u
p flags="
"Processing of a SDP state change event is finished and the status of all affected S
DP Bindings on SDP 101 has been updated."


A:Dut-C# show log log-id 98 | match max-count 1 "service 1001"
"Status of service 1001 (customer 1) changed to administrative state: up,
operational state: up"

A:Dut-C# admin display-config | match post-lines 5 max-
count 2 expression "OSPF.*Config"
echo "OSPFv2 Configuration"
#------------------------------------------------
        ospf
            timers
                spf-wait 1000 1000 1000
            exit
```

```
echo "OSPFv2 (Inst: 1) Configuration"
#--------------------------------------------------
        ospf 1
            asbr
            router-id 10.0.0.1
            export "testall"
*A:Dut# admin display-config | match debug_mirror
            profile "debug_mirror"


*A:Dut# admin display-config | match context parent debug_mirror
#--------------------------------------------------
    system
        security
            profile "debug_mirror"


*A:Dut# admin display-config | match context all debug_mirror
#--------------------------------------------------
    system
        security
            profile "debug_mirror"
                default-action deny-all
                entry 10
                exit


*A:Dut# show log event-control | match ignore-case pre-lines 10 SyncStatus
L  2016 tmnxLogOnlyEventThrottled        MA  gen          0           0
MCPATH:
   2005 tmnxMcPathSrcGrpBlackHole        MI  thr          0           0
   2006 tmnxMcPathSrcGrpBlackHoleCleared MI  thr          0           0
   2007 tmnxMcPathAvailBwLimitExceeded   MI  thr          0           0
   2008 tmnxMcPathAvailBwLimitCleared    MI  thr          0           0
MC_REDUNDANCY:
   2001 tmnxMcRedundancyPeerStateChanged WA  gen          0           0
   2002 tmnxMcRedundancyMismatchDetected WA  gen          0           0
   2003 tmnxMcRedundancyMismatchResolved WA  gen          0           0
   2004 tmnxMcPeerSyncStatusChanged      WA  gen          0           0
```

The following table describes regular expression symbols and their interpretation (similar to what is used for route policy regexp matching). Table 13: Character class expressions describes character class expressions.

Table 12: Regular expression symbols

String	Description
.	Matches any single character
[ ]	Matches a single character that is contained within the brackets [abc] matches "a", "b", or "c" [a-z] matches any lowercase letter [A-Z] matches any uppercase letter [0-9] matches any number
[^]	Matches a single character that is not contained within the brackets [^abc] matches any character other than "a", "b", or "c"

String	Description
	[^a-z] matches any single character that is not a lowercase letter
^	Matches the start of the line (or any line, when applied in multiline mode)
$	Matches the end of the line (or any line, when applied in multiline mode)
()	Defines a "marked subexpression" Every matched instance will be available to the next command as a variable
*	A single character expression followed by "*" matches zero or more copies of the expression
{m,n}	Matches least *m* and at most *n* repetitions of the term
{m}	Matches exactly *m* repetitions of the term
{m,}	Matches *m* or more repetitions of the term
?	The preceding item is optional and matched at most once
+	The preceding item is matched one or more times
-	Used between start and end of a range
\	An escape character to indicate that the following character is a match criteria and not a grouping delimiter

Table 13: Character class expressions

Character class	Characters matched[1]	Description
[:alnum:]	'ABCDEFGHIJKLMNOPQRSTUVWXYZ abcdefghijklmnopqrstuvwxyz 0123456789'	Alphanumeric characters
[:alpha:]	'ABCDEFGHIJKLMNOPQRSTUVWXYZ abcdefghijklmnopqrstuvwxyz'	Alphabetic characters
[:blank:]	' \t'	**Spacebar** and **Tab**
[:cntrl:]	'\007\b\t\n\v\f\r\1\2\3\4\5\6\16\17\20 \21\ 22\23\24\25\26\27\30 \31\32\33\34\35\ 36\37\177'	Control characters
[:digit:]	'0123456789'	Digits

[1] Characters matching the character class are delimited by single quotation marks (').

Character class	Characters matched[1]	Description
[:graph:]	'ABCDEFGHIJKLMNOPQRSTUVWXYZ abcdefghijklmnopqrstuvwxyz 0123456789 !\"#$%&'()*+,-./:;<=>?@[\\ ]^_`{|}~'	Visible characters
[:lower:]	'abcdefghijklmnopqrstuvwxyz'	Lowercase letters
[:print:]	'ABCDEFGHIJKLMNOPQRSTUVWXYZ abcdefghijklmnopqrstuvwxyz 0123456789!\"#$%&'()*+,-./:;<=>?@[\\]^_ `{|}~ '	Visible characters and the Space character
[:punct:]	'!\"#$%&'()*+,-./:;<=>?@[\\]^_`{|}~'	Punctuation characters
[:space:]	'\t\n\v\f\r '	Whitespace (blank) characters
[:upper:]	'ABCDEFGHIJKLMNOPQRSTUVWXYZ'	Uppercase letters
[:xdigit:]	'0123456789ABCDEFabcdef'	Hexadecimal digits

Character class expressions must be enclosed within brackets. The expression [[:digit:]] is treated as a regular expression containing the character class "digit", while [:digit:] is treated as a regular expression matching ":", "d", "i", "g", or "t".

2.8.8 Using the | count output modifier

The **| count** output modifier displays a count of the number of lines that would have otherwise been displayed. The **| count** output modifier is particularly useful when used in conjunction with the **| match** output modifier in order to count the number of output lines that match a specified pattern.

The following example shows usage of the **| count** output modifier.

Example: Using the | count output modifier

```
*A:dut-c# show service service-using vprn

===============================================================================
Services [vprn]
===============================================================================
ServiceId    Type      Adm  Opr  CustomerId Service Name
-------------------------------------------------------------------------------
1            VPRN      Down Down 1
44           VPRN      Up   Up   1
100          VPRN      Down Down 1
102          VPRN      Up   Up   1
235          VPRN      Down Down 1
1000         VPRN      Down Down 1000
-------------------------------------------------------------------------------
Matching Services : 6
-------------------------------------------------------------------------------
===============================================================================
*A:dut-c# show service service-using vprn | match Down | count
```

[1] Characters matching the character class are delimited by single quotation marks (').

```
Count: 4 lines
*A:dut-c#
```

2.8.9 Using the | reverse-dns output modifier

The **| reverse-dns** output modifier performs a reverse DNS lookup on any IPv4 or IPv6 address in the input. The result of the lookup is inserted as the next line in the output on each line where an IP address is identified. If no match is found, no additional output is printed. If the output line is more than 80 characters, the line is truncated.

The following example shows usage of the **| reverse-dns** output modifier.

Example: Using the | reverse-dns output modifier

```
A:node-2# ping 10.184.216.34 | reverse-dns
PING 10.184.216.34 56 data bytes
(10.184.216.34) www.example.com
64 bytes from 10.184.216.34: icmp_seq=1 ttl=61 time=82.4ms.
64 bytes from 10.184.216.34: icmp_seq=2 ttl=61 time=82.5ms.
64 bytes from 10.184.216.34: icmp_seq=3 ttl=61 time=82.4ms.
64 bytes from 10.184.216.34: icmp_seq=4 ttl=61 time=82.3ms.
64 bytes from 10.184.216.34: icmp_seq=5 ttl=61 time=82.2ms.
---- 10.184.216.34 PING Statistics ----
(10.184.216.34) www.example.com
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 82.2ms, avg = 82.4ms, max = 82.5ms, stddev = 0.122ms
```

2.8.10 Redirecting output to a file

SR OS supports output redirection (>), which allows the operator to store the output of a CLI command as a local or remote file.

For example:

```
ping 10.0.0.1 > cf3:/ping/result.txt
```

In some cases, only part of the output might be desirable. The **| match** and output redirection commands can be combined. For example, the following command matches lines with the expression "time.[[:digit:]]+" and redirects the output to the file `cf3:/ping/time.txt`.

```
ping 10.0.0.1 | match expression "time.[[:digit:]]+" > cf3:/ping/time.txt
```

2.9 Configuration rollback

The Configuration Rollback feature provides the ability to undo configurations and revert to previous router configuration states while minimizing impacts to services.

This feature gives the operator better control and visibility over the router configurations and reduces operational risk while increasing flexibility and providing powerful recovery options.

Configuration Rollback is useful in cases where configuration changes are made but the operator later decides not to keep the changes (for example, experimentation or when problems are identified in the configuration during actual network operation).

The advantages of this feature include the following.

- Changes made to the router configuration are performed with minimal impact on services being provided by the 7705 SAR Gen 2 because the router does not need to be rebooted.
- There is no impact in areas of the configuration that did not change.

With the rollback feature, the operator can smoothly revert to previous configurations.

Configuration parameters that are changed or items that the changed configuration have dependencies on are removed (revert to default) and the previous values are restored, which may briefly impact services in changed areas).

A history of changes is preserved using checkpoint IDs, which allow rollback to different points, as well as examination of changes made, as shown in the following figure.

Figure 4: Rollback operation

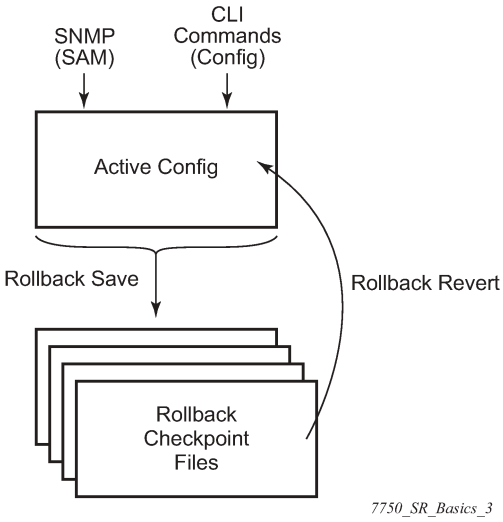

2.9.1 Feature behavior

The following list describes detailed behavior of the rollback feature, including the applicable CLI commands.

- The user can create a rollback checkpoint and later revert to this checkpoint with minimal impact to services.

```
admin>rollback# save [comment comment-string]
```

```
comment-string: a comment associated with the checkpoint, maximum 255 characters
```

- Rollback checkpoints include all current operationally active configurations:
  - changes from direct CLI commands in the configuration branch

  - SNMP sets

- Rollback checkpoints do not include BOF configurations. The BOF file and BOF configuration are not part of a rollback save or rollback. A rollback does not change the BOF configuration. The BOF contains basic information for the node and does not change frequently; changes are mostly made during initial commissioning of the node.
- A rollback save feature can be automatically executed (for example, scheduled monthly) using the CRON facility of SR OS.
- The latest rollback checkpoint file uses the suffix ".rb". The next latest rollback checkpoint file has the suffix ".rb.1", the next oldest has the suffix "rb.2", and so on.

  file-url.rb <--- latest rollback file

  file-url.rb.1

  …

  file-url.rb.9 <--- oldest rollback file
- When a **rollback save** is executed, the system shifts the file suffix of all the previous checkpoints by 1 (new ID = old ID + 1). If there are already as many checkpoint files as the maximum number supported, the last checkpoint file is deleted.
- The maximum number of rollback checkpoints is configurable and defaults to 10, which includes the latest file and files 1 through 9.
- The locations and names of the rollback checkpoint files are configurable to be local (on the compact flash) or remote. The *file-url* must contain a path/directory and filename with no suffix. The .rb suffix for rollback checkpoint files is automatically appended to the rollback checkpoint files.

  ```
  config>system>rollback# rollback-location file-url
  ```
- There is no default rollback location. If a rollback location is not specified, or it is cleared using **no rollback-location**, and a **rollback save** is attempted, the **rollback save** fails and returns an error message.
- The entire set of rollback checkpoint files can be copied from the active CPM CF to the standby CPM CF. This synchronization is done using the following command:

  ```
  admin>redundancy# rollback-sync
  ```
- The operator can enable an automatic synchronization of rollback checkpoint files between the active CPM and standby CPM. When this automatic synchronization is enabled, a **rollback save** causes the new checkpoint file to be saved to both the active and standby CPMs. The suffixes of the old checkpoint files on both active and standby CPMs are incremented.

  Automatic synchronization only causes the new checkpoint file to be copied to both CFs. The older checkpoint files are not automatically copied from active to standby but can be copied manually with **admin redundancy rollback-sync**.
- The **config**>**redundancy**>**synchronize** {**boot-env** | **config**} and **admin**>**redundancy**>**synchronize** {**boot-env** | **config**} commands do not apply to rollback checkpoint files. These commands do not manually or automatically synchronize rollback checkpoint files. The dedicated **rollback-sync** commands must be used to synchronize rollback checkpoint files.

- Rollback files can be deleted using a dedicated rollback checkpoint deletion command.

```
admin>rollback# delete {latest-rb | checkpoint-id}
```

  - Deleting a rollback checkpoint causes the suffixes to be adjusted (decremented) for all checkpoints older that the one that was deleted in order to close the "hole" in the list of checkpoint files and create room to create another checkpoint.
  - If **config**>**redundancy**>**rollback-sync** is enabled, a **rollback delete** also deletes the equivalent checkpoint on the standby CF and shuffles the suffixes on the standby CF.
  - If an operator manually deletes a rollback checkpoint file using **file delete**, the suffixes of the checkpoint files are not shuffled, nor is the equivalent checkpoint file deleted from the standby CF. This manual deletion creates a "hole" in the checkpoint file list until enough new checkpoints have been created to roll the "hole" off the end of the list.
- The following figure shows how a configuration is rolled back to a previous configuration (a saved rollback checkpoint). The previous configuration is loaded and takes operational effect.

```
admin>rollback# revert [latest-rb | checkpoint-id]
```

Figure 5: Configuration rollback

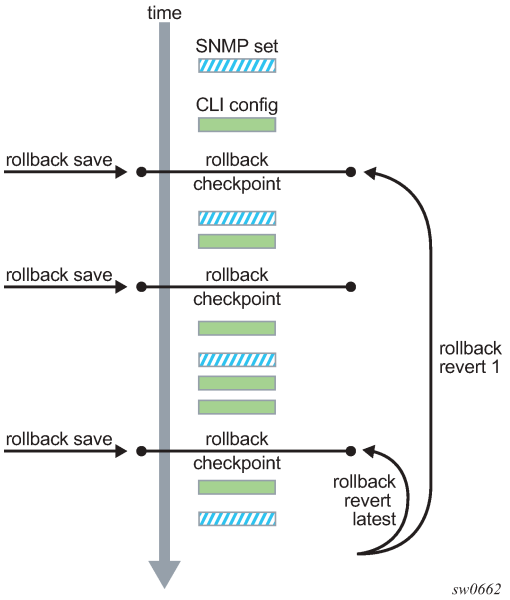

- A rollback revert does not affect the currently stored rollback checkpoint files; files are not deleted or renumbered. This means that if an operator issues the command **rollback revert 3** and then issues the **rollback save** command, the resulting rollback checkpoint files *file-url*.rb and *file-url*.rb.4 contain the same rollback state/configuration.
- The **boot-good-exec** or **boot-bad-exec** command is not automatically executed after a rollback.
- Impacts to the running services are minimized during a rollback.
  - There is no impact in areas of the configuration that did not change.

  - Configuration parameters that are changed or items that the changed configuration has dependencies on are removed (revert to default) and the previous values are restored, which may briefly impact services.
- A rollback undoes any SNMP sets or direct CLI configuration commands that occur after the creation of the last checkpoint.
- When a node is processing a **rollback revert**, both CLI commands from other users and SNMP commands continue to be processed. The only commands that are blocked during a **rollback revert** are other rollback commands, including **revert**, **save**, and **compare**. Only one **rollback** command can be executing at a time on a node.
- Commands are available to view and compare the various rollback checkpoints to current operating and candidate configurations.
- Rollback checkpoint files are not guaranteed to be in any particular format. They are not interchangeable with normal configuration files or executable scripts. A normal configuration file from an **admin save** cannot be renamed as a rollback checkpoint and then referenced for a **rollback revert** operation. Only rollback checkpoint files generated with **rollback save** can be used for rollback revert operations.
- If a hardware change is made after a **rollback save**, then:
  - a rollback can be executed as long as the hardware change was an addition of hardware to the node (for example, a new card or IOM was installed into a previously empty slot)
  - a rollback is not guaranteed to work if hardware was removed or changed (for example, an XCM/IOM was removed, or an XMA/MDA was swapped for a different XMA/MDA type)
- Rollback across a change to the following parameters is not supported:
  - **chassis-mode**
  - **configure isa application-assurance-group minimum-isa-generation**
- Rollback is supported even after an **admin reboot** is performed or the primary configuration in the BOF is changed and an **admin reboot** is performed. The **admin reboot** command does not "break the chain" for rollback.
- Lawful intercept configuration under the **config>li** branch is not affected by a rollback or rescue. LI configuration is not saved in the rollback checkpoint or rescue file, and a rollback revert does not affect any configuration under the **config>li** branch.
- Any configuration or state change performed under the debug branch of the CLI is not saved in the rollback checkpoint file or impacted by a rollback.
- Rollbacks to a checkpoint created in a more recent release are not supported (for example, a node running in 25.3.R1 cannot roll back to a checkpoint created in 25.3.R3).
- The following list captures some side effects and specific behaviors of a rollback revert. Some of these side effects are not related purely to configuration, that is, in the CLI configuration branch, and may have interactions with tools commands, RADIUS, and so on.
  - SAA jobs that are running when a rollback revert is initiated, and need configuration changes due to the rollback, are stopped. If the SAA job is a continuous type, then it is restarted as part of the rollback revert after the configuration changes have been applied (just as if the operator had typed **no shutdown** for the continuous SAA job). Non-continuous SAA jobs that were modified by the rollback would need to be manually restarted if they need to be run again.
  - If **max-nbr-mac-addr** is reduced as part of the revert and the number of MAC addresses in the forwarding database is greater than the **max-nbr-mac-addr**, the rollback is aborted before any

actions are taken and an informative error message is provided. The operator must take actions to remove the MAC addresses if they wish to proceed with the rollback.

– If active subscribers or subscriber hosts or DHCP lease states are present, some associated configuration changes may be blocked, just as those same changes would be blocked if an operator tried to make them using CLI.

– When trying to delete an SLA profile being used by active subscriber hosts, or trying to change a NAT policy in a subscriber profile, if certain configuration changes associated with the hosts or lease states are required as part of the rollback but those changes are blocked, then for each blocked configuration item a warning is printed, that particular configuration item is not changed, and the rollback continues.

– After a multi-chassis peer shutdown or if configuration changes have been made that affect the contents of the distributed database (for example, synchronization tag creation or deletion), further configuration changes related to that peer may be temporarily refused. The duration of the temporary configuration freeze depends on the size of the distributed database. A rollback attempting to make those refused configuration changes fails and an error message is provided to the CLI user.

– If a **force-switchover** command (for example, **tools perform service id 1 endpoint "x" force-switchover spoke-sdp-fec 1**) has been applied to a **spoke-sdp-fec** of a dynamic multi-segment pseudowire, and a rollback revert needs to change the admin state of the **spoke-sdp-fec** (for example, to modify **spoke-sdp-fec** parameters that may be dependent on the admin state), the rollback revert automatically removes the **force-switchover** and the node reverts to whatever is the best spoke SDP in the redundant set.

– Rollback impacts the configuration state of the router, and as with normal operator CLI or SNMP configuration changes, additional actions or steps may need to occur before certain configuration changes take operational effect.

  Configuration changes that require a manual **shutdown** and then **no shutdown** in order to take operational effect also need this manual **shutdown**/**no shutdown** in order to take operational effect after a rollback if the rollback changes those configuration items. Some examples include the following.

  • Changes to autonomous system or confederation values require a BGP **shutdown**/**no shutdown** command.

  • Changes to VPRN **max-routes** require a **shutdown**/**no shutdown** command on the VPRN service.

  • Changes to OSPF or IS-IS **export-limit** require a **shutdown**/**no shutdown** command on OSPF or IS-IS.

  • For configuration changes to an MSAP policy that normally require a **tools perform subscriber-mgmt eval-msap** command to take operational effect on subscribers that are already active, if a rollback changes the MSAP policy configuration, the operator must run the **eval-msap** tools command to have the changes applied to the active subscribers.

– Any uncommitted changes (for example, the **begin** command was entered and some changes were made, but the **commit** command was never entered) in the following areas are lost or cleared when a rollback revert is initiated:

  • **config>app-assure>group policy**

  • **config**>**router**>**policy-options**

  • **config**>**system**>**sync-if-timing**

- Some **card** and **mda** commands require a reboot, removal, or rebuild of an entire card or XMA/MDA. When these commands need to be executed as part of a rollback, the impacted cards and MDAs are listed in a warning and the operator is prompted with a single y/n prompt to proceed. This prompt will not occur for a rollback initiated via SNMP or if the operator uses the **now** keyword with the **rollback revert** command. Some examples of **card** and **mda** commands that may cause a prompt are:
  - **config>card>card-type**
  - **config>card>mda**
  - **config>card>mda>mda-type**
- Although the use of the **Ctrl-C** key combination is not recommended during a rollback revert, it is supported via CLI or SNMP. Interrupting a rollback revert may leave the router in a state that is not necessarily between the old active configuration and the rollback checkpoint, as the rollback processing may have been in the middle of tearing things down or rebuilding configurations. A strong warning is issued in this case to indicate that the operator must examine the configuration and potentially issue another rollback revert to return to a known and coherent configuration.
- An HA CPM switchover during a rollback revert causes the rollback operation to abort. The newly active CPM has an indeterminate configuration. When an HA switchover occurs during a rollback or within a few seconds of a rollback completing, the operator is advised to repeat the rollback revert operation to the same checkpoint.
- A rollback revert operation does not check authorization of each command that is applied during the revert. Permission to execute the revert operation, that is, authorization to execute the **admin rollback revert** command, should only be given to users who are allowed to initiate a rollback revert. It is generally recommended that only system administrators be allowed access to the file system where the rollback files are stored so that they cannot be manually edited.

2.9.2 Rollback and SNMP

The SR OS has SNMP support for rollback status and control. See the TIMETRA-SYSTEM-MIB for details (for example, items such as tmnxSysRollbackStarted).

When the router is doing a rollback revert, SNMP managers see a tmnxSysRollbackStarted trap, then a rapid set of "config change" traps, and then finally, the tmnxSysRollbackStatusChange trap.

During the period when a router is processing a rollback revert, both CLI commands from other users and SNMP commands continue to be processed.

2.9.3 Rescue configuration

A special rescue configuration checkpoint can be created that an operator can revert to at any time. The rescue configuration has its own keyword (**rescue**) and does not use the same rolling suffix indices as the normal rollback checkpoints. This allows the operator to easily return to the rescue configuration state without having to consider a checkpoint index, and ensures that the rescue checkpoint is always available and does not roll off the bottom of the list of checkpoints.

The operator should define a basic rescue configuration that is known to work and give appropriate management access to the node.

The location and filename of the rescue file are configurable. The SR OS appends the .rc suffix to the specified rescue filename.

2.9.4 Operational guidelines

The following points offer some operational guidance on the use of rollback.

- The **admin save** and **admin rollback save** commands should be performed periodically.
- The **admin save** command can be used to back up a complete configuration file that can be used during router reboot, with the following considerations:
  - used with a reboot as a last resort
  - performed after any major hardware changes or major service changes
  - performed after any software upgrade
- The **admin rollback save** command can be used to create a rollback checkpoint as follows:
  - to be used for intermediate checkpoints that can be recovered with minimal impact to services
  - to be performed each time a moderate number of configuration changes have been made
  - to be performed after any hardware changes
  - to be performed after any software upgrade
  - to be scheduled with CRON (for example, once every one or two weeks)
- A new **admin rollback save rescue** must be created when hardware is changed.
- Rollback checkpoint files are not editable, or compatible or interchangeable with configuration files generated with **admin save**.
- The repeated use of the **admin rollback save**, **admin rollback delete**, and **admin rollback revert** commands over the course of weeks or months is not recommended without also executing an occasional **admin save**. In a serious situation, use one of the saved configurations as the primary configuration for an **admin reboot**.
- For a software upgrade, it is recommended that a rollback checkpoint be created using **admin rollback save** in addition to saving the configuration with **admin save**, after an upgrade has been performed and the system is operating as expected. This ensures that a good checkpoint that is fully compatible with the new release is available at a point shortly after the upgrade.
- An operator can create a set of rollback checkpoints to support busy or quiet days or weekends or weekdays and use CRON to shift between them.
- It is beneficial to create a rollback checkpoint before a rollback revert is initiated, especially if significant configuration changes have been applied since the last checkpoint was created. If the rollback is especially significant, that is, there are a lot of major changes, it is also a good practice to perform an **admin save** in case a full reboot is required to recover from an issue.
- A rollback failure may occur in some limited cases where the node needs a long time to complete one of the resulting configuration changes. If a rollback (for example, **rollback revert 5**) fails during execution, it should be attempted again. The second attempt typically completes the remaining configuration changes required to fully revert to the desired checkpoint.
- When a new backup CPM is commissioned, the user executes the **admin redundancy rollback-sync** command to copy the entire set of rollback files from the active CPM CF to the new standby CPM CF. If the operator wants the system to automatically copy new rollback checkpoints to both CFs whenever a new checkpoint is created, the **configure redundancy rollback-sync** command should be enabled.
- An HA CPM switchover during a rollback revert causes the rollback operation to abort. The newly active CPM has an indeterminate configuration. A log event is created in this case to warn the operator.

When an HA switchover occurs during a rollback or within a few seconds of a rollback completing, the operator is advised to repeat the rollback revert operation to the same checkpoint.

- A rollback checkpoint stores the rollback location and the **local-max-checkpoint** and **remote-max-checkpoint** values, and it is possible that a rollback revert operation may change those values. If an operator changes the **local-max-checkpoint** or **remote-max-checkpoint** values, it is recommended that all the existing checkpoints be deleted to prevent a subsequent rollback revert from changing the maximum values to any of the previous values.
- If a warning prompt (**y/n**) is displayed when a rollback revert is initiated, it is highly recommended that the operator respond **no** to the warning prompt the first time, save a rollback checkpoint before attempting this rollback revert, execute the revert again, and respond **yes**. If the rollback encounters problems, a revert to the saved checkpoint can be used to return to the initial configuration state.

2.10 Transactional configuration

Transactional configuration allows an operator to edit a candidate configuration (a set of configuration changes) without causing operational changes in the router (the active or operational configuration). When the candidate configuration is complete, the operator can explicitly commit the changes to make the entire new configuration become active.

Transactional configuration gives the operator better control and visibility over their router configurations and reduces operational risk while increasing flexibility.

The transactional configuration and configuration rollback functions combine to provide the operational model depicted in the following figure.

Figure 6: Router configuration with rollback and transactions

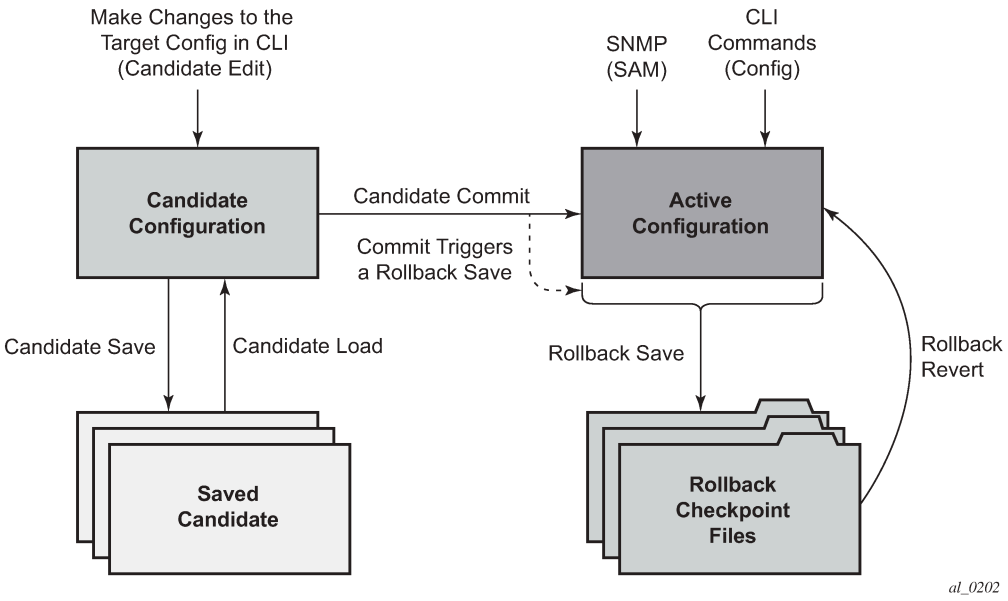

2.10.1 Basic operation

In order to edit the candidate configuration, the operator must first enter the candidate edit mode (edit-cfg) with the **candidate**>**edit** command. The operator can enter and quit the configuration mode as many times as they wish before finally committing the candidate configuration.

In edit-cfg mode, the operator builds a set of candidate configuration changes using the same CLI tree as the standard line-by-line, non-transactional configuration. Tab completion and keyword syntax checking is available.

Just as there is a single operational active configuration that can be modified simultaneously by multiple users in the SR OS, there is also a single global candidate configuration instance. All users make changes in the same global candidate configuration, and a **commit** operation by any user commits the changes made by all users.

Users can exclusively create a candidate configuration by blocking other users and sessions of the same user from entering edit-cfg mode by specifying the **exclusive** parameter. The **config**>**system**>**management-interface**>**cli**>**classic-cli**>**allow-immediate** command can be used to enforce the use of candidate configuration, instead of allowing immediate line-by-line configuration changes.

If a **commit** operation is successful, all of the candidate changes take operational effect and the candidate is cleared. If there is an error in the processing of the commit, or a **commit confirmed** is not confirmed and an auto-revert occurs, the router returns to a configuration state with none of the candidate changes applied. The operator can then continue editing the candidate and try a commit later.

All commands in the candidate configuration must be in the correct order for a commit to be successful. Configuration that depends on other candidate objects must be placed after those objects in the candidate. A set of candidate editing commands (**copy**, **insert**, and so on) are available to correct and reorder the candidate configuration.

The edit-cfg mode is primarily intended for building a candidate configuration while navigating the **configure** branch of the CLI. Many CLI commands in branches other than **configure** are supported while in edit-cfg mode, but access to some CLI branches and commands are blocked, including:

- **exec** command
- **enable-admin** command
- **enable-dynamic-services-config** command
- **admin** branch
- **bof** branch
- **debug** branch
- **tools** branch

The candidate configuration can be saved to a file and subsequently loaded into a candidate configuration. A saved candidate is similar to, but not the same as, an SR OS configuration file generated with an **admin save** command. The saved candidate cannot be used in general as a configuration file and may not **exec** without failures.

There is no SNMP access to the candidate configuration and no SNMP management of candidates, although any configuration changes made using transactional configuration are reported via the standard SR OS SNMP change traps and basic candidate status information is available via SNMP.

A commit may fail for a number of reasons, including:

- **misordering**

  The candidate configuration has changes that are not in the correct order, that is, an object is referred to before it is actually created

- **invalid options and combinations**

  Although many syntax errors are eliminated during the candidate editing process, the candidate configuration may contain combinations of configurations and options that are not valid and are rejected when the SR OS attempts to have them take operational effect

- **out of resources**

  The application of the candidate may exhaust system resources, such as queue resources

Error messages are provided for commit failures to help the operator take the necessary actions to correct the candidate.

Standard line-by-line, non-transactional CLI and SNMP commands are not blocked during the creation or editing of a candidate or the processing of a commit. These commands take immediate operational effect when **Enter** is pressed.

2.10.2 Transactions and rollback

By default, the SR OS automatically creates a new rollback checkpoint after a **commit** operation. The rollback checkpoint includes the new configuration changes made by the commit. An optional **no-checkpoint** keyword can be used to prevent the auto-creation of a rollback checkpoint after a commit. If the commit fails, no new rollback checkpoint is created.

When the **commit confirmed** option is used, a rollback checkpoint is created after the processing of the commit and exists whether the commit is automatically reverted or not.

Transactional configuration relies on the rollback mechanism to operate. Any commands and configurations that are not supported in a rollback revert are also not supported in edit-cfg mode; for example, changes to **chassis-mode**.

2.10.3 Authorization

Authorization works transparently in edit-cfg mode and no unique or new local profile or TACACS+ permissions rules are required other than allowing access to the **candidate** branch. For example, if an operator has permission to access the **configure filter** context, they automatically also have access to the **configure filter** context when in edit-cfg mode.

If the operator’s profile allows access to the candidate **load** and **save** commands, the operations load and save only those items that the user is authorized to access.

The candidate view only displays the items that the user is authorized to access.

The candidate editing commands (such as adding lines, removing lines, and delete operations) only allow operations on items that the user is authorized to access.

The candidate **commit** and **discard** commands, along with **admin rollback revert**, operate on the entire candidate and impact all items; authorization does not apply.

3 Command Trees

3.1 admin Commands

```
- admin
    - certificate
        - clear-ocsp-cache
        - cmpv2
            - cert-request
            - clear-request
            - initial-registration
            - key-update
            - poll
            - show-request
        - convert-file
        - crl-update
        - display
        - est
            - cacert
            - enroll
            - renew
        - export
        - gen-keypair
        - gen-local-cert-req
        - import
        - reload
        - secure-nd-export
        - secure-nd-import
        - update-cert
    - clear
        - lockout
        - password-history
    - compare
    - debug-save
    - disconnect
    - display-config
    - enable-tech
    - ipsec
        - display-key
        - transport-mode
            - display-key
    - reboot
    - redundancy
        - force-switchover
        - rollback-sync
        - synchronize
    - reset-policy-exclusive
    - rollback
        - compare
        - delete
        - revert
        - save
        - view
    - save
    - set-time
    - system
        - license
            - activate
            - clear
            - validate
        - security
            - hash-control
                - custom-hash
```

admin system security os-security

```
            - os-security
                - anti-theft
                    - activate
                    - deactivate
                    - unlock
                - remove-password
                - set-password
            - secure-boot
                - activate
                - revoke-key
                - update-key
                - validate
            - system-password
        - telemetry
            - grpc
                - subscription
    - tech-support
    - view
```

3.2 bof Commands

```
- bof
    - address
    - auto-boot
    - autoconfigure
        - ipv4
            - dhcp
        - ipv6
            - dhcp
    - autonegotiate
    - console-speed
    - dns-domain
    - duplex
    - encrypt
    - encryption-key
    - ip-mtu
    - license-file
    - password
    - persist
    - primary-config
    - primary-dns
    - primary-image
    - save
    - secondary-config
    - secondary-dns
    - secondary-image
    - speed
    - static-route
    - system-base-mac
    - system-profile
    - tertiary-config
    - tertiary-dns
    - tertiary-image
    - wait
```

3.3 candidate Commands

```
- candidate
    - commit
    - confirm
    - copy
    - delete
    - discard
    - edit
    - goto
    - insert
    - load
    - quit
    - redo
    - replace
    - save
    - undo
    - view
```

3.4 configure Commands

```
– configure
```

3.4.1 configure aaa Commands

```
- aaa
    - radius-coa-port
    - radius-server-policy
        - acct-on-off
        - description
        - servers
            - access-algorithm
            - buffering
                - acct-interim
                - acct-start
                - acct-stop
            - disable-stickiness
            - health-check
                - down-timeout
            - hold-down-time
            - ipv6-source-address
            - retry
            - router
            - server
            - source-address
            - timeout
```

3.4.2 configure card Commands

```
- card
    - card-type
    - fail-on-error
    - fp
        - ingress
            - access
                - queue-group
                    - accounting-policy
                    - collect-stats
                    - description
                    - policer-control-override
                        - max-rate
                        - priority-mbs-thresholds
                            - min-thresh-separation
                            - priority
                                - mbs-contribution
                    - policer-control-policy
                    - policer-override
                        - policer
                            - cbs
                            - mbs
                            - packet-byte-offset
                            - rate
                            - stat-mode
            - dist-cpu-protection
                - dynamic-enforcement-policer-pool
            - network
                - pool
                - queue-group
                    - accounting-policy
                    - collect-stats
                    - description
                    - policer-control-override
                        - max-rate
                        - priority-mbs-thresholds
                            - min-thresh-separation
                            - priority
                                - mbs-contribution
                    - policer-control-policy
                    - policer-override
                        - policer
                            - cbs
                            - mbs
                            - packet-byte-offset
                            - rate
                            - stat-mode
                - queue-policy
        - init-extract-prio-mode
    - mda
        - access
            - egress
            - ingress
        - event
            - action
        - fail-on-error
        - mda-type
        - network
            - egress
            - ingress
        - shutdown
```

config card mda sync-e

```
                - sync-e
            - shutdown
```

3.4.3 configure connection-profile-vlan Commands

```
        – connection-profile-vlan
            – description
            – vlan-range
```

3.4.4 configure filter Commands

```
- filter
    - copy
        - ip-filter
        - ipv6-filter
    - dhcp-filter
        - default-action
        - description
        - entry
            - action
            - option
    - dhcp6-filter
        - default-action
        - description
        - entry
            - action
            - option
    - ip-exception
        - description
        - entry
            - description
            - match
                - dst-ip
                - dst-port
                - icmp-code
                - icmp-type
                - src-ip
                - src-port
        - renum
        - scope
    - ip-filter
        - chain-to-system-filter
        - default-action
        - description
        - embed-filter
        - entry
            - action
                - drop
                - drop-extracted-traffic
                - forward
                - ignore-match
                - nat
                - rate-limit
                - reassemble
                - tcp-mss-adjust
            - description
            - log
            - match
                - dscp
                - dst-ip
                - dst-port
                - fragment
                - icmp-code
                - icmp-type
                - ip
                - ip-option
                - multiple-option
                - option-present
                - port
                - src-ip
                - src-port
```

config filter ip-filter entry match src-route-option

```
                    - src-route-option
                    - tcp-ack
                    - tcp-cwr
                    - tcp-ece
                    - tcp-established
                    - tcp-fin
                    - tcp-ns
                    - tcp-psh
                    - tcp-rst
                    - tcp-syn
                    - tcp-urg
                - pbr-down-action-override
                - sticky-dest
            - renum
            - scope
        - ipv6-exception
            - description
            - entry
                - description
                - match
                    - dst-ip
                    - dst-port
                    - icmp-code
                    - icmp-type
                    - port
                    - src-ip
                    - src-port
            - renum
        - ipv6-filter
            - chain-to-system-filter
            - default-action
            - description
            - embed-filter
            - entry
                - action
                    - drop
                    - drop-extracted-traffic
                    - forward
                    - ignore-match
                    - rate-limit
                    - tcp-mss-adjust
                - description
                - log
                - match
                    - ah-ext-hdr
                    - dscp
                    - dst-ip
                    - dst-port
                    - esp-ext-hdr
                    - flow-label
                    - fragment
                    - hop-by-hop-opt
                    - icmp-code
                    - icmp-type
                    - ip
                    - port
                    - routing-type0
                    - src-ip
                    - src-port
                    - tcp-ack
                    - tcp-cwr
                    - tcp-ece
                    - tcp-established
                    - tcp-fin
```

config filter ipv6-filter entry match tcp-ns

```
                        - tcp-ns
                        - tcp-psh
                        - tcp-rst
                        - tcp-syn
                        - tcp-urg
                    - pbr-down-action-override
                    - sticky-dest
                - renum
                - scope
            - log
                - description
                - destination
                - shutdown
                - summary
                    - shutdown
                    - summary-crit
                - wrap-around
            - match-list
                - ip-prefix-list
                    - apply-path
                        - bgp-peers
                    - description
                    - prefix
                    - prefix-exclude
                - ipv6-prefix-list
                    - apply-path
                        - bgp-peers
                    - description
                    - prefix
                    - prefix-exclude
                - port-list
                    - description
                    - port
                - protocol-list
                    - description
                    - protocol
            - md-auto-id
                - filter-id-range
            - policer
                - description
                - mbs
                - pir
                - scope
            - redirect-policy
                - description
                - destination
                    - description
                    - ping-test
                        - drop-count
                        - interval
                        - source-address
                        - timeout
                    - priority
                    - shutdown
                    - unicast-rt-test
                - notify-dest-change
                - router
                - shutdown
                - sticky-dest
            - redirect-policy-binding
                - binding-operator
                - redirect-policy
            - system-filter
                - ip
```

config filter system-filter ipv6

```
                - ipv6
```

3.4.5 configure group-encryption Commands

```
        - group-encryption
            - encryption-keygroup
                - active-outbound-sa
                - description
                - esp-auth-algorithm
                - esp-encryption-algorithm
                - keygroup-name
                - security-association
            - group-encryption-label
```

3.4.6 configure ipsec Commands

```
        - ipsec
            - cert-profile
                - entry
                    - cert
                    - compare-chain-include
                    - key
                    - rsa-signature
                    - send-chain
                        - ca-profile
                - shutdown
            - client-db
                - client
                    - client-identification
                        - idi
                        - peer-ip-prefix
                    - client-name
                    - credential
                        - pre-shared-key
                    - private-interface
                    - private-service
                    - shutdown
                    - ts-negotiation
                    - tunnel-template
                - description
                - match-list
                    - idi
                    - peer-ip-prefix
                - shutdown
            - ike-policy
                - auth-method
                - auto-eap-method
                - auto-eap-own-method
                - description
                - dpd
                - ike-mode
                - ike-transform
                - ike-version
                - ikev1-ph1-responder-delete-notify
                - ikev2-fragment
                - ipsec-lifetime
                - limit-init-exchange
                - lockout
                - match-peer-id-to-cert
                - nat-traversal
                - own-auth-method
                - pfs
                - ppk-required
                - relay-unsolicited-cfg-attribute
                    - internal-ip4-address
                    - internal-ip4-dns
                    - internal-ip4-netmask
                    - internal-ip6-address
                    - internal-ip6-dns
                - send-idr-after-eap-success
            - ike-transform
                - dh-group
                - ike-auth-algorithm
                - ike-encryption-algorithm
                - ike-prf-algorithm
                - isakmp-lifetime
```

config ipsec ipsec-transform

```
            - ipsec-transform
                - esp-auth-algorithm
                - esp-encryption-algorithm
                - extended-sequence-number
                - ipsec-lifetime
                - pfs-dh-group
            - ipsec-transport-mode-profile
                - description
                - dynamic-keying
                    - auto-establish
                    - cert
                        - cert-profile
                        - status-verify
                            - default-result
                            - primary
                        - trust-anchor-profile
                    - ike-policy
                    - local-id
                    - ppk
                    - pre-shared-key
                    - transform
                - max-history-esp-key-records
                - max-history-ike-key-records
                - replay-window
            - ppk-list
                - ppk-id
            - radius-accounting-policy
                - include-radius-attribute
                    - acct-stats
                    - called-station-id
                    - calling-station-id
                    - framed-ip-addr
                    - framed-ipv6-prefix
                    - nas-identifier
                    - nas-ip-addr
                    - nas-port-id
                - radius-server-policy
                - update-interval
            - radius-authentication-policy
                - include-radius-attribute
                    - called-station-id
                    - calling-station-id
                    - client-cert-subject-key-id
                    - nas-identifier
                    - nas-ip-addr
                    - nas-port-id
                - password
                - radius-server-policy
            - show-ipsec-keys
            - static-sa
                - authentication
                - description
                - direction
                - protocol
                - spi
            - trust-anchor-profile
                - trust-anchor
            - ts-list
                - local
                    - entry
                        - address
                        - protocol
                - remote
                    - entry
```

config ipsec ts-list remote entry address

```
                    - address
                    - protocol
            - tunnel-template
                - clear-df-bit
                - copy-traffic-class-upon-decapsulation
                - description
                - encapsulated-ip-mtu
                - icmp-generation
                    - frag-required
                        - interval
                        - message-count
                - icmp6-generation
                    - pkt-too-big
                        - interval
                        - interval
                        - message-count
                - ip-mtu
                - pmtu-discovery-aging
                - ppk-list
                - private-tcp-mss-adjust
                - propagate-pmtu-v4
                - propagate-pmtu-v6
                - public-tcp-mss-adjust
                - replay-window
                - reverse-route
                    - metric
                    - preference
                - sp-reverse-route
                - transform
```

3.4.7 configure isa Commands

```
- isa
    - nat-group
        - active-mda-limit
        - description
        - mda
        - shutdown
        - suppress-lsn-events
        - suppress-lsn-sub-blks-free
    - tunnel-group
        - active-mda-number
        - backup
        - description
        - ipsec-responder-only
        - mda
        - member-pool
        - multi-active
        - primary
        - reassembly
        - shutdown
        - stats-collection
            - isa-dp-cpu-usage
    - tunnel-member-pool
        - description
        - mda
```

3.4.8 configure lag Commands

```
        - lag
            - access
                - adapt-qos
                - bandwidth
                - booking-factor
            - description
            - dynamic-cost
            - encap-type
            - hash-weight-threshold
            - hold-time
            - lacp
            - lacp-fallback
            - lacp-mux-control
            - lacp-xmit-interval
            - lacp-xmit-stdby
            - lldp-member-template
                - dest-mac
                    - admin-status
                    - notification
                    - port-id-subtype
                    - tunnel-nearest-bridge
                    - tunnel-nearest-customer
                    - tunnel-nearest-non-tpmr
                    - tx-mgmt-address
                    - tx-tlvs
            - mac
            - mode
            - monitor-oper-group
            - port
            - port-threshold
            - selection-criteria
            - shutdown
            - standby-signaling
```

3.4.9 configure log Commands

```
        - log
            - accounting-policy
                - align
                - collection-interval
                - custom-record
                    - policer
                        - e-counters
                            - exceed-profile-octets-discarded-count
                            - exceed-profile-octets-forwarded-count
                            - exceed-profile-octets-offered-count
                            - exceed-profile-packets-discarded-count
                            - exceed-profile-packets-forwarded-count
                            - exceed-profile-packets-offered-count
                            - in-plus-profile-octets-discarded-count
                            - in-plus-profile-octets-forwarded-count
                            - in-plus-profile-octets-offered-count
                            - in-plus-profile-packets-discarded-count
                            - in-plus-profile-packets-forwarded-count
                            - in-plus-profile-packets-offered-count
                            - in-profile-octets-discarded-count
                            - in-profile-octets-forwarded-count
                            - in-profile-octets-offered-count
                            - in-profile-packets-discarded-count
                            - in-profile-packets-forwarded-count
                            - in-profile-packets-offered-count
                            - out-profile-octets-discarded-count
                            - out-profile-octets-forwarded-count
                            - out-profile-octets-offered-count
                            - out-profile-packets-discarded-count
                            - out-profile-packets-forwarded-count
                            - out-profile-packets-offered-count
                            - uncoloured-octets-offered-count
                            - uncoloured-packets-offered-count
                        - i-counters
                            - in-profile-octets-discarded-count
                            - in-profile-octets-forwarded-count
                            - in-profile-octets-offered-count
                            - in-profile-packets-discarded-count
                            - in-profile-packets-forwarded-count
                            - in-profile-packets-offered-count
                            - out-profile-octets-discarded-count
                            - out-profile-octets-forwarded-count
                            - out-profile-octets-offered-count
                            - out-profile-packets-discarded-count
                            - out-profile-packets-forwarded-count
                            - out-profile-packets-offered-count
                            - uncoloured-octets-offered-count
                            - uncoloured-packets-offered-count
                    - queue
                        - e-counters
                            - all
                            - in-profile-octets-discarded-count
                            - in-profile-octets-forwarded-count
                            - in-profile-packets-discarded-count
                            - in-profile-packets-forwarded-count
                            - out-profile-octets-discarded-count
                            - out-profile-octets-forwarded-count
                            - out-profile-packets-discarded-count
                            - out-profile-packets-forwarded-count
                        - i-counters
```

configure log accounting-policy custom-record queue i-counters all

```
                        - all
                        - all-octets-offered-count
                        - all-packets-offered-count
                        - high-octets-discarded-count
                        - high-octets-offered-count
                        - high-packets-discarded-count
                        - high-packets-offered-count
                        - in-profile-octets-forwarded-count
                        - in-profile-packets-forwarded-count
                        - low-octets-discarded-count
                        - low-octets-offered-count
                        - low-packets-discarded-count
                        - low-packets-offered-count
                        - out-profile-octets-forwarded-count
                        - out-profile-packets-forwarded-count
                        - uncoloured-octets-offered-count
                        - uncoloured-packets-offered-count
                    - ref-policer
                        - e-counters
                            - exceed-profile-octets-discarded-count
                            - exceed-profile-octets-forwarded-count
                            - exceed-profile-octets-offered-count
                            - exceed-profile-packets-discarded-count
                            - exceed-profile-packets-forwarded-count
                            - exceed-profile-packets-offered-count
                            - in-plus-profile-octets-discarded-count
                            - in-plus-profile-octets-forwarded-count
                            - in-plus-profile-octets-offered-count
                            - in-plus-profile-packets-discarded-count
                            - in-plus-profile-packets-forwarded-count
                            - in-plus-profile-packets-offered-count
                            - in-profile-octets-discarded-count
                            - in-profile-octets-forwarded-count
                            - in-profile-octets-offered-count
                            - in-profile-packets-discarded-count
                            - in-profile-packets-forwarded-count
                            - in-profile-packets-offered-count
                            - out-profile-octets-discarded-count
                            - out-profile-octets-forwarded-count
                            - out-profile-octets-offered-count
                            - out-profile-packets-discarded-count
                            - out-profile-packets-forwarded-count
                            - out-profile-packets-offered-count
                            - uncoloured-octets-offered-count
                            - uncoloured-packets-offered-count
                        - i-counters
                            - in-profile-octets-discarded-count
                            - in-profile-octets-forwarded-count
                            - in-profile-octets-offered-count
                            - in-profile-packets-discarded-count
                            - in-profile-packets-forwarded-count
                            - in-profile-packets-offered-count
                            - out-profile-octets-discarded-count
                            - out-profile-octets-forwarded-count
                            - out-profile-octets-offered-count
                            - out-profile-packets-discarded-count
                            - out-profile-packets-forwarded-count
                            - out-profile-packets-offered-count
                            - uncoloured-octets-offered-count
                            - uncoloured-packets-offered-count
                    - ref-queue
                        - e-counters
                            - all
                            - in-profile-octets-discarded-count
```

config log acct-policy cr ref-queue e-counters in-profile-octets-forwarded-count

```
                        - in-profile-octets-forwarded-count
                        - in-profile-packets-discarded-count
                        - in-profile-packets-forwarded-count
                        - out-profile-octets-discarded-count
                        - out-profile-octets-forwarded-count
                        - out-profile-packets-discarded-count
                        - out-profile-packets-forwarded-count
                    - i-counters
                        - all
                        - all-octets-offered-count
                        - all-packets-offered-count
                        - high-octets-discarded-count
                        - high-octets-offered-count
                        - high-packets-discarded-count
                        - high-packets-offered-count
                        - in-profile-octets-forwarded-count
                        - in-profile-packets-forwarded-count
                        - low-octets-discarded-count
                        - low-octets-offered-count
                        - low-packets-discarded-count
                        - low-packets-offered-count
                        - out-profile-octets-forwarded-count
                        - out-profile-packets-forwarded-count
                        - uncoloured-octets-offered-count
                        - uncoloured-packets-offered-count
                - significant-change
            - default
            - description
            - include-system-info
            - record
            - shutdown
            - to
        - app-route-notifications
            - cold-start-wait
            - route-recovery-wait
        - encryption-key
        - event-control
        - event-damping
        - event-handling
            - handler
                - action-list
                    - entry
                        - description
                        - min-delay
                        - script-policy
                        - shutdown
                - description
                - shutdown
        - event-trigger
            - event
                - description
                - shutdown
                - trigger-entry
                    - debounce
                    - description
                    - event-handler
                    - log-filter
                    - shutdown
        - file-id
            - description
            - location
            - rollover
        - file-storage-control
            - accounting-files-total-size
```

config log storage log-files-total-size

```
            - log-files-total-size
        - filter
            - default-action
            - description
            - entry
                - action
                - action
                - description
                - match
                    - application
                    - message
                    - number
                    - router
                    - severity
                    - subject
        - log-id
            - description
            - filter
            - from
            - netconf-stream
            - shutdown
            - time-format
            - to
        - route-preference
        - services-all-events
            - service
        - snmp-trap-group
            - description
            - trap-target
        - syslog
            - address
            - description
            - facility
            - hostname
            - level
            - log-prefix
            - port
            - timestamp-format
            - tls-client-profile
        - throttle-rate
```

3.4.10 configure macsec Commands

```
        - macsec
            - connectivity-association
                - cipher-suite
                - clear-tag-mode
                - delay-protection
                - description
                - encryption-offset
                - macsec-encrypt
                - replay-protection
                - replay-window-size
                - shutdown
                - static-cak
                    - active-psk
                    - mka-hello-interval
                    - mka-key-server-priority
                    - pre-shared-key
                        - cak
                        - ckn
            - mac-policy
                - dest-mac-address
```

3.4.11 configure mirror Commands

```
- mirror
    - mirror-dest
        - description
        - endpoint
            - description
            - revert-time
        - fc
        - pcap
            - file-url
            - router
        - remote-source
            - far-end
            - spoke-sdp
                - egress
                    - vc-label
                - ingress
                    - vc-label
                - shutdown
        - sap
            - egress
                - ip-mirror
                    - sa-mac
                - qos
        - shutdown
        - slice-size
        - spoke-sdp
            - egress
                - vc-label
            - precedence
            - shutdown
    - mirror-source
        - ip-filter
        - ipv6-filter
        - port
        - sap
        - shutdown
```

3.4.12 configure oam-pm Commands

```
- oam-pm
    - bin-group
        - bin-type
            - bin
                - lower-bound
            - delay-event
            - delay-event-exclusion
            - exclude-from-avg
        - description
        - shutdown
    - session
        - bin-group
        - description
        - ip
            - allow-egress-remark-dscp
            - dest-udp-port
            - destination
            - do-not-fragment
            - dscp
            - fc
            - forwarding
            - pattern
            - profile
            - router
            - router-instance
            - source
            - source-udp-port
            - ttl
            - tunnel
                - mpls
                    - rsvp-te
                        - lsp
                    - rsvp-te-auto
                        - from
                        - lsp-template
                        - to
                    - sr-isis
                        - flex-algo
                        - igp-instance
                        - prefix
                    - sr-ospf
                        - flex-algo
                        - igp-instance
                        - prefix
                    - sr-policy
                        - color
                        - endpoint
                        - segment-list
                    - sr-te
                        - lsp
            - twamp-light
                - allow-ipv6-udp-checksum-zero
                - interval
                - loss
                    - flr-threshold
                    - hli-force-count
                    - timing
                - loss-events
                    - avg-flr-event
                    - chli-event
```

config oam-pm session ip twamp-light loss-events hli-event

```
                        - hli-event
                        - unavailability-event
                        - undet-availability-event
                        - undet-unavailability-event
                    - pad-size
                    - pad-tlv-size
                    - record-stats
                    - session-sender-type
                    - shutdown
                    - test-duration
                    - timestamp-format
                - meas-interval
                    - accounting-policy
                    - boundary-type
                    - clock-offset
                    - event-mon
                        - delay-events
                        - loss-events
                        - shutdown
                    - intervals-stored
```

3.4.13 configure port Commands

```
- port
    - access
        - egress
            - pool
        - ingress
            - pool
    - connector
        - breakout
        - rs-fec-mode
    - ddm-events
    - description
    - dist-cpu-protection
    - dwdm
        - coherent
            - compatibility
            - cpr-window-size
            - dispersion
            - mode
            - report-alarms
            - rx-los-reaction
            - rx-los-thresh
            - sweep
            - target-power
        - frequency
    - ethernet
        - access
            - bandwidth
            - booking-factor
            - egress
                - queue-group
                    - accounting-policy
                    - collect-stats
                    - description
                    - queue-overrides
                        - queue
                            - adaptation-rule
                            - burst-limit
                            - cbs
                            - drop-tail
                                - low
                                    - percent-reduction-from-mbs
                            - mbs
                            - parent
                            - percent-rate
                            - rate
                    - scheduler-override
                        - scheduler
                            - parent
                            - rate
                    - scheduler-policy
            - ingress
                - queue-group
                    - accounting-policy
                    - collect-stats
                    - description
                    - queue-overrides
                        - queue
                            - adaptation-rule
                            - cbs
                            - drop-tail
```

config port ethernet access ing qgrp qover q mbs

```
                                    - low
                                        - percent-reduction-from-mbs
                                - mbs
                                - rate
                        - scheduler-override
                            - scheduler
                                - parent
                                - rate
                        - scheduler-policy
                - accounting-policy
                - autonegotiate
                - collect-stats
                - crc-monitor
                    - sd-threshold
                    - sf-threshold
                    - window-size
                - dampening
                    - half-life
                    - shutdown
                    - suppress-threshold
                - discard-rx-pause-frames
                - dot1q-etype
                - dot1x
                    - macsec
                        - exclude-mac-policy
                        - exclude-protocol
                        - rx-must-be-encrypted
                        - sub-port
                            - ca-name
                            - eapol-destination-address
                            - encap-match
                            - max-peer
                            - shutdown
                    - max-auth-req
                    - per-host-authentication
                        - allowed-source-macs
                            - mac-address
                        - authenticator-init
                        - shutdown
                    - port-control
                    - quiet-period
                    - radius-plcy
                    - radius-server-policy
                    - re-auth-period
                    - re-authentication
                    - server-timeout
                    - shutdown
                    - supplicant-timeout
                    - transmit-period
                    - tunnel-dot1q
                    - tunnel-qinq
                    - tunneling
                - down-on-internal-error
                - duplex
                - egress
                - egress-rate
                - egress-scheduler-override
                    - level
                    - max-rate
                - egress-scheduler-policy
                - encap-type
                - hold-time
                - ingress-rate
                - lacp-tunnel
```

config port ethernet lldp

```
            - lldp
            - lldp
                - dest-mac
                    - admin-status
                    - notification
                    - port-id-subtype
                    - tunnel-nearest-bridge
                    - tunnel-nearest-customer
                    - tunnel-nearest-non-tpmr
                    - tx-mgmt-address
                    - tx-tlvs
            - load-balancing-algorithm
            - mac
            - min-frame-length
            - mode
            - mtu
            - network
                - accounting-policy
                - collect-stats
                - egress
                    - queue-group
                        - accounting-policy
                        - collect-stats
                        - description
                        - policer-control-policy
                        - queue-overrides
                            - queue
                                - adaptation-rule
                                - cbs
                                - drop-tail
                                    - low
                                        - percent-reduction-from-mbs
                                - mbs
                                - percent-rate
                                - rate
                        - scheduler-policy
                - queue-policy
            - qinq-etype
            - report-alarm
            - rs-fec-mode
            - speed
            - ssm
                - code-type
                - shutdown
                - tx-dus
            - util-stats-interval
            - xgig
        - monitor-oper-group
        - network
            - egress
                - pool
        - shutdown
        - transceiver
            - digital-coherent-optics
            - optical-line-system
                - egress-amplifier-gain
```

3.4.14 configure port-xc Commands

```
        - port-xc
            - pxc
                - description
                - port
                - shutdown
```

3.4.15 configure qos Commands

```
- qos
    - match-list
        - ip-prefix-list
            - description
            - prefix
        - ipv6-prefix-list
            - description
            - prefix
    - md-auto-id
        - qos-policy-id-range
    - network
        - description
        - egress
            - dscp
            - fc
                - de-mark
                - dot1p
                - dot1p-in-profile
                - dot1p-out-profile
                - dscp-in-profile
                - dscp-out-profile
                - lsp-exp-in-profile
                - lsp-exp-out-profile
                - port-redirect-group
            - ip-criteria
                - entry
                    - action
                    - description
                    - match
                        - dscp
                        - dst-ip
                        - dst-port
                        - fragment
                        - icmp-type
                        - src-ip
                        - src-port
                - renum
            - ipv6-criteria
                - entry
                    - action
                    - description
                    - match
                        - dscp
                        - dst-ip
                        - dst-port
                        - fragment
                        - icmp-type
                        - src-ip
                        - src-port
                - renum
            - prec
            - remarking
        - ingress
            - default-action
            - dot1p
            - dscp
            - fc
                - fp-redirect-group
            - ip-criteria
                - entry
```

config qos network ingress ip-criteria entry action

```
                        - action
                        - description
                        - match
                            - dscp
                            - dst-ip
                            - dst-port
                            - fragment
                            - src-ip
                            - src-port
                    - renum
                - ipv6-criteria
                    - entry
                        - action
                        - description
                        - match
                            - dscp
                            - dst-ip
                            - dst-port
                            - fragment
                            - src-ip
                            - src-port
                    - renum
                - ler-use-dscp
                - lsp-exp
            - scope
        - network-queue
            - description
            - fc
                - multicast-queue
                - queue
            - queue
                - avg-frame-overhead
                - cbs
                - drop-tail
                    - low
                        - percent-reduction-from-mbs
                - mbs
                - port-parent
                - rate
        - policer-control-policy
            - description
            - root
                - max-percent-rate
                - max-rate
                - priority-mbs-thresholds
                    - min-thresh-separation
                    - priority
                        - mbs-contribution
                - profile-preferred
            - tier
                - arbiter
                    - description
                    - parent
                    - percent-rate
                    - rate
        - port-scheduler-policy
            - description
            - level
            - max-rate
            - orphan-override
        - queue-group-templates
            - egress
                - queue-group
                    - description
```

cfg qos qgrps egr qgrp fc

```
                    - fc
                        - queue
                    - policer
                        - cbs
                        - description
                        - high-prio-only
                        - mbs
                        - packet-byte-offset
                        - parent
                        - percent-rate
                        - profile-capped
                        - rate
                        - stat-mode
                    - queue
                        - burst-limit
                        - cbs
                        - drop-tail
                            - exceed
                                - percent-reduction-from-mbs
                            - high
                                - percent-reduction-from-mbs
                            - highplus
                                - percent-reduction-from-mbs
                            - low
                                - percent-reduction-from-mbs
                        - dynamic-mbs
                        - mbs
                        - packet-byte-offset
                        - parent
                        - percent-rate
                        - port-parent
                        - queue-delay
                        - rate
                        - wred-queue
                - ingress
                    - queue-group
                        - description
                        - policer
                            - adaptation-rule
                            - cbs
                            - description
                            - high-prio-only
                            - mbs
                            - packet-byte-offset
                            - parent
                            - percent-rate
                            - profile-capped
                            - rate
                            - stat-mode
                        - queue
                            - burst-limit
                            - cbs
                            - drop-tail
                                - low
                                    - percent-reduction-from-mbs
                            - mbs
                            - packet-byte-offset
                            - parent
                            - rate
            - sap-egress
                - description
                - dot1p
                - dscp
                - ethernet-ctag
```

config qos sap-egress fc

```
                - fc
                    - de-mark
                    - de-mark-inner
                    - de-mark-outer
                    - dot1p
                    - dot1p-inner
                    - dot1p-outer
                    - dscp
                    - policer
                    - prec
                    - queue
                - ip-criteria
                    - entry
                        - action
                        - description
                        - match
                            - dscp
                            - dst-ip
                            - dst-port
                            - fragment
                            - src-ip
                            - src-port
                    - renum
                - ipv6-criteria
                    - entry
                        - action
                        - description
                        - match
                            - dscp
                            - dst-ip
                            - dst-port
                            - src-ip
                            - src-port
                    - renum
                - policer
                    - cbs
                    - description
                    - high-prio-only
                    - mbs
                    - packet-byte-offset
                    - parent
                    - percent-rate
                    - profile-capped
                    - profile-out-preserve
                    - rate
                    - stat-mode
                - prec
                - queue
                    - avg-frame-overhead
                    - burst-limit
                    - cbs
                    - drop-tail
                        - exceed
                            - percent-reduction-from-mbs
                        - high
                            - percent-reduction-from-mbs
                        - highplus
                            - percent-reduction-from-mbs
                        - low
                            - percent-reduction-from-mbs
                    - mbs
                    - packet-byte-offset
                    - parent
                    - percent-rate
```

config qos sap-egress queue port-parent

```
                - port-parent
                - rate
                - wred-queue
            - scope
        - sap-ingress
            - default-fc
            - default-priority
            - description
            - dot1p
            - dscp
            - fc
                - broadcast-policer
                - broadcast-queue
                - de-1-out-profile
                - egress-fc
                - in-remark
                - multicast-policer
                - multicast-queue
                - out-remark
                - policer
                - profile
                - queue
                - unknown-policer
                - unknown-queue
            - ip-criteria
                - entry
                    - action
                    - description
                    - match
                        - dscp
                        - dst-ip
                        - dst-port
                        - fragment
                        - src-ip
                        - src-port
                - renum
            - ipv6-criteria
                - entry
                    - action
                    - description
                    - match
                        - dscp
                        - dst-ip
                        - dst-port
                        - fragment
                        - src-ip
                        - src-port
                - renum
            - mac-criteria
                - entry
                    - action
                    - description
                    - match
                        - dot1p
                        - dsap
                        - dst-mac
                        - etype
                        - inner-tag
                        - outer-tag
                        - snap-oui
                        - snap-pid
                        - src-mac
                        - ssap
                - renum
```

config qos sap-ingress mac-criteria type

```
                - type
            - policer
                - cbs
                - description
                - high-prio-only
                - mbs
                - packet-byte-offset
                - parent
                - percent-rate
                - profile-capped
                - rate
                - stat-mode
            - prec
            - queue
                - burst-limit
                - cbs
                - drop-tail
                    - low
                        - percent-reduction-from-mbs
                - mbs
                - packet-byte-offset
                - parent
                - percent-rate
                - rate
            - scope
        - scheduler-policy
            - description
            - frame-based-accounting
            - tier
                - parent-location
                - scheduler
                    - description
                    - parent
                    - percent-rate
                    - port-parent
                    - rate
```

3.4.16 configure redundancy Commands

```
- redundancy
    - bgp-multi-homing
        - boot-timer
        - site-activation-timer
        - site-activation-timer
        - site-min-down-timer
    - cert-sync
    - multi-chassis
        - ipsec-domain
            - designated-role
            - priority
            - revertive
            - shutdown
            - tunnel-group
        - peer
            - authentication-key
            - description
            - mc-endpoint
                - bfd-enable
                - boot-timer
                - hold-on-neighbor-failure
                - keep-alive-interval
                - passive-mode
                - shutdown
                - system-priority
            - mc-ipsec
                - bfd-enable
                - discovery-interval
                - domain
                    - shutdown
                - hold-on-neighbor-failure
                - keep-alive-interval
                - tunnel-group
                - tunnel-group
                    - peer-group
                    - priority
                    - shutdown
            - mc-lag
                - hold-on-neighbor-failure
                - keep-alive-interval
                - lag
                - shutdown
            - mc-ring
                - l3-ring
                    - in-band-control-path
                        - debounce
                        - dst-ip
                        - interface
                        - max-debounce-time
                        - service-id
                        - service-name
                    - ring-node
                        - connectivity-verify
                            - dst-ip
                            - interval
                            - service-id
                            - service-name
                            - shutdown
                            - src-ip
                            - src-mac
```

config redundancy mc peer mcr l3ring node cv vlan

```
                                - vlan
                        - shutdown
                - peer-name
                - shutdown
                - source-address
                - sync
                    - ipsec
                    - l2tp
                    - local-dhcp-server
                    - mc-ring
                    - port
                        - range
                    - shutdown
                    - transport-encryption
                        - application
                    - tunnel-group
        - rollback-sync
        - synchronize
```

3.4.17 configure router Commands

```
        – router
```

3.4.17.1 configure router admin-tags Commands

```
            - admin-tags
                - admin-tag
                - route-admin-tag-policy
                    - exclude
                    - include
```

3.4.17.2 configure router aggregate Commands

```
            - aggregate
```

3.4.17.3 configure router allow-bgp-to-igp-export Commands

```
            - allow-bgp-to-igp-export
```

3.4.17.4 configure router allow-icmp-redirect Commands

```
            - allow-icmp-redirect
```

3.4.17.5 configure router allow-icmp6-redirect Commands

```
            - allow-icmp6-redirect
```

3.4.17.6 configure router autonomous-system Commands

```
            - autonomous-system
```

3.4.17.7 configure router bfd Commands

```
            – bfd
                – abort
                – begin
                – bfd-template
                    – echo-receive
                    – multiplier
                    – receive-interval
                    – transmit-interval
                    – type
                – commit
```

3.4.17.8 configure router bgp Commands

```
            - bgp
                - add-paths
                    - evpn
                    - ipv4
                    - ipv6
                    - label-ipv4
                    - label-ipv6
                    - vpn-ipv4
                    - vpn-ipv6
                - advertise-external
                - advertise-inactive
                - advertise-ipv6-next-hops
                - aggregator-id-zero
                - auth-keychain
                - authentication-key
                - best-path-selection
                    - always-compare-med
                    - as-path-ignore
                    - compare-origin-validation-state
                    - d-path-length-ignore
                    - deterministic-med
                    - ebgp-ibgp-equal
                    - ignore-nh-metric
                    - ignore-router-id
                    - origin-invalid-unusable
                - bfd-enable
                - bfd-strict-mode
                    - advertise
                    - next-hop-reachability
                - bgp-tunnel-metric
                - bgp-tunnel-preference
                - block-prefix-sid
                - cluster
                - connect-retry
                - convergence
                    - family
                        - max-wait-to-advertise
                    - min-wait-to-advertise
                - damp-peer-oscillations
                - damping
                - def-recv-evpn-encap
                - default-label-preference
                - default-preference
                - description
                - disable-4byte-asn
                - disable-client-reflect
                - disable-communities
                - disable-fast-external-failover
                - disable-route-table-install
                - dynamic-neighbor-limit
                - ebgp-default-reject-policy
                - egress-peer-engineering
                    - shutdown
                - enable-inter-as-vpn
                - enable-peer-tracking
                - enable-rr-vpn-forwarding
                - enable-subconfed-vpn-forwarding
                - enforce-first-as
                - error-handling
                    - legacy-mode
```

config router bgp error-handling update-fault-tolerance

```
            - update-fault-tolerance
        - export
        - extended-nh-encoding
        - family
        - graceful-restart
            - enable-notification
            - long-lived
                - advertise-stale-to-all-neighbors
                - advertised-stale-time
                - family
                    - advertised-stale-time
                    - helper-override-stale-time
                - forwarding-bits-set
                - helper-override-restart-time
                - helper-override-stale-time
            - restart-time
            - stale-routes-time
        - group
            - add-paths
                - evpn
                - ipv4
                - ipv6
                - label-ipv4
                - label-ipv6
                - vpn-ipv4
                - vpn-ipv6
            - advertise-inactive
            - advertise-ipv6-next-hops
            - aggregator-id-zero
            - aigp
            - as-override
            - auth-keychain
            - authentication-key
            - bfd-enable
            - bfd-strict-mode
                - advertise
                - next-hop-reachability
            - block-prefix-sid
            - cluster
            - connect-retry
            - damp-peer-oscillations
            - damping
            - def-recv-evpn-encap
            - default-label-preference
            - default-preference
            - default-route-target
            - description
            - disable-4byte-asn
            - disable-capability-negotiation
            - disable-client-reflect
            - disable-communities
            - disable-fast-external-failover
            - dynamic-neighbor
                - interface
                    - allowed-peer-as
                    - max-sessions
                - match
                    - prefix
                        - allowed-peer-as
            - dynamic-neighbor-limit
            - ebgp-default-reject-policy
            - egress-engineering
                - shutdown
            - egress-peer-engineering-label-unicast
```

config router bgp group enable-origin-validation

```
                    - enable-origin-validation
                    - enable-peer-tracking
                    - enforce-first-as
                    - error-handling
                        - update-fault-tolerance
                    - export
                    - extended-nh-encoding
                    - family
                    - graceful-restart
                        - enable-notification
                        - long-lived
                            - advertise-stale-to-all-neighbors
                            - advertised-stale-time
                            - family
                                - advertised-stale-time
                                - helper-override-stale-time
                            - forwarding-bits-set
                            - helper-override-restart-time
                            - helper-override-stale-time
                        - restart-time
                        - stale-routes-time
                    - hold-time
                    - import
                    - initial-send-delay-zero
                    - keepalive
                    - label-preference
                    - link-bandwidth
                        - accept-from-ebgp
                        - add-to-received-ebgp
                        - aggregate-used-paths
                        - send-to-ebgp
                    - local-address
                    - local-as
                    - local-preference
                    - loop-detect
                    - loop-detect-threshold
                    - med-out
                    - min-route-advertisement
                    - multihop
                    - multipath-eligible
                    - neighbor
                        - add-paths
                            - evpn
                            - ipv4
                            - ipv6
                            - label-ipv4
                            - label-ipv6
                            - vpn-ipv4
                            - vpn-ipv6
                        - advertise-inactive
                        - advertise-ipv6-next-hops
                        - advertise-ldp-prefix
                        - aggregator-id-zero
                        - aigp
                        - as-override
                        - auth-keychain
                        - authentication-key
                        - bfd-enable
                        - bfd-strict-mode
                            - advertise
                            - next-hop-reachability
                        - block-prefix-sid
                        - cluster
                        - connect-retry
```

config router bgp group neighbor damp-peer-oscillations

```
                        - damp-peer-oscillations
                        - damping
                        - def-recv-evpn-encap
                        - default-label-preference
                        - default-preference
                        - default-route-target
                        - description
                        - disable-4byte-asn
                        - disable-capability-negotiation
                        - disable-client-reflect
                        - disable-communities
                        - disable-fast-external-failover
                        - ebgp-default-reject-policy
                        - egress-engineering
                            - shutdown
                        - egress-peer-engineering-label-unicast
                        - enable-origin-validation
                        - enable-peer-tracking
                        - enforce-first-as
                        - error-handling
                            - update-fault-tolerance
                        - export
                        - extended-nh-encoding
                        - family
                        - graceful-restart
                            - enable-notification
                            - long-lived
                                - advertise-stale-to-all-neighbors
                                - advertised-stale-time
                                - family
                                    - advertised-stale-time
                                    - helper-override-stale-time
                                - forwarding-bits-set
                                - helper-override-restart-time
                                - helper-override-stale-time
                            - restart-time
                            - stale-routes-time
                        - hold-time
                        - import
                        - initial-send-delay-zero
                        - keepalive
                        - l2vpn-cisco-interop
                        - label-preference
                        - link-bandwidth
                            - accept-from-ebgp
                            - add-to-received-ebgp
                            - aggregate-used-paths
                            - send-to-ebgp
                        - local-address
                        - local-as
                        - local-preference
                        - loop-detect
                        - loop-detect-threshold
                        - med-out
                        - min-route-advertisement
                        - multihop
                        - multipath-eligible
                        - next-hop-self
                        - next-hop-unchanged
                        - outbound-route-filtering
                            - extended-community
                                - accept-orf
                                - send-orf
                        - passive
```

config router bgp group neighbor path-mtu-discovery

```
                        - path-mtu-discovery
                        - peer-as
                        - preference
                        - prefix-limit
                        - remove-private
                        - selective-label-ipv4-install
                        - send-default
                        - shutdown
                        - split-horizon
                        - tcp-mss
                        - third-party-nexthop
                        - ttl-security
                        - type
                        - vpn-apply-export
                        - vpn-apply-import
                    - next-hop-self
                    - next-hop-unchanged
                    - outbound-route-filtering
                        - extended-community
                            - accept-orf
                            - send-orf
                    - passive
                    - path-mtu-discovery
                    - peer-as
                    - preference
                    - prefix-limit
                    - remove-private
                    - selective-label-ipv4-install
                    - send-default
                    - shutdown
                    - split-horizon
                    - tcp-mss
                    - third-party-nexthop
                    - ttl-security
                    - type
                    - vpn-apply-export
                    - vpn-apply-import
                - hold-time
                - ibgp-multipath
                - import
                - initial-send-delay-zero
                - keepalive
                - label-allocation
                    - label-ipv6
                        - disable-explicit-null
                - label-preference
                - link-state-export-enable
                - link-state-import-enable
                - local-as
                - local-preference
                - loop-detect
                - loop-detect-threshold
                - med-out
                - min-route-advertisement
                - mp-bgp-keep
                - multi-path
                    - ipv4
                    - ipv6
                    - label-ipv4
                    - label-ipv6
                    - maximum-paths
                - multihop
                - neighbor-trust
                - next-hop-resolution
```

config router bgp next-hop-res allow-unresolved-leaking

```
                    - allow-unresolved-leaking
                    - labeled-routes
                        - allow-static
                        - rr-use-route-table
                        - transport-tunnel
                            - family
                                - allow-flex-algo-fallback
                                - enforce-strict-tunnel-tagging
                                - enforce-untagged-route
                                - resolution
                                - resolution-filter
                                    - bgp
                                    - ldp
                                    - rsvp
                                    - sr-isis
                                    - sr-ospf
                                    - sr-ospf3
                                    - sr-policy
                                    - sr-te
                                    - udp
                        - use-bgp-routes
                            - label-ipv6-explicit-null
                    - policy
                    - shortcut-tunnel
                        - family
                            - allow-flex-algo-fallback
                            - disallow-igp
                            - enforce-strict-tunnel-tagging
                            - enforce-untagged-route
                            - resolution
                            - resolution-filter
                                - bgp
                                - ldp
                                - rsvp
                                - sr-isis
                                - sr-ospf
                                - sr-ospf3
                                - sr-policy
                                - sr-te
                    - use-bgp-routes
                    - use-leaked-routes
                        - static
                    - vpn-family-policy
                    - weighted-ecmp
                - optimal-route-reflection
                    - location
                        - primary-ip-address
                        - primary-ipv6-address
                        - secondary-ip-address
                        - secondary-ipv6-address
                        - tertiary-ip-address
                        - tertiary-ipv6-address
                    - spf-wait
                - outbound-route-filtering
                    - extended-community
                        - accept-orf
                        - send-orf
                - path-mtu-discovery
                - peer-tracking-policy
                - preference
                - purge-timer
                - rapid-update
                - rapid-withdrawal
                - remove-private
```

config router bgp rib-management

```
            - rib-management
                - ipv4
                    - leak-import
                    - route-table-import
                - ipv6
                    - leak-import
                    - route-table-import
                - label-ipv4
                    - leak-import
                    - route-table-import
                - label-ipv6
                    - route-table-import
            - route-target-list
            - router-id
            - segment-routing
                - prefix-sid-range
                - shutdown
            - selective-label-ip
            - selective-label-ip-prioritization
            - selective-label-ipv4-install
            - send-default
            - shutdown
            - split-horizon
            - sr-policy-import
            - tcp-mss
            - third-party-nexthop
            - vpn-apply-export
            - vpn-apply-import
```

3.4.17.9 configure router confederation Commands

```
            - confederation
```

3.4.17.10 configure router description Commands

```
            – description
```

3.4.17.11 configure router dhcp Commands

```
- dhcp
    - local-dhcp-server
        - description
        - failover
            - ignore-mclt-on-takeover
            - maximum-client-lead-time
            - partner-down-delay
            - peer
            - shutdown
            - startup-wait-time
        - force-renews
        - lease-hold-time
        - lease-hold-time-for
            - internal-lease-ipsec
            - solicited-release
        - pool
            - description
            - failover
                - ignore-mclt-on-takeover
                - maximum-client-lead-time
                - partner-down-delay
                - peer
                - shutdown
                - startup-wait-time
            - max-lease-time
            - min-lease-time
            - minimum-free
            - nak-non-matching-subnet
            - offer-time
            - options
                - custom-option
                - dns-server
                - domain-name
                - lease-rebind-time
                - lease-renew-time
                - lease-time
                - netbios-name-server
                - netbios-node-type
            - subnet
                - address-range
                - drain
                - exclude-addresses
                - maximum-declined
                - minimum-free
                - options
                    - custom-option
                    - default-router
                    - subnet-mask
        - shutdown
        - use-gi-address
        - use-pool-from-client
        - user-db
        - user-ident
```

3.4.17.12 configure router dhcp6 Commands

```
            - dhcp6
                - local-dhcp-server
                    - allow-lease-query
                    - description
                    - failover
                        - ignore-mclt-on-takeover
                        - maximum-client-lead-time
                        - partner-down-delay
                        - peer
                        - shutdown
                        - startup-wait-time
                    - ignore-rapid-commit
                    - interface-id-mapping
                    - lease-hold-time
                    - lease-hold-time-for
                        - internal-lease-ipsec
                        - solicited-release
                    - pool
                        - delegated-prefix-length
                        - description
                        - exclude-prefix
                        - failover
                            - ignore-mclt-on-takeover
                            - maximum-client-lead-time
                            - partner-down-delay
                            - peer
                            - shutdown
                            - startup-wait-time
                        - options
                            - custom-option
                            - dns-server
                            - domain-name
                        - prefix
                            - drain
                            - options
                                - custom-option
                                - dns-server
                                - domain-name
                            - preferred-lifetime
                            - rebind-timer
                            - renew-timer
                            - thresholds
                                - minimum-free
                                    - depleted-event
                                    - minimum
                            - valid-lifetime
                        - thresholds
                            - minimum-free
                                - depleted-event
                                - minimum
                    - server-id
                    - shutdown
                    - use-link-address
                    - use-pool-from-client
                    - user-db
                    - user-ident
```

3.4.17.13 configure router dns Commands

```
            - dns
                - redirect-vprn
                    - service
```

3.4.17.14 configure router ecmp Commands

```
            - ecmp
```

3.4.17.15 configure router fib-priority Commands

```
            - fib-priority
```

3.4.17.16 configure router fib-telemetry Commands

```
            - fib-telemetry
```

3.4.17.17 configure router flexible-algorithm-definitions Commands

```
            - flexible-algorithm-definitions
                - flex-algo
                    - description
                    - exclude
                        - admin-group
                    - flags-tlv
                    - include-all
                        - admin-group
                    - include-any
                        - admin-group
                    - metric-type
                    - priority
                    - shutdown
```

3.4.17.18 configure router icmp-tunneling Commands

```
            - icmp-tunneling
```

3.4.17.19 configure router if-attribute Commands

```
            - if-attribute
                - admin-group
                - srlg-group
```

3.4.17.20 configure router igmp Commands

```
            - igmp
                - grp-if-query-src-ip
                - interface
                    - disable-router-alert-check
                    - import
                    - max-groups
                    - max-grp-sources
                    - max-sources
                    - query-interval
                    - query-last-member-interval
                    - query-response-interval
                    - redundant-multicast
                    - shutdown
                    - ssm-translate
                        - grp-range
                            - source
                    - static
                        - group
                            - source
                            - starg
                    - subnet-check
                    - version
                - query-interval
                - query-last-member-interval
                - query-response-interval
                - robust-count
                - shutdown
                - ssm-translate
                    - grp-range
                        - source
```

3.4.17.21 configure router interface Commands

```
            - interface
                - address
                - allow-directed-broadcasts
                - arp-learn-unsolicited
                - arp-limit
                - arp-proactive-refresh
                - arp-retry-timer
                - arp-timeout
                - autoconfigure
                    - dhcp-client
                        - class-id
                        - client-id
                        - lease-time
                        - request-options
                            - dns-server
                            - router
                            - static-route
                        - shutdown
                - bfd
                - description
                - dhcp
                    - description
                    - gi-address
                    - option
                        - action
                        - circuit-id
                        - remote-id
                        - vendor-specific-option
                            - client-mac-address
                            - pool-name
                            - port-id
                            - service-id
                            - string
                            - system-id
                    - relay-plain-bootp
                    - server
                    - shutdown
                    - trusted
                - dist-cpu-protection
                - egress
                    - filter
                - enable-mac-accounting
                - group-encryption
                    - encryption-keygroup
                    - ip-exception
                - hold-time
                    - down
                    - up
                - icmp
                    - mask-reply
                    - param-problem
                    - redirects
                    - ttl-expired
                    - unreachables
                - if-attribute
                    - admin-group
                    - delay
                        - static
                    - srlg-group
                - ingress
```

config router if ingress filter

```
            - filter
        - ip-helper-address
        - ip-mtu
        - ipsec
            - ip-exception
            - ipsec-tunnel
                - bfd-designate
                - bfd-enable
                - clear-df-bit
                - copy-traffic-class-upon-decapsulation
                - description
                - dynamic-keying
                    - auto-establish
                    - cert
                        - cert-profile
                        - status-verify
                            - default-result
                            - primary
                        - trust-anchor-profile
                    - ike-policy
                    - local-id
                    - ppk
                    - pre-shared-key
                    - transform
                - encapsulated-ip-mtu
                - icmp-generation
                    - frag-required
                        - interval
                        - message-count
                - icmp6-generation
                    - pkt-too-big
                        - interval
                        - message-count
                - ip-mtu
                - local-gateway-address
                - manual-keying
                    - security-association
                - max-history-esp-key-records
                - max-history-ike-key-records
                - pmtu-discovery-aging
                - private-tcp-mss-adjust
                - propagate-pmtu-v4
                - propagate-pmtu-v6
                - public-tcp-mss-adjust
                - remote-gateway-address
                - replay-window
                - security-policy
                - shutdown
            - ipv6-exception
            - shutdown
        - ipv6
            - address
            - bfd
            - dad-disable
            - forward-ipv4-packets
            - icmp6
                - packet-too-big
                - param-problem
                - redirects
                - time-exceeded
                - unreachables
            - link-local-address
            - local-dhcp-server
            - local-proxy-nd
```

config router if ipv6 nd-learn-unsolicited

```
                    - nd-learn-unsolicited
                    - nd-proactive-refresh
                    - neighbor
                    - neighbor-limit
                    - proxy-nd-policy
                    - reachable-time
                    - secure-nd
                        - allow-unsecured-msgs
                        - link-local-modifier
                        - public-key-min-bits
                        - security-parameter
                        - shutdown
                    - stale-time
                    - tcp-mss
                    - urpf-check
                        - ignore-default
                        - mode
                    - vrrp
                        - backup
                        - bfd-enable
                        - init-delay
                        - mac
                        - master-int-inherit
                        - message-interval
                        - ntp-reply
                        - oper-group
                        - ping-reply
                        - policy
                        - preempt
                        - priority
                        - shutdown
                        - standby-forwarding
                        - telnet-reply
                        - traceroute-reply
                - ldp-sync-timer
                - load-balancing
                    - egr-ip-load-balancing
                    - lsr-load-balancing
                - local-dhcp-server
                - local-proxy-arp
                - loopback
                - mac
                - network-domain
                - port
                - proxy-arp-policy
                - qos
                - remote-proxy-arp
                - secondary
                - shutdown
                - static-arp
                - tcp-mss
                - tos-marking-state
                - unnumbered
                - untrusted
                - urpf-check
                    - ignore-default
                    - mode
                - urpf-selected-vprns
                - vrrp
                    - authentication-key
                    - backup
                    - bfd-enable
                    - init-delay
                    - mac
```

config router if vrrp master-int-inherit

```
                    - master-int-inherit
                    - message-interval
                    - ntp-reply
                    - oper-group
                    - ping-reply
                    - policy
                    - preempt
                    - priority
                    - shutdown
                    - ssh-reply
                    - standby-forwarding
                    - telnet-reply
                    - traceroute-reply
```

3.4.17.22 configure router ipsec Commands

```
            - ipsec
                - multi-chassis-shunt-interface
                    - next-hop
                - multi-chassis-shunting-profile
                    - peer
                        - multi-chassis-shunt-interface
                - security-policy
                    - entry
                        - local-ip
                        - local-v6-ip
                        - remote-ip
                        - remote-v6-ip
```

3.4.17.23 configure router ipv6 Commands

```
            - ipv6
                - reachable-time
                - stale-time
```

3.4.17.24 configure router ipv6-te-router-id Commands

```
            - ipv6-te-router-id
```

3.4.17.25 configure router isis Commands

```
            - isis
                - advertise-passive-only
                - advertise-router-capability
                - advertise-tunnel-link
                - all-l1isis
                - all-l2isis
                - area-id
                - auth-keychain
                - authentication-check
                - authentication-key
                - authentication-type
                - csnp-authentication
                - csnp-on-p2p
                - database-export
                - default-route-tag
                - disable-ldp-sync
                - export
                - export-limit
                - flexible-algorithms
                    - advertise-admin-group
                    - flex-algo
                        - advertise
                        - loopfree-alternates
                        - micro-loop-avoidance
                        - participate
                    - shutdown
                - graceful-restart
                    - helper-disable
                - hello-authentication
                - hello-padding
                - ignore-attached-bit
                - ignore-lsp-errors
                - ignore-narrow-metric
                - igp-shortcut
                    - allow-sr-over-srte
                    - shutdown
                    - tunnel-next-hop
                        - family
                            - resolution
                            - resolution-filter
                                - rsvp
                                - sr-te
                - iid-tlv-enable
                - import
                - interface
                    - adjacency-set
                    - bfd-enable
                    - csnp-interval
                    - default-instance
                    - delay-normalization
                        - delay-tolerance-interval
                        - minimum-delay
                    - flex-algo
                        - ipv4-node-sid
                        - ipv6-node-sid
                    - hello-auth-keychain
                    - hello-authentication
                    - hello-authentication-key
                    - hello-authentication-type
                    - hello-padding
```

config router isis interface interface-type

```
                    - interface-type
                    - ipv4-adjacency-sid
                    - ipv4-multicast-disable
                    - ipv4-node-sid
                    - ipv6-adjacency-sid
                    - ipv6-multicast-disable
                    - ipv6-node-sid
                    - ipv6-unicast-disable
                    - level
                        - hello-auth-keychain
                        - hello-authentication-key
                        - hello-authentication-type
                        - hello-interval
                        - hello-multiplier
                        - hello-padding
                        - ipv4-multicast-metric
                        - ipv6-multicast-metric
                        - ipv6-unicast-metric
                        - metric
                        - passive
                        - priority
                        - sd-offset
                        - sf-offset
                    - level-capability
                    - lfa-policy-map
                    - load-balancing-weight
                    - loopfree-alternate-exclude
                    - lsp-pacing-interval
                    - mesh-group
                    - passive
                    - retransmit-interval
                    - shutdown
                    - sid-protection
                    - tag
                    - tag
                - ipv4-multicast-routing
                - ipv4-routing
                - ipv6-multicast-routing
                - ipv6-routing
                - ldp-over-rsvp
                - level
                    - advertise-router-capability
                    - auth-keychain
                    - authentication-key
                    - authentication-type
                    - csnp-authentication
                    - database-export-exclude
                    - default-ipv4-multicast-metric
                    - default-ipv6-multicast-metric
                    - default-ipv6-unicast-metric
                    - default-metric
                    - external-preference
                    - hello-authentication
                    - hello-padding
                    - loopfree-alternate-exclude
                    - lsp-mtu-size
                    - preference
                    - psnp-authentication
                    - wide-metrics-only
                - level-capability
                - link-group
                    - description
                    - level
                        - ipv4-multicast-metric-offset
```

config router isis link-group level ipv4-unicast-metric-offset

```
                    - ipv4-unicast-metric-offset
                    - ipv6-multicast-metric-offset
                    - ipv6-unicast-metric-offset
                    - member
                    - oper-members
                    - revert-members
                - loopfree-alternates
                    - augment-route-table
                    - exclude
                        - prefix-policy
                    - multi-homed-prefix
                        - preference
                    - remote-lfa
                        - node-protect
                    - ti-lfa
                        - node-protect
                - lsp-lifetime
                - lsp-minimum-remaining-lifetime
                - lsp-mtu-size
                - lsp-refresh-interval
                - multi-topology
                    - ipv4-multicast
                    - ipv6-multicast
                    - ipv6-unicast
                - multicast-import
                - overload
                - overload-export-external
                - overload-export-interlevel
                - overload-fib-error-notify-only
                - overload-include-locators
                - overload-on-boot
                - poi-tlv-enable
                - prefix-attributes-tlv
                - prefix-limit
                - prefix-unreachable
                    - maximum-number-upas
                    - process-received-upa
                    - upa-lifetime
                    - upa-metric
                - psnp-authentication
                - reference-bandwidth
                - rib-priority
                - router-id
                - segment-routing
                    - adj-sid-hold
                    - adjacency-set
                        - family
                        - parallel
                        - sid
                    - adjacency-sid
                        - allocate-dual-sids
                    - export-tunnel-table
                    - mapping-server
                        - shutdown
                        - sid-map
                    - maximum-sid-depth
                        - override-bmi
                        - override-erld
                    - micro-loop-avoidance
                    - multi-topology
                    - prefix-sid-range
                    - shutdown
                    - srlb
                    - tunnel-mtu
```

config router isis segment-routing tunnel-table-pref

```
                    - tunnel-table-pref
                - shutdown
                - standard-multi-instance
                - strict-adjacency-check
                - summary-address
                - suppress-attached-bit
                - system-id
                - timers
                    - lsp-wait
                    - spf-wait
                - traffic-engineering
                - traffic-engineering-options
                    - advertise-delay
                    - application-link-attributes
                        - legacy
                    - ipv6
                - unicast-import-disable
```

3.4.17.26 configure router ldp Commands

```
            - ldp
                - aggregate-prefix-match
                    - prefix-exclude
                    - shutdown
                - consider-system-ip-in-gep
                - export
                - export-tunnel-table
                - fast-reroute
                - fec-originate
                - graceful-restart
                    - maximum-recovery-time
                    - neighbor-liveness-time
                - implicit-null-label
                - import
                - import-pmsi-routes
                - import-tunnel-table
                - interface-parameters
                    - interface
                        - bfd-enable
                        - ipv4
                            - fec-type-capability
                                - prefix-ipv4
                                - prefix-ipv6
                            - hello
                            - keepalive
                            - local-lsr-id
                            - shutdown
                            - transport-address
                        - ipv6
                            - fec-type-capability
                                - prefix-ipv4
                                - prefix-ipv6
                            - hello
                            - keepalive
                            - local-lsr-id
                            - shutdown
                            - transport-address
                        - load-balancing-weight
                        - shutdown
                    - ipv4
                        - hello
                        - keepalive
                        - transport-address
                    - ipv6
                        - hello
                        - keepalive
                        - transport-address
                - label-withdrawal-delay
                - legacy-ipv4-lsr-interop
                - max-ecmp-routes
                - prefer-protocol-stitching
                - prefer-tunnel-in-tunnel
                - session-parameters
                    - peer
                        - adv-adj-addr-only
                        - adv-local-lsr-id
                        - community
                        - dod-label-distribution
                        - export-addresses
                        - export-prefixes
```

config router ldp session-params peer fec-limit

```
                        - fec-limit
                        - fec-type-capability
                            - prefix-ipv4
                            - prefix-ipv6
                        - fec129-cisco-interop
                        - import-prefixes
                        - pe-id-mac-flush-interop
                - shortcut-local-ttl-propagate
                - shortcut-transit-ttl-propagate
                - shutdown
                - targeted-session
                    - auto-rx
                        - ipv4
                            - shutdown
                            - tunneling
                    - auto-tx
                        - ipv4
                            - shutdown
                            - tunneling
                    - disable-targeted-session
                    - export-prefixes
                    - import-prefixes
                    - ipv4
                        - hello
                        - hello-reduction
                        - keepalive
                    - ipv6
                        - hello
                        - hello-reduction
                        - keepalive
                    - peer
                        - bfd-enable
                        - hello
                        - hello-reduction
                        - keepalive
                        - local-lsr-id
                        - shutdown
                        - tunneling
                            - lsp
                    - peer-template
                        - adv-local-lsr-id
                        - bfd-enable
                        - community
                        - hello
                        - hello-reduction
                        - keepalive
                        - local-lsr-id
                        - shutdown
                        - tunneling
                    - peer-template-map
                    - resolve-v6-prefix-over-shortcut
                - tcp-session-parameters
                    - auth-keychain
                    - authentication-key
                    - peer-transport
                        - auth-keychain
                        - authentication-key
                        - path-mtu-discovery
                        - ttl-security
                - tunnel-down-damp-time
                - tunnel-table-pref
                - weighted-ecmp
```

3.4.17.27 configure router ldp-shortcut Commands

```
            - ldp-shortcut
```

3.4.17.28 configure router leak-export Commands

```
            - leak-export
```

3.4.17.29 configure router leak-export-limit Commands

```
            - leak-export-limit
```

3.4.17.30 configure router mc-maximum-routes Commands

```
            – mc-maximum-routes
```

3.4.17.31 configure router mld Commands

```
            - mld
                - grp-if-query-src-ip
                - interface
                    - disable-router-alert-check
                    - import
                    - max-groups
                    - max-grp-sources
                    - max-sources
                    - query-interval
                    - query-last-listener-interval
                    - query-response-interval
                    - shutdown
                    - ssm-translate
                        - grp-range
                            - source
                    - static
                        - group
                            - source
                            - starg
                    - version
                - query-interval
                - query-last-listener-interval
                - query-response-interval
                - robust-count
                - shutdown
                - ssm-translate
                    - grp-range
                        - source
```

3.4.17.32 configure router mpls Commands

```
            - mpls
                - admin-group-frr
                - auto-lsp
                - bypass-resignal-timer
                - cspf-on-loose-hop
                - dynamic-bypass
                - exponential-backoff-retry
                - forwarding-policies
                    - forwarding-policy
                        - binding-label
                        - egress-statistics
                            - shutdown
                        - ingress-statistics
                            - shutdown
                        - next-hop-group
                            - backup-next-hop
                                - next-hop
                            - primary-next-hop
                                - next-hop
                            - shutdown
                        - preference
                        - revert-timer
                        - shutdown
                    - reserved-label-block
                    - shutdown
                - frr-object
                - hold-timer
                - interface
                    - admin-group
                    - label-map
                        - pop
                        - shutdown
                        - swap
                    - shutdown
                    - srlg-group
                    - te-metric
                - least-fill-min-thd
                - least-fill-reoptim-thd
                - logger-event-bundling
                - lsp
                    - adaptive
                    - admin-tag
                    - adspec
                    - bgp-shortcut
                    - bgp-transport-tunnel
                    - binding-sid
                    - class-type
                    - exclude
                    - exclude-node
                    - fallback-path-computation-method
                    - fast-reroute
                        - hop-limit
                        - node-protect
                        - propagate-admin-group
                    - from
                    - hop-limit
                    - igp-shortcut
                    - include
                    - label-stack-reduction
                    - ldp-over-rsvp
```

config router mpls lsp least-fill

```
                    - least-fill
                    - load-balancing-weight
                    - local-sr-protection
                    - main-ct-retry-limit
                    - max-sr-labels
                    - metric
                    - metric-type
                    - path-computation-method
                    - path-profile
                    - pce-associations
                        - diversity
                        - policy
                    - pce-control
                    - pce-report
                    - primary
                        - adaptive
                        - backup-class-type
                        - bandwidth
                        - class-type
                        - delay-metric-limit
                        - exclude
                        - hop-limit
                        - include
                        - priority
                        - record
                        - record-label
                        - shutdown
                    - propagate-admin-group
                    - retry-limit
                    - retry-timer
                    - revert-timer
                    - rsvp-resv-style
                    - secondary
                        - adaptive
                        - bandwidth
                        - class-type
                        - delay-metric-limit
                        - exclude
                        - hop-limit
                        - include
                        - path-preference
                        - priority
                        - record
                        - record-label
                        - shutdown
                        - srlg
                        - standby
                    - shutdown
                    - to
                    - vprn-auto-bind
                - lsp-bsid-block
                - lsp-history
                    - shutdown
                - lsp-init-retry-timeout
                - lsp-template
                    - adaptive
                    - admin-tag
                    - adspec
                    - backup-class-type
                    - bandwidth
                    - bgp-shortcut
                    - bgp-transport-tunnel
                    - binding-sid
                    - class-type
```

config router mpls lsp-template default-path

```
                    - default-path
                    - delay-metric-limit
                    - exclude
                    - fallback-path-computation-method
                    - family
                    - fast-reroute
                        - hop-limit
                        - node-protect
                        - propagate-admin-group
                    - from
                    - hop-limit
                    - igp-shortcut
                    - include
                    - label-stack-reduction
                    - ldp-over-rsvp
                    - least-fill
                    - load-balancing-weight
                    - local-sr-protection
                    - lsp-self-ping
                    - main-ct-retry-limit
                    - max-sr-labels
                    - metric
                    - metric-type
                    - path-computation-method
                    - path-profile
                    - pce-associations
                        - diversity
                        - policy
                    - pce-control
                    - pce-report
                    - priority
                    - propagate-admin-group
                    - record
                    - record-label
                    - retry-limit
                    - retry-timer
                    - shutdown
                    - vprn-auto-bind
                - max-bypass-associations
                - max-bypass-plr-associations
                - mbb-prefer-current-hops
                - p2p-active-path-fast-retry
                - path
                    - hop
                    - shutdown
                - pce-initiated-lsp
                    - sr-te
                        - shutdown
                - pce-report
                - resignal-on-igp-event
                - resignal-on-igp-overload
                - resignal-timer
                - retry-on-igp-overload
                - secondary-fast-retry-timer
                - shortcut-local-ttl-propagate
                - shortcut-transit-ttl-propagate
                - shutdown
                - sr-te-resignal
                    - resignal-on-igp-event
                    - resignal-on-igp-overload
                    - resignal-timer
                - srlg-database
                    - router-id
                        - interface
```

configure router mpls srlg-database router-id shutdown

```
                        - shutdown
                - srlg-frr
                - static-lsp
                    - metric
                    - push
                    - shutdown
                    - to
                - static-lsp-fast-retry
                - strict-ero-nhop-direct-resolution
                - tunnel-table-pref
                    - rsvp-te
                    - sr-te
                - user-srlg-db
```

3.4.17.33 configure router mpls-labels Commands

```
            - mpls-labels
                - bgp-labels-hold-timer
                - reserved-label-block
                    - start-label
                - sr-labels
                - static-label-range
```

3.4.17.34 configure router mss-adjust-group Commands

```
            - mss-adjust-group
```

3.4.17.35 configure router nat Commands

```
            - nat
                - inside
                    - classic-lsn-max-subscriber-limit
                    - destination-prefix
                    - deterministic
                        - address-map
                            - outside-range
                            - shutdown
                        - prefix-map
                            - map
                            - shutdown
                    - nat-policy
                - outside
                    - downstream-ip-filter
                    - mtu
                    - pool
                        - address-range
                            - description
                            - drain
                        - description
                        - icmp-echo-reply
                        - mode
                        - port-forwarding-dyn-block-reservation
                        - port-forwarding-range
                        - port-reservation
                        - shutdown
                        - subscriber-limit
                        - watermarks
                    - upstream-ip-filter
```

3.4.17.36 configure router network-domains Commands

```
            - network-domains
                - network-domain
                    - description
```

3.4.17.37 configure router origin-validation Commands

```
            - origin-validation
                - rpki-session
                    - connect-retry
                    - description
                    - local-address
                    - port
                    - refresh-time
                    - shutdown
                    - stale-time
                - static-entry
```

3.4.17.38 configure router ospf Commands

```
            - ospf
                - advertise-router-capability
                - advertise-tunnel-link
                - area
                    - advertise-router-capability
                    - area-range
                    - blackhole-aggregate
                    - database-export-exclude
                    - export
                    - import
                    - interface
                        - adjacency-set
                        - adjacency-sid
                        - advertise-router-capability
                        - advertise-subnet
                        - auth-keychain
                        - authentication-key
                        - authentication-type
                        - bfd-enable
                        - dead-interval
                        - delay-normalization
                            - delay-tolerance-interval
                            - minimum-delay
                        - flex-algo
                            - node-sid
                        - hello-interval
                        - interface-type
                        - lfa-policy-map
                        - load-balancing-weight
                        - loopfree-alternate-exclude
                        - lsa-filter-out
                        - message-digest-key
                        - metric
                        - mtu
                        - neighbor
                        - node-sid
                        - passive
                        - poll-interval
                        - priority
                        - retransmit-interval
                        - rib-priority
                        - shutdown
                        - sid-protection
                        - transit-delay
                    - loopfree-alternate-exclude
                    - nssa
                        - area-range
                        - originate-default-route
                        - redistribute-external
                        - summaries
                    - stub
                        - default-metric
                        - summaries
                    - virtual-link
                        - auth-keychain
                        - authentication-key
                        - authentication-type
                        - dead-interval
                        - hello-interval
                        - message-digest-key
```

config router ospf area virtual-link retransmit-interval

```
                        - retransmit-interval
                        - shutdown
                        - transit-delay
                - asbr
                - compatible-rfc1583
                - database-export
                - disable-ldp-sync
                - export
                - export-limit
                - external-db-overflow
                - external-preference
                - flexible-algorithms
                    - advertise-admin-group
                    - flex-algo
                        - advertise
                        - loopfree-alternates
                        - micro-loop-avoidance
                        - participate
                    - shutdown
                - graceful-restart
                    - helper-disable
                    - strict-lsa-checking
                - igp-shortcut
                    - allow-sr-over-srte
                    - shutdown
                    - tunnel-next-hop
                        - family
                            - resolution
                            - resolution-filter
                                - rsvp
                                - sr-te
                - import
                - ldp-over-rsvp
                - loopfree-alternates
                    - augment-route-table
                    - exclude
                        - prefix-policy
                    - multi-homed-prefix
                        - preference
                    - remote-lfa
                        - node-protect
                    - ti-lfa
                        - node-protect
                - multi-instance
                - multicast-import
                - overload
                - overload-include-ext-1
                - overload-include-ext-2
                - overload-include-stub
                - overload-on-boot
                - preference
                - reference-bandwidth
                - rib-priority
                - router-id
                - rtr-adv-lsa-limit
                - segment-routing
                    - adj-sid-hold
                    - adjacency-set
                        - parallel
                        - sid
                    - adjacency-sid
                        - allocate-dual-sids
                    - backup-node-sid
                    - export-tunnel-table
```

config router ospf segm-rtng mapping-server

```
                    - mapping-server
                        - shutdown
                        - sid-map
                    - micro-loop-avoidance
                    - prefix-sid-range
                    - shutdown
                    - srlb
                    - tunnel-mtu
                    - tunnel-table-pref
                - shutdown
                - timers
                    - incremental-spf-wait
                    - lsa-accumulate
                    - lsa-arrival
                    - lsa-generate
                    - redistribute-delay
                    - spf-wait
                - traffic-engineering
                - traffic-engineering-options
                    - advertise-delay
                    - sr-te
                - unicast-import-disable
```

3.4.17.39 configure router ospf3 Commands

```
            - ospf3
                - advertise-router-capability
                - area
                    - advertise-router-capability
                    - area-range
                    - blackhole-aggregate
                    - database-export-exclude
                    - export
                    - extended-lsa
                    - import
                    - interface
                        - advertise-router-capability
                        - authentication
                        - bfd-enable
                        - dead-interval
                        - hello-interval
                        - interface-type
                        - lfa-policy-map
                        - load-balancing-weight
                        - loopfree-alternate-exclude
                        - lsa-filter-out
                        - metric
                        - mtu
                        - neighbor
                        - passive
                        - poll-interval
                        - priority
                        - retransmit-interval
                        - rib-priority
                        - shutdown
                        - transit-delay
                    - key-rollover-interval
                    - loopfree-alternate-exclude
                    - nssa
                        - area-range
                        - originate-default-route
                        - redistribute-external
                        - summaries
                    - stub
                        - default-metric
                        - summaries
                    - virtual-link
                        - authentication
                        - dead-interval
                        - hello-interval
                        - retransmit-interval
                        - shutdown
                        - transit-delay
                - asbr
                - database-export
                - disable-ldp-sync
                - export
                - export-limit
                - extended-lsa
                - external-db-overflow
                - external-preference
                - graceful-restart
                    - helper-disable
                    - strict-lsa-checking
                - igp-shortcut
```

configure router ospf3 igp-shortcut shutdown

```
                    - shutdown
                    - tunnel-next-hop
                        - family
                            - resolution
                            - resolution-filter
                                - rsvp
                                - sr-te
                - import
                - loopfree-alternates
                    - exclude
                        - prefix-policy
                - multicast-import
                - overload
                - overload-include-ext-1
                - overload-include-ext-2
                - overload-include-stub
                - overload-on-boot
                - preference
                - reference-bandwidth
                - rib-priority
                - router-id
                - rtr-adv-lsa-limit
                - shutdown
                - timers
                    - incremental-spf-wait
                    - lsa-accumulate
                    - lsa-arrival
                    - lsa-generate
                    - redistribute-delay
                    - spf-wait
                - unicast-import-disable
```

3.4.17.40 configure router pcep Commands

```
            - pcep
                - pcc
                    - dead-timer
                    - keepalive
                    - local-address
                    - local-address-ipv6
                    - max-srte-pce-init-lsps
                    - pce-associations
                        - diversity
                            - association-id
                            - association-source
                            - disjointness-reference
                            - disjointness-type
                            - diversity-type
                        - policy
                            - association-id
                            - association-source
                    - peer
                        - auth-keychain
                        - route-preference
                        - shutdown
                        - tls-client-profile
                        - tls-wait-timer
                    - redelegation-timer
                    - report-path-constraints
                    - shutdown
                    - state-timer
                    - unknown-message-rate
```

3.4.17.41 configure router pim Commands

```
            - pim
                - apply-to
                - enable-mdt-spt
                - import
                - interface
                    - assert-period
                    - bfd-enable
                    - bsm-check-rtr-alert
                    - hello-interval
                    - hello-multiplier
                    - improved-assert
                    - instant-prune-echo
                    - ipv4-multicast-disable
                    - ipv6-multicast-disable
                    - max-groups
                    - monitor-oper-group
                    - multicast-senders
                    - priority
                    - shutdown
                    - sticky-dr
                    - three-way-hello
                    - tracking-support
                - ipv4-multicast-disable
                - ipv6-multicast-disable
                - non-dr-attract-traffic
                - pim-ssm-scaling
                - rp
                    - anycast
                        - rp-set-peer
                    - auto-rp-discovery
                    - bootstrap-export
                    - bootstrap-import
                    - bsr-candidate
                        - address
                        - hash-mask-len
                        - priority
                        - shutdown
                    - ipv6
                        - anycast
                            - rp-set-peer
                        - bsr-candidate
                            - address
                            - hash-mask-len
                            - priority
                            - shutdown
                        - embedded-rp
                            - group-range
                            - shutdown
                        - rp-candidate
                            - address
                            - group-range
                            - holdtime
                            - priority
                            - shutdown
                        - static
                            - address
                                - group-prefix
                                - override
                    - rp-candidate
                        - address
```

config router pim rp rp-candidate group-range

```
                    - group-range
                    - holdtime
                    - priority
                    - shutdown
                - static
                    - address
                        - group-prefix
                        - override
            - rpf-table
            - rpf6-table
            - rpfv
            - shutdown
            - source-address
                - register-message
            - spt-switchover-threshold
            - ssm-assert-compatible-mode
            - ssm-default-range-disable
            - ssm-groups
                - group-range
```

3.4.17.42 configure router policy-options Commands

```
- policy-options
    - abort
    - as-path
        - expression
    - as-path-group
        - entry
    - begin
    - commit
    - community
        - expression
        - members
    - damping
        - half-life
        - max-suppress
        - reuse
        - suppress
    - exclusive-lock-time
    - global-variables
        - name
    - policy-statement
        - default-action
            - add-paths-send-limit
            - admin-tag-policy
            - advertise-label
            - aigp-metric
            - as-path
            - as-path-prepend
            - bgp-high-priority
            - bgp-leak
            - bgp-med
            - bgp-tunnel-metric
            - community
            - create-mpls-tunnel
            - create-udp-tunnel
            - damping
            - disable-route-table-install
            - flex-algo
            - install-backup-path
            - local-preference
            - metric
            - next-hop
            - next-hop-self
            - origin
            - origin-validation-state
            - preference
            - resolve-static
            - sr-label-index
            - sr-maintenance-policy
            - sticky-ecmp
            - tag
            - type
        - description
        - entry
            - action
                - add-paths-send-limit
                - admin-tag-policy
                - advertise-label
                - aigp-metric
                - as-path
                - as-path-prepend
```

config router policy-options policy-statement entry action bgp-high-priority

```
                - bgp-high-priority
                - bgp-leak
                - bgp-med
                - bgp-tunnel-metric
                - community
                - create-mpls-tunnel
                - create-udp-tunnel
                - damping
                - disable-route-table-install
                - flex-algo
                - install-backup-path
                - local-preference
                - metric
                - next-hop
                - next-hop-self
                - origin
                - origin-validation-state
                - preference
                - resolve-static
                - sr-label-index
                - sr-maintenance-policy
                - sticky-ecmp
                - tag
                - type
            - conditional-expression
                - route-exists
            - description
            - from
                - aggregate-contributor
                - area
                - as-path
                - as-path-group
                - as-path-length
                - cluster-id
                - color
                - community
                - community-count
                - distinguisher
                - endpoint
                - evpn-type
                - external
                - family
                - group-address
                - host-ip
                - interface
                - interface-subnets
                - level
                - local-preference
                - metric
                - neighbor
                - next-hop
                - origin
                - origin-validation-state
                - path-type
                - policy
                - policy-variables
                    - name
                - prefix-list
                - prefix-list-override
                - protocol
                - route-distinguisher-list
                - source-address
                - state
                - tag
```

config router policy-options policy-statement entry from type

```
                        - type
                    - to
                        - level
                        - neighbor
                        - prefix-list
                        - protocol
                - renumber
            - prefix-list
                - prefix
            - route-distinguisher-list
                - rd-entry
```

3.4.17.43 configure router policy-reference-checks Commands

```
            – policy-reference-checks
```

3.4.17.44 configure router radius-server Commands

```
            - radius-server
                - server
                    - accept-coa
                    - acct-port
                    - auth-port
                    - coa-script-policy
                    - description
                    - pending-requests-limit
```

3.4.17.45 configure router reassembly-group Commands

```
            - reassembly-group
```

3.4.17.46 configure router rip Commands

```
            - rip
                - authentication-key
                - authentication-type
                - bfd-enable
                - check-zero
                - description
                - export
                - export-limit
                - group
                    - authentication-key
                    - authentication-type
                    - bfd-enable
                    - check-zero
                    - description
                    - export
                    - import
                    - message-size
                    - metric-in
                    - metric-out
                    - neighbor
                        - authentication-key
                        - authentication-type
                        - bfd-enable
                        - check-zero
                        - description
                        - export
                        - import
                        - message-size
                        - metric-in
                        - metric-out
                        - preference
                        - receive
                        - send
                        - shutdown
                        - split-horizon
                        - timers
                        - unicast-address
                    - preference
                    - receive
                    - send
                    - shutdown
                    - split-horizon
                    - timers
                - import
                - message-size
                - metric-in
                - metric-out
                - preference
                - receive
                - send
                - shutdown
                - split-horizon
                - timers
```

3.4.17.47 configure router ripng Commands

```
            - ripng
                - bfd-enable
                - check-zero
                - description
                - export
                - export-limit
                - group
                    - bfd-enable
                    - check-zero
                    - description
                    - export
                    - import
                    - message-size
                    - metric-in
                    - metric-out
                    - neighbor
                        - bfd-enable
                        - check-zero
                        - description
                        - export
                        - import
                        - message-size
                        - metric-in
                        - metric-out
                        - preference
                        - receive
                        - send
                        - shutdown
                        - split-horizon
                        - timers
                        - unicast-address
                    - preference
                    - receive
                    - send
                    - shutdown
                    - split-horizon
                    - timers
                - import
                - message-size
                - metric-in
                - metric-out
                - preference
                - receive
                - send
                - shutdown
                - split-horizon
                - timers
```

3.4.17.48 configure router route-next-hop-policy Commands

```
            - route-next-hop-policy
                - abort
                - begin
                - commit
                - template
                    - description
                    - exclude-group
                    - include-group
                    - nh-type
                    - protection-type
                    - srlg-enable
```

3.4.17.49 configure router router-advertisement Commands

```
            - router-advertisement
                - dns-options
                    - rdnss-lifetime
                    - server
                - interface
                    - current-hop-limit
                    - dns-options
                        - include-dns
                        - rdnss-lifetime
                        - server
                    - managed-configuration
                    - max-advertisement-interval
                    - min-advertisement-interval
                    - mtu
                    - nd-router-preference
                    - other-stateful-configuration
                    - prefix
                        - autonomous
                        - on-link
                        - preferred-lifetime
                        - valid-lifetime
                    - reachable-time
                    - retransmit-time
                    - router-lifetime
                    - shutdown
                    - use-virtual-mac
```

3.4.17.50 configure router router-id Commands

```
            - router-id
```

3.4.17.51 configure router rsvp Commands

```
            - rsvp
                - authentication-over-bypass
                - diffserv-te
                    - class-type-bw
                    - fc
                    - te-class
                - gr-helper-time
                - graceful-shutdown
                - implicit-null-label
                - interface
                    - auth-keychain
                    - authentication-key
                    - bfd-enable
                    - class-type-bw
                    - gr-helper
                    - graceful-shutdown
                    - hello-interval
                    - implicit-null-label
                    - refresh-reduction
                        - reliable-delivery
                    - shutdown
                    - subscription
                    - te-down-threshold
                    - te-up-threshold
                - keep-multiplier
                - msg-pacing
                    - max-burst
                    - period
                - node-id-in-rro
                - p2p-merge-point-abort-timer
                - preemption-timer
                - rapid-retransmit-time
                - rapid-retry-limit
                - refresh-reduction-over-bypass
                - refresh-time
                - shutdown
                - te-down-threshold
                - te-threshold-update
                    - on-cac-failure
                    - update-timer
                - te-up-threshold
```

3.4.17.52 configure router segment-routing Commands

```
- segment-routing
    - maintenance-policy
        - bfd-enable
        - bfd-template
        - hold-down-timer
        - mode
        - return-path-label
        - revert-timer
        - shutdown
        - threshold
    - sr-mpls
        - prefix-sids
            - ipv4-sid
            - ipv6-sid
            - node-sid
    - sr-policies
        - reserved-label-block
        - shutdown
        - static-policy
            - binding-sid
            - color
            - distinguisher
            - endpoint
            - head-end
            - maintenance-policy
            - preference
            - segment-list
                - segment
                    - mpls-label
                - shutdown
                - weight
            - shutdown
```

3.4.17.53 configure router sgt-qos Commands

```
            - sgt-qos
                - application
                - dscp
```

3.4.17.54 configure router single-sfm-overload Commands

```
            - single-sfm-overload
```

3.4.17.55 configure router static-route-entry Commands

```
            - static-route-entry
                - black-hole
                    - community
                    - description
                    - dynamic-bgp
                    - generate-icmp
                    - metric
                    - preference
                    - prefix-list
                    - shutdown
                    - tag
                - community
                - indirect
                    - community
                    - cpe-check
                        - drop-count
                        - interval
                        - log
                        - padding-size
                    - description
                    - metric
                    - preference
                    - prefix-list
                    - shutdown
                    - tag
                    - tunnel-next-hop
                        - disallow-igp
                        - flex-algo
                        - resolution
                        - resolution-filter
                            - ldp
                            - rsvp-te
                                - lsp
                            - sr-isis
                            - sr-ospf
                            - sr-ospf3
                            - sr-te
                                - lsp
                - ipsec-tunnel
                    - community
                    - description
                    - metric
                    - preference
                    - shutdown
                    - tag
                - leak-destination
                    - router-instance
                - next-hop
                    - bfd-enable
                    - community
                    - cpe-check
                        - drop-count
                        - interval
                        - log
                        - padding-size
                    - description
                    - ldp-sync
                    - load-balancing-weight
                    - metric
                    - preference
```

config router static-route-entry next-hop prefix-list

```
                    - prefix-list
                    - shutdown
                    - tag
                    - validate-next-hop
                - tag
```

3.4.17.56 configure router static-route-hold-down Commands

```
            - static-route-hold-down
```

3.4.17.57 configure router subscriber-mgmt Commands

```
            – subscriber-mgmt
```

3.4.17.58 configure router triggered-policy Commands

```
            - triggered-policy
```

3.4.17.59 configure router ttl-propagate Commands

```
            - ttl-propagate
                - label-route-local
                - label-route-transit
                - lsr-label-route
                - sr-mpls-local
                - sr-mpls-transit
                - vprn-local
                - vprn-transit
```

3.4.17.60 configure router tunnel-interface Commands

```
            - tunnel-interface
                - description
```

3.4.17.61 configure router twamp-light Commands

```
            - twamp-light
                - reflector
                    - allow-ipv6-udp-checksum-zero
                    - description
                    - prefix
                        - description
                    - shutdown
                    - type
```

3.4.17.62 configure router weighted-ecmp Commands

```
            - weighted-ecmp
```

3.4.18 configure saa Commands

```
- saa
    - test
        - accounting-policy
        - continuous
        - description
        - jitter-event
        - latency-event
        - loss-event
        - probe-history
        - shutdown
        - trap-gen
            - probe-fail-enable
            - probe-fail-threshold
            - test-completion-enable
            - test-fail-enable
            - test-fail-threshold
        - type
            - icmp-ping
            - lsp-ping
            - sdp-ping
        - type-multi-line
            - lsp-ping
                - fc
                - interval
                - profile
                - send-count
                - size
                - sr-policy
                    - fc
                    - interval
                    - path-destination
                    - profile
                    - segment-list
                    - send-count
                    - size
                    - src-ip-address
                    - timeout
                    - ttl
                - src-ip-address
                - timeout
                - ttl
```

3.4.19 configure service Commands

```
        – service
```

3.4.19.1 configure service customer Commands

```
            - customer
                - contact
                - description
                - multi-service-site
                    - assignment
                    - description
                    - egress
                        - policer-control-policy
                        - scheduler-override
                            - scheduler
                                - parent
                                - rate
                        - scheduler-policy
                    - ingress
                        - policer-control-policy
                        - scheduler-override
                            - scheduler
                                - parent
                                - rate
                        - scheduler-policy
                - phone
```

3.4.19.2 configure service epipe Commands

```
            – epipe
                – bgp
                    – adv-service-mtu
                    – pw-template-binding
                    – route-distinguisher
                    – route-target
                    – vsi-export
                    – vsi-import
                – bgp-evpn
                    – evi
                    – local-attachment-circuit
                        – eth-tag
                    – mpls
                        – auto-bind-tunnel
                            – allow-flex-algo-fallback
                            – ecmp
                            – enforce-strict-tunnel-tagging
                            – enforce-untagged-route
                            – resolution
                            – resolution-filter
                                – bgp
                                – ldp
                                – rsvp
                                – sr-isis
                                – sr-ospf
                                – sr-ospf3
                                – sr-policy
                                – sr-te
                            – weighted-ecmp
                        – control-word
                        – default-route-tag
                        – domain-id
                        – dynamic-egress-label-limit
                        – ecmp
                        – evi-three-byte-auto-rt
                        – force-vlan-vc-forwarding
                        – hash-label
                        – mh-mode
                        – oper-group
                        – route-next-hop
                        – send-tunnel-encap
                        – shutdown
                    – remote-attachment-circuit
                        – eth-tag
                – bgp-vpws
                    – remote-ve-name
                        – ve-id
                    – shutdown
                    – ve-name
                        – ve-id
                – description
                – endpoint
                    – active-hold-delay
                    – description
                    – revert-time
                    – standby-signaling-master
                    – standby-signaling-slave
                – ignore-l2vpn-mtu-mismatch
                – load-balancing
                – oper-group
```

config service epipe sap

```
            - sap
                - accounting-policy
                - bandwidth
                - collect-stats
                - description
                - dist-cpu-protection
                - egress
                    - agg-rate
                        - queue-frame-based-accounting
                        - rate
                    - filter
                    - policer-control-override
                        - max-rate
                        - priority-mbs-thresholds
                            - min-thresh-separation
                            - priority
                                - mbs-contribution
                    - policer-control-policy
                    - policer-override
                        - policer
                            - cbs
                            - mbs
                            - packet-byte-offset
                            - percent-rate
                            - rate
                            - stat-mode
                    - qinq-mark-top-only
                    - qos
                    - queue-override
                        - queue
                            - adaptation-rule
                            - avg-frame-overhead
                            - burst-limit
                            - cbs
                            - drop-tail
                                - low
                                    - percent-reduction-from-mbs
                            - mbs
                            - parent
                            - percent-rate
                            - rate
                    - scheduler-override
                        - scheduler
                            - parent
                            - rate
                    - scheduler-policy
                - ignore-oper-down
                - ingress
                    - aggregate-policer
                        - cbs
                        - mbs
                        - rate
                    - filter
                    - match-qinq-dot1p
                    - policer-control-override
                        - max-rate
                        - priority-mbs-thresholds
                            - min-thresh-separation
                            - priority
                                - mbs-contribution
                    - policer-control-policy
                    - policer-override
                        - policer
                            - cbs
```

config service epipe sap ingress policer-over plcr mbs

```
                            - mbs
                            - packet-byte-offset
                            - percent-rate
                            - rate
                            - stat-mode
                        - qos
                        - queue-override
                            - queue
                                - adaptation-rule
                                - cbs
                                - drop-tail
                                    - low
                                        - percent-reduction-from-mbs
                                - mbs
                                - parent
                                - percent-rate
                                - rate
                        - scheduler-override
                            - scheduler
                                - parent
                                - rate
                        - scheduler-policy
                    - monitor-oper-group
                    - multi-service-site
                    - oper-group
                    - ring-node
                    - shutdown
                - service-mtu
                - shutdown
                - site
                    - boot-timer
                    - sap
                    - shutdown
                    - site-activation-timer
                    - site-id
                    - site-min-down-timer
                    - site-preference
                - spoke-sdp
                    - accounting-policy
                    - adv-service-mtu
                    - bandwidth
                    - block-on-peer-fault
                    - collect-stats
                    - control-channel-status
                        - acknowledgment
                        - refresh-timer
                        - request-timer
                        - shutdown
                    - control-word
                    - description
                    - egress
                        - filter
                        - qos
                        - vc-label
                    - force-vlan-vc-forwarding
                    - hash-label
                    - ingress
                        - filter
                        - qos
                        - vc-label
                    - monitor-oper-group
                    - oper-group
                    - precedence
                    - pw-status-signaling
```

config service epipe spoke-sdp shutdown

```
                    - shutdown
                    - standby-signaling-slave
                    - vlan-vc-tag
                - spoke-sdp-fec
                    - auto-config
                    - path
                    - precedence
                    - pw-template-bind
                    - retry-count
                    - retry-timer
                    - saii-type2
                    - shutdown
                    - signaling
                    - standby-signaling-slave
                    - taii-type2
```

3.4.19.3 configure service ies Commands

```
            - ies
                - description
                - interface
                    - address
                    - allow-directed-broadcasts
                    - arp-host-route
                        - populate
                    - arp-learn-unsolicited
                    - arp-limit
                    - arp-populate
                    - arp-proactive-refresh
                    - arp-retry-timer
                    - arp-timeout
                    - bfd
                    - description
                    - dhcp
                        - description
                        - gi-address
                        - lease-populate
                        - option
                            - action
                            - circuit-id
                            - remote-id
                            - vendor-specific-option
                                - client-mac-address
                                - pool-name
                                - sap-id
                                - service-id
                                - string
                                - system-id
                        - proxy-server
                            - emulated-server
                            - lease-time
                            - shutdown
                        - relay-plain-bootp
                        - relay-proxy
                        - server
                        - shutdown
                        - trusted
                        - use-arp
                    - dynamic-tunnel-redundant-next-hop
                    - enable-mac-accounting
                    - hold-time
                        - down
                        - up
                    - icmp
                        - mask-reply
                        - param-problem
                        - redirects
                        - ttl-expired
                        - unreachables
                    - if-attribute
                        - admin-group
                        - srlg-group
                    - ingress
                    - ip-helper-address
                    - ip-mtu
                    - ipsec
                        - ip-exception
                        - ipsec-tunnel
```

config service ies if ipsec ipsec-tunnel bfd-designate

```
                        - bfd-designate
                        - bfd-enable
                        - clear-df-bit
                        - copy-traffic-class-upon-decapsulation
                        - description
                        - dynamic-keying
                            - auto-establish
                            - cert
                                - cert-profile
                                - status-verify
                                    - default-result
                                    - primary
                                - trust-anchor-profile
                            - ike-policy
                            - local-id
                            - ppk
                            - pre-shared-key
                            - transform
                        - encapsulated-ip-mtu
                        - icmp-generation
                            - frag-required
                                - interval
                                - message-count
                        - icmp6-generation
                            - pkt-too-big
                                - interval
                                - message-count
                        - ip-mtu
                        - local-gateway-address
                        - manual-keying
                            - security-association
                        - max-history-esp-key-records
                        - max-history-ike-key-records
                        - pmtu-discovery-aging
                        - private-tcp-mss-adjust
                        - propagate-pmtu-v4
                        - propagate-pmtu-v6
                        - public-tcp-mss-adjust
                        - remote-gateway-address
                        - replay-window
                        - security-policy
                        - shutdown
                    - ipv6-exception
                    - shutdown
                - ipv6
                    - address
                    - bfd
                    - dad-disable
                    - dhcp6-relay
                        - description
                        - lease-populate
                        - lease-populate
                        - option
                            - interface-id
                            - remote-id
                        - server
                        - shutdown
                        - source-address
                    - forward-ipv4-packets
                    - icmp6
                        - packet-too-big
                        - param-problem
                        - redirects
                        - time-exceeded
```

config service ies if ipv6 icmp6 unreachables

```
                    - unreachables
                - link-local-address
                - local-dhcp-server
                - local-proxy-nd
                - nd-host-route
                    - populate
                - nd-learn-unsolicited
                - nd-proactive-refresh
                - neighbor
                - neighbor-limit
                - proxy-nd-policy
                - reachable-time
                - secure-nd
                    - allow-unsecured-msgs
                    - link-local-modifier
                    - public-key-min-bits
                    - security-parameter
                    - shutdown
                - stale-time
                - tcp-mss
                - urpf-check
                    - ignore-default
                    - mode
                - vrrp
                    - backup
                    - bfd-enable
                    - init-delay
                    - mac
                    - master-int-inherit
                    - message-interval
                    - ntp-reply
                    - oper-group
                    - ping-reply
                    - policy
                    - preempt
                    - priority
                    - shutdown
                    - standby-forwarding
                    - telnet-reply
                    - traceroute-reply
            - load-balancing
                - egr-ip-load-balancing
            - local-dhcp-server
            - local-proxy-arp
            - loopback
            - mac
            - monitor-oper-group
            - multi-chassis-shunting-profile
            - multicast-network-domain
            - proxy-arp-policy
            - remote-proxy-arp
            - sap
                - accounting-policy
                - bandwidth
                - collect-stats
                - description
                - dist-cpu-protection
                - egress
                    - agg-rate
                        - queue-frame-based-accounting
                        - rate
                    - filter
                    - policer-control-override
                        - max-rate
```

config service ies if sap egress policer-ctrl-over priority-mbs-thresholds

```
                        - priority-mbs-thresholds
                            - min-thresh-separation
                            - priority
                                - mbs-contribution
                    - policer-control-policy
                    - policer-override
                        - policer
                            - cbs
                            - mbs
                            - packet-byte-offset
                            - percent-rate
                            - rate
                            - stat-mode
                    - qinq-mark-top-only
                    - qos
                    - queue-override
                        - queue
                            - adaptation-rule
                            - avg-frame-overhead
                            - burst-limit
                            - cbs
                            - drop-tail
                                - low
                                    - percent-reduction-from-mbs
                            - mbs
                            - parent
                            - percent-rate
                            - rate
                    - scheduler-override
                        - scheduler
                            - parent
                            - rate
                    - scheduler-policy
                - ingress
                    - aggregate-policer
                        - cbs
                        - mbs
                        - rate
                    - filter
                    - match-qinq-dot1p
                    - policer-control-override
                        - max-rate
                        - priority-mbs-thresholds
                            - min-thresh-separation
                            - priority
                                - mbs-contribution
                    - policer-control-policy
                    - policer-override
                        - policer
                            - cbs
                            - mbs
                            - packet-byte-offset
                            - percent-rate
                            - rate
                            - stat-mode
                    - qos
                    - queue-override
                        - queue
                            - adaptation-rule
                            - cbs
                            - drop-tail
                                - low
                                    - percent-reduction-from-mbs
                            - mbs
```

config service ies if sap ingress queue-override queue parent

```
                        - parent
                        - percent-rate
                        - rate
                - scheduler-override
                    - scheduler
                        - parent
                        - rate
                - scheduler-policy
            - ip-tunnel
                - backup-remote-ip
                - clear-df-bit
                - delivery-service
                - description
                - dest-ip
                - dscp
                - encapsulated-ip-mtu
                - gre-header
                - icmp-generation
                    - frag-required
                        - interval
                        - message-count
                - icmp6-generation
                    - packet-too-big
                - ip-mtu
                - ipsec-transport-mode-profile
                - pmtu-discovery-aging
                - private-tcp-mss-adjust
                - propagate-pmtu-v4
                - propagate-pmtu-v6
                - public-tcp-mss-adjust
                - reassembly
                - remote-ip
                - shutdown
                - source
            - ipsec-gw
                - cert
                    - cert-profile
                    - status-verify
                        - default-result
                        - primary
                    - trust-anchor-profile
                - client-db
                - default-secure-service
                - default-tunnel-template
                - dhcp
                    - gi-address
                    - send-release
                    - server
                    - shutdown
                - dhcp6
                    - link-address
                    - send-release
                    - server
                    - shutdown
                - ike-policy
                - local-address-assignment
                    - ipv4
                        - address-source
                    - ipv6
                        - address-source
                    - shutdown
                - local-gateway-address
                - local-id
                - max-history-esp-key-records
```

config service ies if sap ipsec-gw max-history-ike-key-records

```
                        - max-history-ike-key-records
                        - pre-shared-key
                        - radius-accounting-policy
                        - radius-authentication-policy
                        - shutdown
                        - ts-negotiation
                    - multi-service-site
                    - shutdown
                - secondary
                - shutdown
                - spoke-sdp
                    - accounting-policy
                    - collect-stats
                    - control-channel-status
                        - acknowledgment
                        - refresh-timer
                        - request-timer
                        - shutdown
                    - description
                    - egress
                        - filter
                        - qos
                        - vc-label
                    - hash-label
                    - ingress
                        - filter
                        - qos
                        - vc-label
                    - shutdown
                - static-arp
                - static-tunnel-redundant-next-hop
                - tcp-mss
                - tos-marking-state
                - unnumbered
                - unnumbered
                - urpf-check
                    - ignore-default
                    - mode
                - vpls
                    - egress
                        - reclassify-using-qos
                        - v4-routed-override-filter
                        - v6-routed-override-filter
                    - evpn
                        - arp
                            - advertise
                            - flood-garp-and-unknown-req
                            - learn-dynamic
                        - nd
                            - advertise
                            - learn-dynamic
                    - ingress
                        - v4-routed-override-filter
                        - v6-routed-override-filter
                - vrrp
                    - authentication-key
                    - backup
                    - bfd-enable
                    - init-delay
                    - mac
                    - master-int-inherit
                    - message-interval
                    - ntp-reply
                    - oper-group
```

config service ies if vrrp ping-reply

```
                        - ping-reply
                        - policy
                        - preempt
                        - priority
                        - shutdown
                        - ssh-reply
                        - standby-forwarding
                        - telnet-reply
                        - traceroute-reply
                - shutdown
```

3.4.19.4 configure service mac-list Commands

```
            - mac-list
                - description
                - mac
```

3.4.19.5 configure service md-auto-id Commands

```
            - md-auto-id
                - customer-id-range
                - pw-template-id-range
                - service-id-range
```

3.4.19.6 configure service nat Commands

```
- nat
    - nat-policy
        - block-limit
        - description
        - filtering
        - pool
        - port-limits
            - forwarding
            - watermarks
        - session-limits
            - max
            - watermarks
        - tcp-mss-adjust
        - timeouts
            - icmp-query
            - tcp-established
            - tcp-syn
            - tcp-time-wait
            - tcp-transitory
            - udp
            - udp-dns
            - udp-initial
        - udp-inbound-refresh
    - port-forwarding
        - lsn
```

3.4.19.7 configure service oper-group Commands

```
            - oper-group
                - bfd-enable
                - hold-time
                    - group
```

3.4.19.8 configure service proxy-arp-nd Commands

```
            - proxy-arp-nd
                - mac-list
                    - mac
```

3.4.19.9 configure service pw-routing Commands

```
            - pw-routing
                - boot-timer
                - local-prefix
                    - advertise-bgp
                - path
                    - hop
                    - shutdown
                - retry-count
                - retry-timer
                - spe-address
                - static-route
```

3.4.19.10 configure service pw-template Commands

```
            - pw-template
                - accounting-policy
                - allow-fragmentation
                - auto-learn-mac-protect
                - block-on-peer-fault
                - collect-stats
                - controlword
                - disable-aging
                - disable-learning
                - discard-unknown-source
                - egress
                    - filter
                    - filter-name
                        - ip
                        - ipv6
                        - mac
                    - mfib-allowed-mda-destinations
                        - mda
                    - qos
                - encryption-keygroup
                - force-vlan-vc-forwarding
                - hash-label
                - igmp-snooping
                    - fast-leave
                    - import
                    - last-member-query-interval
                    - max-num-groups
                    - query-interval
                    - query-response-interval
                    - robust-count
                    - send-queries
                    - version
                - ingress
                    - filter
                    - filter-name
                        - ip
                        - ipv6
                        - mac
                    - qos
                - l2pt-termination
                - limit-mac-move
                - mac-pinning
                - max-nbr-mac-addr
                - path-mtu
                - restrict-protected-src
                - sdp-exclude
                - sdp-include
                - split-horizon-group
                    - auto-learn-mac-protect
                    - description
                    - restrict-protected-src
                    - restrict-unprotected-dst
                - stp
                    - auto-edge
                    - edge-port
                    - link-type
                    - path-cost
                    - priority
                    - root-guard
                    - shutdown
```

config service pw-template vc-type

```
                - vc-type
                - vlan-vc-tag
```

3.4.19.11 configure service sdp Commands

```
            - sdp
                - accounting-policy
                - adv-mtu-override
                - allow-fragmentation
                - bgp-tunnel
                - booking-factor
                - collect-stats
                - description
                - encryption-keygroup
                - far-end
                - keep-alive
                    - hello-time
                    - hold-down-time
                    - max-drop-count
                    - message-length
                    - shutdown
                    - timeout
                - ldp
                - local-end
                - lsp
                - metric
                - mixed-lsp-mode
                    - revert-time
                - network-domain
                - path-mtu
                - sdp-group
                - shutdown
                - signaling
                - sr-isis
                - sr-ospf
                - sr-te-lsp
                - tunnel-far-end
                - vlan-vc-etype
                - weighted-ecmp
```

3.4.19.12 configure service sdp-group Commands

```
            - sdp-group
                - group-name
```

3.4.19.13 configure service system Commands

```
- system
    - bgp-auto-rd-range
    - bgp-evpn
        - ad-per-evi-routes
            - attribute-propagation
            - bgp-path-selection
            - d-path-ignore
        - evpn-etree-leaf-label
        - ip-prefix-routes
            - interface-ful
                - attribute-uniform-propagation
                - bgp-path-selection
        - multicast-leave-sync-propagation
        - route-distinguisher
    - fdb-table-size
    - vpn-gre-source-ip
```

3.4.19.14 configure service template Commands

```
            - template
                - vpls-sap-template
                    - bpdu-translation
                    - collect-stats
                    - disable-aging
                    - disable-learning
                    - discard-unknown-source
                    - dist-cpu-protection
                    - egress
                        - filter
                        - filter-name
                            - ip
                            - ipv6
                            - mac
                        - policer-control-policy
                        - qinq-mark-top-only
                        - qos
                        - scheduler-policy
                    - ingress
                        - filter
                        - filter-name
                            - ip
                            - ipv6
                            - mac
                        - match-qinq-dot1p
                        - policer-control-policy
                        - qos
                        - scheduler-policy
                    - l2pt-termination
                    - limit-mac-move
                    - mac-move-level
                    - max-nbr-mac-addr
                    - stp
                        - auto-edge
                        - edge-port
                        - link-type
                        - path-cost
                        - priority
                        - root-guard
                        - shutdown
                - vpls-template
                    - customer
                    - disable-aging
                    - disable-learning
                    - discard-unknown
                    - fdb-table-high-wmark
                    - fdb-table-low-wmark
                    - fdb-table-size
                    - load-balancing
                    - local-age
                    - mac-move
                        - move-frequency
                        - number-retries
                        - primary-ports
                            - cumulative-factor
                        - retry-timeout
                        - secondary-ports
                            - cumulative-factor
                        - shutdown
                    - remote-age
```

config service template vpls-template service-mtu

```
                    - service-mtu
                    - stp
                        - forward-delay
                        - hello-time
                        - hold-count
                        - max-age
                        - mode
                        - priority
                        - shutdown
                    - temp-flooding
```

3.4.19.15 configure service vpls Commands

```
            - vpls
                - allow-ip-int-bind
                    - evpn-mpls-ecmp
                    - mld-snooping
                        - mrouter-port
                - bgp
                    - adv-service-mtu
                    - pw-template-binding
                        - monitor-oper-group
                        - oper-group
                    - route-distinguisher
                    - route-target
                    - vsi-export
                    - vsi-import
                - bgp-ad
                    - shutdown
                    - vpls-id
                    - vsi-id
                        - prefix
                - bgp-evpn
                    - arp-nd-extended-community-advertisement
                    - arp-nd-only-with-fdb-advertisement
                    - evi
                    - ignore-mtu-mismatch
                    - incl-mcast-l2-attributes-advertisement
                    - incl-mcast-orig-ip
                    - ingress-repl-inc-mcast-advertisement
                    - ip-route-advertisement
                    - ip-route-link-bandwidth
                        - advertise
                        - weighted-ecmp
                    - mac-advertisement
                    - mac-duplication
                        - black-hole-dup-mac
                        - detect
                        - retry
                        - trusted-mac-time
                    - mpls
                        - auto-bind-tunnel
                            - allow-flex-algo-fallback
                            - ecmp
                            - enforce-strict-tunnel-tagging
                            - enforce-untagged-route
                            - resolution
                            - resolution-filter
                                - bgp
                                - ldp
                                - rsvp
                                - sr-isis
                                - sr-ospf
                                - sr-ospf3
                                - sr-policy
                                - sr-te
                            - weighted-ecmp
                        - control-word
                        - default-route-tag
                        - dynamic-egress-label-limit
                        - evi-three-byte-auto-rt
                        - force-vlan-vc-forwarding
                        - hash-label
```

config service vpls bgp-evpn mpls ingress-replication-bum-label

```
                    - ingress-replication-bum-label
                    - mh-mode
                    - oper-group
                    - restrict-protected-src
                    - route-next-hop
                    - send-tunnel-encap
                    - shutdown
                    - split-horizon-group
                - sel-mcast-advertisement
                - unknown-mac-route
                - vlan-aware-bundle
            - bgp-vpls
                - max-ve-id
                - shutdown
                - ve-name
                    - ve-id
            - description
            - disable-aging
            - disable-learning
            - discard-unknown
            - endpoint
                - auto-learn-mac-protect
                - block-on-mesh-failure
                - description
                - ignore-standby-signaling
                - mac-pinning
                - max-nbr-mac-addr
                - mc-endpoint
                    - mc-ep-peer
                - restrict-protected-src
                - revert-time
                - static-mac
                - suppress-standby-signaling
            - fdb-table-high-wmark
            - fdb-table-low-wmark
            - fdb-table-size
            - igmp-snooping
                - query-interval
                - query-src-ip
                - report-src-ip
                - robust-count
                - shutdown
            - ignore-l2vpn-mtu-mismatch
            - interface
                - address
                - arp-timeout
                - description
                - hold-time
                    - down
                    - up
                - mac
                - shutdown
                - static-arp
            - local-age
            - mac-move
                - move-frequency
                - number-retries
                - primary-ports
                    - cumulative-factor
                    - sap
                    - spoke-sdp
                - retry-timeout
                - secondary-ports
                    - cumulative-factor
```

config service vpls mac-move secondary-ports sap

```
                    - sap
                    - spoke-sdp
                - shutdown
            - mac-protect
                - mac
            - mac-subnet-length
            - mcast-ipv6-snooping-scope
            - mcr-default-gtw
                - ip
                - mac
            - mesh-sdp
                - accounting-policy
                - adv-service-mtu
                - auto-learn-mac-protect
                - collect-stats
                - control-word
                - description
                - dhcp
                    - description
                    - snoop
                - egress
                    - filter
                    - mfib-allowed-mda-destinations
                        - mda
                    - qos
                    - vc-label
                - force-vlan-vc-forwarding
                - hash-label
                - igmp-snooping
                    - disable-router-alert-check
                    - fast-leave
                    - import
                    - last-member-query-interval
                    - max-num-groups
                    - max-num-grp-sources
                    - max-num-sources
                    - mrouter-port
                    - query-interval
                    - query-response-interval
                    - robust-count
                    - send-queries
                    - static
                        - group
                            - source
                            - starg
                    - version
                - ingress
                    - filter
                    - qos
                    - vc-label
                - mac-pinning
                - mld-snooping
                    - disable-router-alert-check
                    - fast-leave
                    - import
                    - last-member-query-interval
                    - max-num-groups
                    - mrouter-port
                    - query-interval
                    - query-response-interval
                    - robust-count
                    - send-queries
                    - static
                        - group
```

config service vpls mesh-sdp mld-snooping static group source

```
                        - source
                        - starg
                    - version
                - restrict-protected-src
                - shutdown
                - static-mac
                - vlan-vc-tag
            - mfib-table-high-wmark
            - mfib-table-low-wmark
            - mfib-table-size
            - mld-snooping
                - query-interval
                - query-src-ip
                - report-src-ip
                - robust-count
                - shutdown
            - propagate-mac-flush
            - proxy-arp
                - age-time
                - dup-detect
                - dynamic
                    - mac-list
                    - resolve
                    - sap
                - dynamic-arp-populate
                - evpn-route-tag
                - garp-flood-evpn
                - process-arp-probes
                - received-garp-flood
                - received-unknown-arp-request-flood
                - restrict-non-configured-ip-address
                - send-refresh
                - shutdown
                - static
                - table-size
                - unknown-arp-request-flood-evpn
            - proxy-nd
                - age-time
                - dup-detect
                - dynamic
                    - mac-list
                    - resolve
                    - sap
                - dynamic-nd-populate
                - evpn-nd-advertise
                - evpn-route-tag
                - host-unsolicited-na-flood-evpn
                - process-dad-neighbor-solicitations
                - received-host-unsolicited-na-flood
                - received-router-unsolicited-na-flood
                - received-unknown-ns-flood
                - restrict-non-configured-ip-address
                - router-unsolicited-na-flood-evpn
                - send-refresh
                - shutdown
                - static
                - table-size
                - unknown-ns-flood-evpn
            - remote-age
            - sap
                - accounting-policy
                - auto-learn-mac-protect
                - bandwidth
                - bgp-vpls-mh-ve-id
```

config service vpls sap bpdu-translation

```
                    - bpdu-translation
                    - collect-stats
                    - description
                    - dhcp
                        - description
                        - lease-populate
                        - option
                            - action
                            - circuit-id
                            - remote-id
                            - vendor-specific-option
                                - client-mac-address
                                - sap-id
                                - service-id
                                - string
                                - system-id
                        - proxy-server
                            - emulated-server
                            - lease-time
                            - shutdown
                        - shutdown
                        - snoop
                    - dhcp-user-db
                    - dhcp6
                        - description
                        - ldra
                            - interface-type
                            - options
                                - interface-id
                                - remote-id
                            - shutdown
                    - dhcp6-user-db
                    - disable-aging
                    - disable-learning
                    - discard-unknown-source
                    - dist-cpu-protection
                    - egress
                    - egress
                        - agg-rate
                            - queue-frame-based-accounting
                            - rate
                        - dest-mac-rewrite
                        - filter
                        - policer-control-override
                            - max-rate
                            - priority-mbs-thresholds
                                - min-thresh-separation
                                - priority
                                    - mbs-contribution
                        - policer-control-policy
                        - policer-override
                            - policer
                                - cbs
                                - mbs
                                - packet-byte-offset
                                - percent-rate
                                - rate
                                - stat-mode
                        - qinq-mark-top-only
                        - qos
                        - qos
                        - queue-override
                            - queue
                                - adaptation-rule
```

config service vpls sap egress queue-override queue avg-frame-overhead

```
                        - avg-frame-overhead
                        - burst-limit
                        - cbs
                        - drop-tail
                            - low
                                - percent-reduction-from-mbs
                        - mbs
                        - parent
                        - percent-rate
                        - rate
                    - scheduler-override
                        - scheduler
                            - parent
                            - rate
                    - scheduler-policy
                - force-l2pt-boundary
                - igmp-snooping
                    - disable-router-alert-check
                    - fast-leave
                    - import
                    - last-member-query-interval
                    - max-num-groups
                    - max-num-grp-sources
                    - max-num-sources
                    - mrouter-port
                    - query-interval
                    - query-response-interval
                    - robust-count
                    - send-queries
                    - static
                        - group
                            - source
                            - starg
                    - version
                - ingress
                    - aggregate-policer
                        - cbs
                        - mbs
                        - rate
                    - filter
                    - match-qinq-dot1p
                    - policer-control-override
                        - max-rate
                        - priority-mbs-thresholds
                            - min-thresh-separation
                            - priority
                                - mbs-contribution
                    - policer-control-policy
                    - policer-override
                        - policer
                            - cbs
                            - mbs
                            - packet-byte-offset
                            - percent-rate
                            - rate
                            - stat-mode
                    - qos
                    - queue-override
                        - queue
                            - adaptation-rule
                            - cbs
                            - drop-tail
                                - low
                                    - percent-reduction-from-mbs
```

config service vpls sap ingress queue-override queue mbs

```
                            - mbs
                            - parent
                            - percent-rate
                            - rate
                        - scheduler-override
                            - scheduler
                                - parent
                                - rate
                        - scheduler-policy
                    - l2pt-termination
                    - limit-mac-move
                    - mac-pinning
                    - managed-vlan-list
                        - default-sap
                        - range
                    - max-nbr-mac-addr
                    - mld-snooping
                        - disable-router-alert-check
                        - fast-leave
                        - import
                        - last-member-query-interval
                        - max-num-groups
                        - mrouter-port
                        - query-interval
                        - query-response-interval
                        - robust-count
                        - send-queries
                        - static
                            - group
                                - source
                                - starg
                        - version
                    - monitor-oper-group
                    - msap-defaults
                        - group-interface
                        - policy
                        - service
                    - multi-service-site
                    - oper-group
                    - process-cpm-traffic-on-sap-down
                    - restrict-protected-src
                    - restrict-unprotected-dst
                    - ring-node
                    - shutdown
                    - static-mac
                    - stp
                        - auto-edge
                        - edge-port
                        - link-type
                        - mst-instance
                            - mst-path-cost
                            - mst-port-priority
                        - path-cost
                        - port-num
                        - priority
                        - root-guard
                        - shutdown
                - selective-learned-fdb
                - send-flush-on-failure
                - service-mtu
                - shutdown
                - site
                    - boot-timer
                    - failed-threshold
```

config service vpls site mesh-sdp-binding

```
                - mesh-sdp-binding
                - monitor-oper-group
                - sap
                - shutdown
                - site-activation-timer
                - site-id
                - site-min-down-timer
                - split-horizon-group
                - spoke-sdp
            - split-horizon-group
                - auto-learn-mac-protect
                - description
                - restrict-protected-src
                - restrict-unprotected-dst
            - spoke-sdp
                - accounting-policy
                - adv-service-mtu
                - auto-learn-mac-protect
                - block-on-mesh-failure
                - bpdu-translation
                - collect-stats
                - control-channel-status
                    - acknowledgment
                    - refresh-timer
                    - request-timer
                    - request-timer
                    - shutdown
                - control-word
                - control-word
                - description
                - dhcp
                    - description
                    - snoop
                - disable-aging
                - disable-learning
                - discard-unknown-source
                - egress
                    - filter
                    - mfib-allowed-mda-destinations
                        - mda
                    - qos
                    - vc-label
                - force-vlan-vc-forwarding
                - hash-label
                - igmp-snooping
                    - disable-router-alert-check
                    - fast-leave
                    - import
                    - last-member-query-interval
                    - max-num-groups
                    - max-num-grp-sources
                    - max-num-sources
                    - mrouter-port
                    - query-interval
                    - query-response-interval
                    - robust-count
                    - send-queries
                    - static
                        - group
                            - source
                            - starg
                    - version
                - ignore-standby-signaling
                - ingress
```

config service vpls spoke-sdp ingress filter

```
                    - filter
                    - qos
                    - vc-label
                - l2pt-termination
                - limit-mac-move
                - mac-pinning
                - max-nbr-mac-addr
                - mld-snooping
                    - disable-router-alert-check
                    - fast-leave
                    - import
                    - last-member-query-interval
                    - max-num-groups
                    - mrouter-port
                    - query-interval
                    - query-response-interval
                    - robust-count
                    - send-queries
                    - static
                        - group
                            - source
                            - starg
                    - version
                - monitor-oper-group
                - oper-group
                - precedence
                - pw-status-signaling
                - restrict-protected-src
                - shutdown
                - shutdown
                - static-mac
                - stp
                    - auto-edge
                    - edge-port
                    - link-type
                    - path-cost
                    - port-num
                    - priority
                    - root-guard
                    - shutdown
                - vlan-vc-tag
            - static-mac
                - mac
            - stp
                - forward-delay
                - hello-time
                - hold-count
                - max-age
                - mode
                - mst-instance
                    - mst-priority
                    - vlan-range
                - mst-max-hops
                - mst-name
                - mst-revision
                - priority
                - shutdown
            - temp-flooding
            - vpls-group
                - mvrp-control
                - sap-template-binding
                - service-range
                - shutdown
                - vpls-template-binding
```

3.4.19.16 configure service vprn Commands

```
            - vprn
                - aaa
                    - remote-servers
                        - radius
                            - access-algorithm
                            - accounting
                            - accounting-port
                            - authorization
                            - interactive-authentication
                            - port
                            - retry
                            - server
                            - shutdown
                            - timeout
                            - use-default-template
                        - tacplus
                            - accounting
                            - authorization
                            - ignore-unknown-mandatory-vsas
                            - interactive-authentication
                            - priv-lvl-map
                                - priv-lvl
                            - request-format
                                - access-operation-cmd
                            - retry-timeout
                            - server
                            - service-request
                                - nokia-grpc-rpc-authorization
                                - nokia-netconf-base-op-authorization
                                - nokia-user
                                - nokia-user-profile
                            - shutdown
                            - timeout
                            - use-default-template
                - aggregate
                - allow-export-bgp-vpn
                - auto-bind-tunnel
                    - allow-flex-algo-fallback
                    - ecmp
                    - enforce-strict-tunnel-tagging
                    - resolution
                    - resolution
                    - resolution-filter
                        - bgp
                        - gre
                        - ldp
                        - rsvp
                        - rsvp
                        - sr-isis
                        - sr-ospf
                        - sr-ospf3
                        - sr-policy
                        - sr-te
                    - weighted-ecmp
                - autonomous-system
                - bgp
                    - advertise-inactive
                    - advertise-ipv6-next-hops
                    - aggregator-id-zero
                    - attribute-set
```

config service vprn bgp attribute-set remove

```
            - remove
        - auth-keychain
        - authentication-key
        - best-path-selection
            - always-compare-med
            - as-path-ignore
            - compare-origin-validation-state
            - deterministic-med
            - ebgp-ibgp-equal
            - ignore-nh-metric
            - ignore-router-id
            - origin-invalid-unusable
        - bfd-enable
        - bfd-strict-mode
            - advertise
            - next-hop-reachability
        - cluster
        - connect-retry
        - convergence
            - family
                - max-wait-to-advertise
            - min-wait-to-advertise
        - damp-peer-oscillations
        - damping
        - default-label-preference
        - default-preference
        - description
        - disable-4byte-asn
        - disable-client-reflect
        - disable-communities
        - disable-fast-external-failover
        - dynamic-neighbor-limit
        - ebgp-default-reject-policy
        - eibgp-loadbalance
        - enable-peer-tracking
        - enforce-first-as
        - error-handling
            - legacy-mode
            - update-fault-tolerance
        - export
        - extended-nh-encoding
        - family
        - graceful-restart
            - enable-notification
            - long-lived
                - advertise-stale-to-all-neighbors
                - advertised-stale-time
                - family
                    - advertised-stale-time
                    - helper-override-stale-time
                - forwarding-bits-set
                - helper-override-restart-time
                - helper-override-stale-time
            - restart-time
            - stale-routes-time
        - group
            - advertise-inactive
            - advertise-ipv6-next-hops
            - aggregator-id-zero
            - as-override
            - auth-keychain
            - authentication-key
            - bfd-enable
            - bfd-strict-mode
```

configure service vprn bgp group bfd-strict-mode advertise

```
                        - advertise
                        - next-hop-reachability
                    - cluster
                    - connect-retry
                    - damp-peer-oscillations
                    - damping
                    - default-label-preference
                    - default-preference
                    - description
                    - disable-4byte-asn
                    - disable-capability-negotiation
                    - disable-client-reflect
                    - disable-communities
                    - disable-fast-external-failover
                    - dynamic-neighbor
                        - interface
                            - allowed-peer-as
                            - max-sessions
                        - match
                            - prefix
                                - allowed-peer-as
                    - dynamic-neighbor-limit
                    - ebgp-default-reject-policy
                    - enable-origin-validation
                    - enable-peer-tracking
                    - enforce-first-as
                    - error-handling
                        - update-fault-tolerance
                    - evpn-link-bandwidth
                        - add-to-received-bgp
                    - export
                    - extended-nh-encoding
                    - family
                    - graceful-restart
                        - enable-notification
                        - long-lived
                            - advertise-stale-to-all-neighbors
                            - advertised-stale-time
                            - family
                                - advertised-stale-time
                                - helper-override-stale-time
                            - forwarding-bits-set
                            - helper-override-restart-time
                            - helper-override-stale-time
                        - restart-time
                        - stale-routes-time
                    - hold-time
                    - import
                    - initial-send-delay-zero
                    - keepalive
                    - label-preference
                    - link-bandwidth
                        - accept-from-ebgp
                        - add-to-received-ebgp
                        - aggregate-used-paths
                        - send-to-ebgp
                    - local-address
                    - local-as
                    - local-preference
                    - loop-detect
                    - loop-detect-threshold
                    - med-out
                    - min-route-advertisement
                    - multihop
```

config service vprn bgp group multipath-eligible

```
                        - multipath-eligible
                        - neighbor
                            - advertise-inactive
                            - advertise-ipv6-next-hops
                            - aggregator-id-zero
                            - as-override
                            - auth-keychain
                            - authentication-key
                            - bfd-enable
                            - bfd-strict-mode
                                - advertise
                                - next-hop-reachability
                            - cluster
                            - connect-retry
                            - damp-peer-oscillations
                            - damping
                            - default-label-preference
                            - default-preference
                            - description
                            - disable-4byte-asn
                            - disable-capability-negotiation
                            - disable-client-reflect
                            - disable-communities
                            - disable-fast-external-failover
                            - ebgp-default-reject-policy
                            - enable-origin-validation
                            - enable-peer-tracking
                            - enforce-first-as
                            - error-handling
                                - update-fault-tolerance
                            - evpn-link-bandwidth
                                - add-to-received-bgp
                            - export
                            - extended-nh-encoding
                            - family
                            - graceful-restart
                                - enable-notification
                                - long-lived
                                    - advertise-stale-to-all-neighbors
                                    - advertised-stale-time
                                    - family
                                        - advertised-stale-time
                                        - helper-override-stale-time
                                    - forwarding-bits-set
                                    - helper-override-restart-time
                                    - helper-override-stale-time
                                - restart-time
                                - stale-routes-time
                            - hold-time
                            - import
                            - initial-send-delay-zero
                            - keepalive
                            - label-preference
                            - link-bandwidth
                                - accept-from-ebgp
                                - add-to-received-ebgp
                                - aggregate-used-paths
                                - send-to-ebgp
                            - local-address
                            - local-as
                            - local-preference
                            - loop-detect
                            - loop-detect-threshold
                            - med-out
```

config service vprn bgp group neighbor min-route-advertisement

```
                        - min-route-advertisement
                        - multihop
                        - multipath-eligible
                        - next-hop-self
                        - passive
                        - path-mtu-discovery
                        - peer-as
                        - preference
                        - prefix-limit
                        - remove-private
                        - send-default
                        - shutdown
                        - split-horizon
                        - tcp-mss
                        - third-party-nexthop
                        - ttl-security
                        - type
                    - next-hop-self
                    - passive
                    - path-mtu-discovery
                    - peer-as
                    - preference
                    - prefix-limit
                    - remove-private
                    - send-default
                    - shutdown
                    - split-horizon
                    - tcp-mss
                    - third-party-nexthop
                    - ttl-security
                    - type
                - hold-time
                - ibgp-multipath
                - import
                - initial-send-delay-zero
                - keepalive
                - label-preference
                - local-as
                - local-preference
                - loop-detect
                - loop-detect-threshold
                - med-out
                - min-route-advertisement
                - multi-path
                    - ipv4
                    - ipv6
                    - label-ipv4
                    - label-ipv6
                    - maximum-paths
                - multihop
                - next-hop-resolution
                    - policy
                    - use-bgp-routes
                    - use-leaked-routes
                        - static
                - path-mtu-discovery
                - peer-tracking-policy
                - preference
                - rapid-withdrawal
                - remove-private
                - rib-management
                    - ipv4
                        - leak-import
                        - route-table-import
```

configure service vprn bgp rib-management ipv6

```
                    - ipv6
                        - leak-import
                        - route-table-import
                    - label-ipv4
                        - leak-import
                        - route-table-import
                    - label-ipv6
                        - leak-import
                - router-id
                - send-default
                - shutdown
                - split-horizon
                - tcp-mss
                - third-party-nexthop
            - bgp-evpn
                - mpls
                    - auto-bind-tunnel
                        - allow-flex-algo-fallback
                        - ecmp
                        - enforce-strict-tunnel-tagging
                        - enforce-untagged-route
                        - resolution
                        - resolution-filter
                            - bgp
                            - ldp
                            - rsvp
                            - sr-isis
                            - sr-ospf
                            - sr-ospf3
                            - sr-policy
                            - sr-te
                    - default-route-tag
                    - domain-id
                    - dynamic-egress-label-limit
                    - evi
                    - evpn-link-bandwidth
                        - advertise
                        - weighted-ecmp
                    - route-distinguisher
                    - send-tunnel-encap
                    - shutdown
                    - vrf-export
                    - vrf-import
                    - vrf-target
            - bgp-ipvpn
                - attribute-set
                    - export
                    - import
                - mpls
                    - auto-bind-tunnel
                        - allow-flex-algo-fallback
                        - ecmp
                        - enforce-strict-tunnel-tagging
                        - enforce-untagged-route
                        - resolution
                        - resolution-filter
                            - bgp
                            - gre
                            - ldp
                            - rsvp
                            - sr-isis
                            - sr-ospf
                            - sr-ospf3
                            - sr-policy
```

config service vprn bgp-ipvpn mpls auto-bind-tunnel resolution-filter sr-te

```
                        - sr-te
                    - static-blackhole-first
                    - weighted-ecmp
                - domain-id
                - dynamic-egress-label-limit
                - route-distinguisher
                - shutdown
                - vrf-export
                - vrf-import
                - vrf-target
            - bgp-shared-queue
            - carrier-carrier-vpn
            - confederation
            - d-path-length-ignore
            - description
            - dhcp
                - local-dhcp-server
                    - description
                    - failover
                        - ignore-mclt-on-takeover
                        - maximum-client-lead-time
                        - partner-down-delay
                        - peer
                        - shutdown
                        - startup-wait-time
                    - force-renews
                    - lease-hold-time
                    - lease-hold-time-for
                        - internal-lease-ipsec
                        - solicited-release
                    - pool
                        - description
                        - failover
                            - ignore-mclt-on-takeover
                            - maximum-client-lead-time
                            - partner-down-delay
                            - peer
                            - shutdown
                            - startup-wait-time
                        - max-lease-time
                        - min-lease-time
                        - minimum-free
                        - nak-non-matching-subnet
                        - offer-time
                        - options
                            - custom-option
                            - dns-server
                            - domain-name
                            - lease-rebind-time
                            - lease-renew-time
                            - lease-time
                            - netbios-name-server
                            - netbios-node-type
                        - subnet
                            - address-range
                            - drain
                            - exclude-addresses
                            - maximum-declined
                            - minimum-free
                            - options
                                - custom-option
                                - default-router
                                - subnet-mask
                    - shutdown
```

config service vprn dhcp server use-gi-address

```
                - use-gi-address
                - use-pool-from-client
                - user-db
                - user-ident
            - dhcp6
                - local-dhcp-server
                    - allow-lease-query
                    - description
                    - failover
                        - ignore-mclt-on-takeover
                        - maximum-client-lead-time
                        - partner-down-delay
                        - peer
                        - shutdown
                        - startup-wait-time
                    - ignore-rapid-commit
                    - interface-id-mapping
                    - lease-hold-time
                    - lease-hold-time-for
                        - internal-lease-ipsec
                        - solicited-release
                    - pool
                        - delegated-prefix-length
                        - description
                        - exclude-prefix
                        - failover
                            - ignore-mclt-on-takeover
                            - maximum-client-lead-time
                            - partner-down-delay
                            - peer
                            - shutdown
                            - startup-wait-time
                        - options
                            - custom-option
                            - dns-server
                            - domain-name
                        - prefix
                            - drain
                            - options
                                - custom-option
                                - dns-server
                                - domain-name
                            - preferred-lifetime
                            - rebind-timer
                            - renew-timer
                            - thresholds
                                - minimum-free
                                    - depleted-event
                                    - minimum
                            - valid-lifetime
                        - thresholds
                            - minimum-free
                                - depleted-event
                                - minimum
                    - server-id
                    - shutdown
                    - use-link-address
                    - use-pool-from-client
                    - user-db
                    - user-ident
            - dns
                - default-domain
                - ipv4-source-address
                - ipv6-source-address
```

config service vprn dns primary-dns

```
            - primary-dns
            - secondary-dns
            - shutdown
            - tertiary-dns
        - ecmp
        - ecmp-unequal-cost
        - encryption-keygroup
        - export-inactive-bgp
        - export-inactive-bgp-enhanced
        - fib-priority
        - grt-lookup
            - enable-grt
                - allow-local-management
            - export-grt
            - export-limit
            - export-v6-limit
            - import-grt
        - hash-label
        - igmp
            - grp-if-query-src-ip
            - interface
                - disable-router-alert-check
                - import
                - max-groups
                - max-grp-sources
                - max-sources
                - query-interval
                - query-last-member-interval
                - query-response-interval
                - redundant-multicast
                - shutdown
                - ssm-translate
                    - grp-range
                        - source
                - static
                    - group
                        - source
                        - starg
                - subnet-check
                - version
            - query-interval
            - query-last-member-interval
            - query-response-interval
            - robust-count
            - shutdown
            - ssm-translate
                - grp-range
                    - source
        - ignore-nh-metric
        - interface
            - address
            - allow-directed-broadcasts
            - arp-host-route
                - populate
            - arp-learn-unsolicited
            - arp-limit
            - arp-populate
            - arp-proactive-refresh
            - arp-retry-timer
            - arp-timeout
            - bfd
            - description
            - dhcp
                - description
```

config service vprn if dhcp gi-address

```
                    - gi-address
                    - lease-populate
                    - option
                        - action
                        - circuit-id
                        - remote-id
                        - vendor-specific-option
                            - client-mac-address
                            - pool-name
                            - sap-id
                            - service-id
                            - string
                            - system-id
                    - proxy-server
                        - emulated-server
                        - lease-time
                        - shutdown
                    - relay-plain-bootp
                    - relay-proxy
                    - server
                    - shutdown
                    - trusted
                    - use-arp
                - dynamic-tunnel-redundant-next-hop
                - enable-mac-accounting
                - hold-time
                    - down
                    - up
                - icmp
                    - mask-reply
                    - param-problem
                    - redirects
                    - ttl-expired
                    - unreachables
                - if-attribute
                    - admin-group
                    - srlg-group
                - ingress
                - ip-helper-address
                - ip-mtu
                - ipsec
                    - ip-exception
                    - ipsec-tunnel
                        - bfd-designate
                        - bfd-enable
                        - clear-df-bit
                        - copy-traffic-class-upon-decapsulation
                        - description
                        - dynamic-keying
                            - auto-establish
                            - cert
                                - cert-profile
                                - status-verify
                                    - default-result
                                    - primary
                                - trust-anchor-profile
                            - ike-policy
                            - local-id
                            - ppk
                            - pre-shared-key
                            - transform
                        - encapsulated-ip-mtu
                        - icmp-generation
                            - frag-required
```

config service vprn if ipsec ipsec-tunnel icmp-generation frag-required interval

```
                                - interval
                                - message-count
                            - icmp6-generation
                                - pkt-too-big
                                - interval
                                - message-count
                            - ip-mtu
                            - local-gateway-address
                            - manual-keying
                                - security-association
                            - max-history-esp-key-records
                            - max-history-ike-key-records
                            - pmtu-discovery-aging
                            - private-tcp-mss-adjust
                            - propagate-pmtu-v4
                            - propagate-pmtu-v6
                            - public-tcp-mss-adjust
                            - remote-gateway-address
                            - replay-window
                            - security-policy
                            - shutdown
                        - ipv6-exception
                        - shutdown
                    - ipv6
                        - address
                        - bfd
                        - dad-disable
                        - dhcp6-relay
                            - description
                            - lease-populate
                            - option
                                - interface-id
                                - remote-id
                            - server
                            - shutdown
                            - source-address
                        - forward-ipv4-packets
                        - icmp6
                            - packet-too-big
                            - param-problem
                            - redirects
                            - time-exceeded
                            - unreachables
                        - link-local-address
                        - local-dhcp-server
                        - local-proxy-nd
                        - nd-host-route
                            - populate
                        - nd-learn-unsolicited
                        - nd-proactive-refresh
                        - neighbor
                        - neighbor-limit
                        - proxy-nd-policy
                        - reachable-time
                        - secure-nd
                            - allow-unsecured-msgs
                            - link-local-modifier
                            - public-key-min-bits
                            - security-parameter
                            - shutdown
                        - stale-time
                        - tcp-mss
                        - urpf-check
                            - ignore-default
```

config service vprn if ipv6 urpf-check mode

```
                        - mode
                    - vrrp
                        - backup
                        - bfd-enable
                        - init-delay
                        - mac
                        - master-int-inherit
                        - message-interval
                        - ntp-reply
                        - oper-group
                        - ping-reply
                        - policy
                        - preempt
                        - priority
                        - shutdown
                        - standby-forwarding
                        - telnet-reply
                        - traceroute-reply
                - load-balancing
                    - egr-ip-load-balancing
                - local-dhcp-server
                - local-proxy-arp
                - loopback
                - mac
                - monitor-oper-group
                - multi-chassis-shunting-profile
                - proxy-arp-policy
                - remote-proxy-arp
                - sap
                    - accounting-policy
                    - bandwidth
                    - collect-stats
                    - description
                    - dist-cpu-protection
                    - egress
                        - agg-rate
                            - queue-frame-based-accounting
                            - rate
                        - filter
                        - policer-control-override
                            - max-rate
                            - priority-mbs-thresholds
                                - min-thresh-separation
                                - priority
                                    - mbs-contribution
                        - policer-control-policy
                        - policer-override
                            - policer
                                - cbs
                                - mbs
                                - packet-byte-offset
                                - percent-rate
                                - rate
                                - stat-mode
                        - qinq-mark-top-only
                        - qos
                        - queue-override
                            - queue
                                - adaptation-rule
                                - avg-frame-overhead
                                - burst-limit
                                - cbs
                                - drop-tail
                                    - low
```

config service vprn if sap egress queue-override queue mbs

```
                                        - percent-reduction-from-mbs
                                - mbs
                                - mbs
                                - parent
                                - percent-rate
                                - rate
                        - scheduler-override
                            - scheduler
                                - parent
                                - rate
                        - scheduler-policy
                    - ingress
                        - aggregate-policer
                            - cbs
                            - mbs
                            - rate
                        - filter
                        - match-qinq-dot1p
                        - policer-control-override
                            - max-rate
                            - priority-mbs-thresholds
                                - min-thresh-separation
                                - priority
                                    - mbs-contribution
                        - policer-control-policy
                        - policer-override
                            - policer
                                - cbs
                                - mbs
                                - packet-byte-offset
                                - percent-rate
                                - rate
                                - stat-mode
                        - qos
                        - queue-override
                            - queue
                                - adaptation-rule
                                - cbs
                                - drop-tail
                                    - low
                                        - percent-reduction-from-mbs
                                - mbs
                                - mbs
                                - parent
                                - percent-rate
                                - rate
                        - scheduler-override
                            - scheduler
                                - parent
                                - rate
                        - scheduler-policy
                    - ip-tunnel
                        - backup-remote-ip
                        - clear-df-bit
                        - delivery-service
                        - description
                        - description
                        - dest-ip
                        - dscp
                        - encapsulated-ip-mtu
                        - encapsulated-ip-mtu
                        - gre-header
                        - icmp-generation
                            - frag-required
```

config service vprn if sap ip-tunnel icmp-generation frag-required interval

```
                        - interval
                        - message-count
                    - icmp6-generation
                        - packet-too-big
                    - ip-mtu
                    - ipsec-transport-mode-profile
                    - pmtu-discovery-aging
                    - private-tcp-mss-adjust
                    - propagate-pmtu-v4
                    - propagate-pmtu-v6
                    - public-tcp-mss-adjust
                    - reassembly
                    - remote-ip
                    - shutdown
                    - source
                - ipsec-gw
                    - cert
                        - cert-profile
                        - status-verify
                            - default-result
                            - primary
                        - trust-anchor-profile
                    - client-db
                    - default-secure-service
                    - default-tunnel-template
                    - dhcp
                        - gi-address
                        - send-release
                        - server
                        - shutdown
                    - dhcp6
                        - link-address
                        - send-release
                        - server
                        - shutdown
                    - ike-policy
                    - local-address-assignment
                        - ipv4
                            - address-source
                        - ipv6
                            - address-source
                        - shutdown
                    - local-gateway-address
                    - local-id
                    - max-history-esp-key-records
                    - max-history-ike-key-records
                    - pre-shared-key
                    - radius-accounting-policy
                    - radius-authentication-policy
                    - shutdown
                    - ts-negotiation
                - ipsec-tunnel
                    - bfd-designate
                    - bfd-enable
                    - clear-df-bit
                    - copy-traffic-class-upon-decapsulation
                    - description
                    - dest-ip
                    - dynamic-keying
                        - auto-establish
                        - cert
                            - cert-profile
                            - status-verify
                                - default-result
```

configure service vprn interface sap ipsec-tunnel dynamic-keying cert trust-anchor-profile

```
                                    - primary
                                - trust-anchor-profile
                            - ike-policy
                            - local-id
                            - ppk
                            - pre-shared-key
                            - transform
                        - encapsulated-ip-mtu
                        - icmp-generation
                            - frag-required
                                - interval
                                - message-count
                        - icmp6-generation
                            - pkt-too-big
                                - interval
                                - message-count
                        - ip-mtu
                        - local-gateway-address
                        - manual-keying
                            - security-association
                        - max-history-esp-key-records
                        - max-history-ike-key-records
                        - pmtu-discovery-aging
                        - private-tcp-mss-adjust
                        - propagate-pmtu-v4
                        - propagate-pmtu-v6
                        - public-tcp-mss-adjust
                        - replay-window
                        - security-policy
                        - shutdown
                    - multi-service-site
                    - shutdown
                - secondary
                - shutdown
                - spoke-sdp
                    - accounting-policy
                    - collect-stats
                    - control-channel-status
                        - acknowledgment
                        - refresh-timer
                        - request-timer
                        - shutdown
                    - description
                    - egress
                        - filter
                        - qos
                        - vc-label
                    - hash-label
                    - ingress
                        - filter
                        - qos
                        - vc-label
                    - shutdown
                - static-arp
                - static-tunnel-redundant-next-hop
                - tcp-mss
                - tos-marking-state
                - unnumbered
                - unnumbered
                - urpf-check
                    - ignore-default
                    - mode
                - vpls
                    - egress
```

config service vprn if vpls egress reclassify-using-qos

```
                        - reclassify-using-qos
                        - v4-routed-override-filter
                        - v6-routed-override-filter
                    - evpn
                        - arp
                            - advertise
                            - flood-garp-and-unknown-req
                            - learn-dynamic
                        - nd
                            - advertise
                            - learn-dynamic
                    - evpn-tunnel
                    - ingress
                        - v4-routed-override-filter
                        - v6-routed-override-filter
                - vrrp
                - vrrp
                    - authentication-key
                    - backup
                    - bfd-enable
                    - init-delay
                    - mac
                    - master-int-inherit
                    - message-interval
                    - ntp-reply
                    - oper-group
                    - ping-reply
                    - policy
                    - preempt
                    - priority
                    - shutdown
                    - ssh-reply
                    - standby-forwarding
                    - telnet-reply
                    - traceroute-reply
            - ip-mirror-interface
                - description
                - shutdown
                - spoke-sdp
                    - description
                    - ingress
                        - filter
                        - vc-label
                    - shutdown
            - ipsec
                - allow-reverse-route-override
                - multi-chassis-shunt-interface
                    - next-hop
                - multi-chassis-shunting-profile
                    - peer
                        - multi-chassis-shunt-interface
                - overlapping-reverse-route
                - security-policy
                    - entry
                        - local-ip
                        - local-v6-ip
                        - remote-ip
                        - remote-v6-ip
            - ipv6
                - reachable-time
                - stale-time
            - isis
                - advertise-passive-only
                - advertise-router-capability
```

config service vprn isis all-l1isis

```
                    - all-l1isis
                    - all-l2isis
                    - area-id
                    - auth-keychain
                    - authentication-check
                    - authentication-key
                    - authentication-type
                    - csnp-authentication
                    - csnp-on-p2p
                    - default-route-tag
                    - export
                    - export-limit
                    - graceful-restart
                        - helper-disable
                    - hello-authentication
                    - hello-padding
                    - ignore-attached-bit
                    - ignore-lsp-errors
                    - ignore-narrow-metric
                    - iid-tlv-enable
                    - import
                    - interface
                        - bfd-enable
                        - csnp-interval
                        - default-instance
                        - hello-auth-keychain
                        - hello-authentication
                        - hello-authentication-key
                        - hello-authentication-type
                        - hello-padding
                        - interface-type
                        - ipv4-multicast-disable
                        - ipv6-unicast-disable
                        - level
                            - hello-auth-keychain
                            - hello-authentication-key
                            - hello-authentication-type
                            - hello-interval
                            - hello-multiplier
                            - hello-padding
                            - ipv4-multicast-metric
                            - ipv6-unicast-metric
                            - metric
                            - passive
                            - priority
                            - sd-offset
                            - sf-offset
                        - level-capability
                        - lfa-policy-map
                        - load-balancing-weight
                        - loopfree-alternate-exclude
                        - lsp-pacing-interval
                        - mesh-group
                        - passive
                        - retransmit-interval
                        - shutdown
                        - tag
                    - ipv4-multicast-routing
                    - ipv4-routing
                    - ipv6-routing
                    - level
                        - advertise-router-capability
                        - auth-keychain
                        - authentication-key
```

config service vprn isis level authentication-type

```
                    - authentication-type
                    - csnp-authentication
                    - default-ipv4-multicast-metric
                    - default-ipv6-unicast-metric
                    - default-metric
                    - external-preference
                    - hello-authentication
                    - hello-padding
                    - loopfree-alternate-exclude
                    - lsp-mtu-size
                    - preference
                    - psnp-authentication
                    - wide-metrics-only
                - level-capability
                - link-group
                    - description
                    - level
                        - ipv4-multicast-metric-offset
                        - ipv4-unicast-metric-offset
                        - ipv6-unicast-metric-offset
                        - member
                        - oper-members
                        - revert-members
                - loopfree-alternates
                    - exclude
                        - prefix-policy
                - lsp-lifetime
                - lsp-minimum-remaining-lifetime
                - lsp-mtu-size
                - lsp-refresh-interval
                - multi-topology
                    - ipv4-multicast
                    - ipv6-unicast
                - multicast-import
                - overload
                - overload-export-external
                - overload-export-interlevel
                - overload-fib-error-notify-only
                - overload-on-boot
                - poi-tlv-enable
                - prefix-attributes-tlv
                - prefix-limit
                - psnp-authentication
                - reference-bandwidth
                - rib-priority
                - router-id
                - shutdown
                - standard-multi-instance
                - strict-adjacency-check
                - summary-address
                - suppress-attached-bit
                - system-id
                - timers
                    - lsp-wait
                    - spf-wait
                - unicast-import-disable
            - label-mode
            - local-routes-domain-id
            - log
                - filter
                    - default-action
                    - description
                    - entry
                        - action
```

configure service vprn log filter entry description

```
                        - description
                        - match
                            - application
                            - message
                            - number
                            - severity
                            - subject
                    - log-id
                        - description
                        - filter
                        - from
                        - netconf-stream
                        - shutdown
                        - time-format
                        - to
                    - snmp-trap-group
                        - description
                        - trap-target
                    - syslog
                        - address
                        - description
                        - facility
                        - hostname
                        - level
                        - log-prefix
                        - port
                        - timestamp-format
                        - tls-client-profile
                - management
                    - allow-ftp
                    - allow-grpc
                    - allow-netconf
                    - allow-ssh
                    - allow-telnet
                    - allow-telnet6
                - maximum-ipv6-routes
                - maximum-routes
                - mc-maximum-routes
                - mld
                    - grp-if-query-src-ip
                    - interface
                        - disable-router-alert-check
                        - import
                        - max-groups
                        - max-grp-sources
                        - max-sources
                        - query-interval
                        - query-last-listener-interval
                        - query-response-interval
                        - shutdown
                        - ssm-translate
                            - grp-range
                                - source
                        - static
                            - group
                                - source
                                - starg
                        - version
                    - query-interval
                    - query-last-listener-interval
                    - query-response-interval
                    - robust-count
                    - shutdown
                    - ssm-translate
```

config service vprn mld ssm-translate grp-range

```
                    - grp-range
                        - source
                - mss-adjust-group
                - nat
                    - inside
                        - classic-lsn-max-subscriber-limit
                        - destination-prefix
                        - deterministic
                            - address-map
                                - outside-range
                                - shutdown
                            - prefix-map
                                - map
                                - shutdown
                        - nat-policy
                    - outside
                        - downstream-ip-filter
                        - mtu
                        - pool
                            - address-range
                                - description
                                - drain
                            - description
                            - icmp-echo-reply
                            - mode
                            - port-forwarding-dyn-block-reservation
                            - port-forwarding-range
                            - port-reservation
                            - shutdown
                            - subscriber-limit
                            - watermarks
                        - upstream-ip-filter
                - network
                    - ingress
                        - filter
                        - qos
                        - urpf-check
                - network-interface
                    - address
                    - allow-directed-broadcasts
                    - arp-retry-timer
                    - arp-timeout
                    - bfd
                    - description
                    - dist-cpu-protection
                    - egress
                    - egress
                        - filter
                    - hold-time
                        - down
                        - up
                    - icmp
                        - mask-reply
                        - param-problem
                        - redirects
                        - ttl-expired
                        - unreachables
                    - ingress
                        - filter
                    - ip-mtu
                    - lag
                    - load-balancing
                        - egr-ip-load-balancing
                        - lsr-load-balancing
```

config service vprn nw-if loopback

```
                - loopback
                - mac
                - port
                - qos
                - secondary
                - shutdown
                - static-arp
                - tcp-mss
                - tos-marking-state
                - urpf-check
                    - ignore-default
                    - mode
            - ntp
                - authenticate
                - authentication-check
                - authentication-key
                - authentication-keychain
                - broadcast
                - shutdown
            - ospf
                - advertise-router-capability
                - area
                    - advertise-ne-profile
                    - advertise-router-capability
                    - area-range
                    - blackhole-aggregate
                    - export
                    - import
                    - interface
                        - advertise-router-capability
                        - advertise-subnet
                        - auth-keychain
                        - authentication-key
                        - authentication-type
                        - bfd-enable
                        - dead-interval
                        - hello-interval
                        - interface-type
                        - lfa-policy-map
                        - load-balancing-weight
                        - loopfree-alternate-exclude
                        - lsa-filter-out
                        - message-digest-key
                        - metric
                        - mtu
                        - neighbor
                        - passive
                        - poll-interval
                        - priority
                        - retransmit-interval
                        - rib-priority
                        - shutdown
                        - transit-delay
                    - loopfree-alternate-exclude
                    - nssa
                        - area-range
                        - originate-default-route
                        - redistribute-external
                        - summaries
                    - sham-link
                        - auth-keychain
                        - authentication-key
                        - authentication-type
                        - dead-interval
```

config service vprn ospf area sham-link hello-interval

```
                        - hello-interval
                        - message-digest-key
                        - metric
                        - retransmit-interval
                        - shutdown
                        - transit-delay
                    - stub
                        - default-metric
                        - summaries
                    - virtual-link
                        - auth-keychain
                        - authentication-key
                        - authentication-type
                        - dead-interval
                        - hello-interval
                        - message-digest-key
                        - retransmit-interval
                        - shutdown
                        - transit-delay
                - compatible-rfc1583
                - export
                - export-limit
                - external-db-overflow
                - external-preference
                - graceful-restart
                    - helper-disable
                    - strict-lsa-checking
                - ignore-dn-bit
                - import
                - loopfree-alternates
                    - exclude
                        - prefix-policy
                - multicast-import
                - overload
                - overload-include-ext-1
                - overload-include-ext-2
                - overload-include-stub
                - overload-on-boot
                - preference
                - reference-bandwidth
                - rib-priority
                - router-id
                - rtr-adv-lsa-limit
                - shutdown
                - super-backbone
                - suppress-dn-bit
                - timers
                    - incremental-spf-wait
                    - lsa-accumulate
                    - lsa-arrival
                    - lsa-generate
                    - redistribute-delay
                    - spf-wait
                - unicast-import-disable
                - vpn-domain
                - vpn-tag
            - ospf3
                - advertise-router-capability
                - area
                    - advertise-router-capability
                    - area-range
                    - blackhole-aggregate
                    - export
                    - import
```

config service vprn ospf3 area interface

```
                    - interface
                        - advertise-router-capability
                        - authentication
                        - bfd-enable
                        - dead-interval
                        - hello-interval
                        - interface-type
                        - lfa-policy-map
                        - load-balancing-weight
                        - loopfree-alternate-exclude
                        - lsa-filter-out
                        - metric
                        - mtu
                        - neighbor
                        - passive
                        - poll-interval
                        - priority
                        - retransmit-interval
                        - rib-priority
                        - shutdown
                        - transit-delay
                    - key-rollover-interval
                    - loopfree-alternate-exclude
                    - nssa
                        - area-range
                        - originate-default-route
                        - redistribute-external
                        - summaries
                    - stub
                        - default-metric
                        - summaries
                    - virtual-link
                        - authentication
                        - dead-interval
                        - hello-interval
                        - retransmit-interval
                        - shutdown
                        - transit-delay
                - export
                - export-limit
                - external-db-overflow
                - external-preference
                - graceful-restart
                    - helper-disable
                    - strict-lsa-checking
                - ignore-dn-bit
                - import
                - loopfree-alternates
                    - exclude
                        - prefix-policy
                - multicast-import
                - overload
                - overload-include-ext-1
                - overload-include-ext-2
                - overload-include-stub
                - overload-on-boot
                - preference
                - reference-bandwidth
                - rib-priority
                - router-id
                - rtr-adv-lsa-limit
                - shutdown
                - suppress-dn-bit
                - timers
```

configure service vprn ospf3 timers incremental-spf-wait

```
                    - incremental-spf-wait
                    - lsa-accumulate
                    - lsa-arrival
                    - lsa-generate
                    - redistribute-delay
                    - spf-wait
                - unicast-import-disable
            - pim
                - apply-bgp-nh-override
                - apply-to
                - import
                - interface
                    - assert-period
                    - bfd-enable
                    - bsm-check-rtr-alert
                    - hello-interval
                    - hello-multiplier
                    - improved-assert
                    - instant-prune-echo
                    - ipv4-multicast-disable
                    - ipv6-multicast-disable
                    - max-groups
                    - monitor-oper-group
                    - multicast-senders
                    - priority
                    - shutdown
                    - sticky-dr
                    - three-way-hello
                    - tracking-support
                - ipv4-multicast-disable
                - ipv6-multicast-disable
                - mtu-over-head
                - non-dr-attract-traffic
                - rp
                    - anycast
                        - rp-set-peer
                    - auto-rp-discovery
                    - bootstrap-export
                    - bootstrap-import
                    - bsr-candidate
                        - address
                        - hash-mask-len
                        - priority
                        - shutdown
                    - ipv6
                        - anycast
                            - rp-set-peer
                        - bsr-candidate
                            - address
                            - hash-mask-len
                            - priority
                            - shutdown
                        - embedded-rp
                            - group-range
                            - shutdown
                        - rp-candidate
                            - address
                            - group-range
                            - holdtime
                            - priority
                            - shutdown
                        - static
                            - address
                                - group-prefix
```

configure service vprn pim rp ipv6 static address override

```
                                - override
                        - rp-candidate
                            - address
                            - group-range
                            - holdtime
                            - priority
                            - shutdown
                        - static
                            - address
                                - group-prefix
                                - override
                    - rpf-table
                    - rpf6-table
                    - shutdown
                    - source-address
                        - register-message
                    - spt-switchover-threshold
                    - ssm-assert-compatible-mode
                    - ssm-default-range-disable
                    - ssm-groups
                        - group-range
                - radius-server
                    - server
                        - accept-coa
                        - acct-port
                        - auth-port
                        - coa-script-policy
                        - description
                        - pending-requests-limit
                - reassembly-group
                - rip
                    - authentication-key
                    - authentication-type
                    - bfd-enable
                    - check-zero
                    - description
                    - export
                    - export-limit
                    - group
                        - authentication-key
                        - authentication-type
                        - bfd-enable
                        - check-zero
                        - description
                        - export
                        - import
                        - message-size
                        - metric-in
                        - metric-out
                        - neighbor
                            - authentication-key
                            - authentication-type
                            - bfd-enable
                            - check-zero
                            - description
                            - export
                            - import
                            - message-size
                            - metric-in
                            - metric-out
                            - preference
                            - receive
                            - send
                            - shutdown
```

config service vprn rip group neighbor split-horizon

```
                        - split-horizon
                        - timers
                        - unicast-address
                    - preference
                    - receive
                    - send
                    - shutdown
                    - split-horizon
                    - timers
                - import
                - message-size
                - metric-in
                - metric-out
                - preference
                - propagate-metric
                - receive
                - send
                - shutdown
                - split-horizon
                - timers
            - ripng
                - bfd-enable
                - check-zero
                - description
                - export
                - export-limit
                - group
                    - bfd-enable
                    - check-zero
                    - description
                    - export
                    - import
                    - message-size
                    - metric-in
                    - metric-out
                    - neighbor
                        - bfd-enable
                        - check-zero
                        - description
                        - export
                        - import
                        - message-size
                        - metric-in
                        - metric-out
                        - preference
                        - receive
                        - send
                        - shutdown
                        - split-horizon
                        - timers
                        - unicast-address
                    - preference
                    - receive
                    - send
                    - shutdown
                    - split-horizon
                    - timers
                - import
                - message-size
                - metric-in
                - metric-out
                - preference
                - receive
                - send
```

configure service vprn ripng shutdown

```
                    - shutdown
                    - split-horizon
                    - timers
                - route-distinguisher
                - router-advertisement
                    - dns-options
                        - rdnss-lifetime
                        - server
                    - interface
                        - current-hop-limit
                        - dns-options
                            - include-dns
                            - rdnss-lifetime
                            - server
                        - managed-configuration
                        - max-advertisement-interval
                        - min-advertisement-interval
                        - mtu
                        - nd-router-preference
                        - other-stateful-configuration
                        - prefix
                            - autonomous
                            - on-link
                            - preferred-lifetime
                            - valid-lifetime
                        - reachable-time
                        - retransmit-time
                        - router-lifetime
                        - shutdown
                        - use-virtual-mac
                - router-id
                - sgt-qos
                    - application
                    - dscp
                - shutdown
                - shutdown
                - single-sfm-overload
                - snmp
                    - access
                    - community
                - source-address
                    - application
                    - application6
                - spoke-sdp
                    - description
                - static-route-entry
                    - black-hole
                        - community
                        - description
                        - generate-icmp
                        - metric
                        - preference
                        - prefix-list
                        - shutdown
                        - tag
                    - community
                    - grt
                        - description
                        - metric
                        - preference
                        - shutdown
                    - indirect
                        - community
                        - cpe-check
```

config service vprn static-route-entry indirect cpe-check drop-count

```
                        - drop-count
                        - interval
                        - log
                        - padding-size
                    - description
                    - metric
                    - preference
                    - prefix-list
                    - shutdown
                    - tag
                - ipsec-tunnel
                    - community
                    - description
                    - metric
                    - preference
                    - shutdown
                    - tag
                - next-hop
                    - bfd-enable
                    - community
                    - cpe-check
                        - drop-count
                        - interval
                        - log
                        - padding-size
                    - description
                    - load-balancing-weight
                    - metric
                    - preference
                    - prefix-list
                    - shutdown
                    - tag
                    - validate-next-hop
                - tag
            - static-route-hold-down
            - ttl-propagate
                - local
                - transit
            - twamp-light
                - reflector
                    - allow-ipv6-udp-checksum-zero
                    - description
                    - prefix
                        - description
                    - shutdown
                    - type
            - type
            - vrf-export
            - vrf-import
            - vrf-target
            - weighted-ecmp
            - weighted-ecmp
```

3.4.20 configure subscriber-mgmt Commands

```
        – subscriber-mgmt
            – local-user-db
                – description
                – ipoe
                    – host
                        – address
                        – gi-address
                        – host-identification
                            – circuit-id
                            – duid-en
                            – duid-ll-llt
                            – mac
                            – option60
                            – remote-id
                            – sap-id
                            – service-id
                            – string
                            – system-id
                        – options
                            – custom-option
                            – default-router
                            – dns-server
                            – domain-name
                            – lease-rebind-time
                            – lease-renew-time
                            – lease-time
                            – netbios-name-server
                            – netbios-node-type
                            – subnet-mask
                        – options6
                            – boot-file-param
                            – boot-file-url
                            – dns-server
                        – rip-policy
                        – shutdown
                    – mask
                    – match-list
                – shutdown
            – rip-policy
                – authentication-key
                – authentication-type
                – description
```

3.4.21 configure system Commands

```
        – system
            – alarms
                – max-cleared
                – shutdown
            – allow-boot-license-violations
            – boot-bad-exec
            – boot-good-exec
            – clli-code
            – config-backup
            – congestion-management
            – contact
            – coordinates
            – cron
                – schedule
                    – count
                    – day-of-month
                    – description
                    – end-time
                    – hour
                    – interval
                    – minute
                    – month
                    – script-policy
                    – shutdown
                    – type
                    – weekday
            – dhcp6
                – adv-noaddrs-global
            – dns
                – address-pref
                – dnssec
                    – ad-validation
            – enable-icmp-vse
            – file-transmission-profile
                – http-version
                – ipv4-source-address
                – ipv6-source-address
                – redirection
                – retry
                – router
                – timeout
            – grpc
                – allow-unsecure-connection
                – delay-on-boot
                – gnmi
                    – auto-config-save
                    – proto-version
                    – shutdown
                – gnoi
                    – cert-mgmt
                        – shutdown
                    – file
                        – shutdown
                    – system
                        – shutdown
                – listening-port
                – max-msg-size
                – md-cli
                    – shutdown
                – shutdown
```

config system grpc tcp-keepalive

```
            - tcp-keepalive
                - idle-time
                - interval
                - retries
                - shutdown
            - tls-server-profile
        - grpc-tunnel
            - delay-on-boot
            - destination-group
                - allow-unsecure-connection
                - description
                - destination
                    - local-source-address
                    - originated-qos-marking
                    - router-instance
                - tcp-keepalive
                    - idle-time
                    - interval
                    - retries
                    - shutdown
                - tls-client-profile
            - tunnel
                - description
                - destination-group
                - handler
                    - port
                    - shutdown
                    - target-type
                - shutdown
                - target-name
        - ip
            - buffer-unresolved-packets
            - enforce-unique-if-index
            - forward-6in4
            - forward-ip-over-gre
            - ipv6-eh
        - lacp-system-priority
        - lldp
            - message-fast-tx
            - message-fast-tx-init
            - notification-interval
            - reinit-delay
            - shutdown
            - tx-credit-max
            - tx-hold-multiplier
            - tx-interval
        - load-balancing
            - l4-load-balancing
            - lsr-load-balancing
            - service-id-lag-hashing
        - location
        - login-control
            - exponential-backoff
            - ftp
                - inbound-max-sessions
            - idle-timeout
            - login-banner
            - login-scripts
                - global
                - per-user
            - motd
            - pre-login-message
            - ssh
                - disable-graceful-shutdown
```

config system login-control ssh inbound-max-sessions

```
                - inbound-max-sessions
                - max-channels-per-connection
                - outbound-max-sessions
                - ttl-security
            - telnet
                - enable-graceful-shutdown
                - inbound-max-sessions
                - outbound-max-sessions
                - ttl-security
        - management-interface
            - cli
                - classic-cli
                    - allow-immediate
                - cli-engine
                - md-cli
                    - auto-config-save
                    - environment
                        - command-completion
                            - enter
                            - space
                            - tab
                        - commit-options
                            - comment
                            - confirm
                        - console
                            - length
                            - width
                        - history
                            - recall
                            - size
                        - info-output
                            - always-display
                                - admin-state
                        - message-severity-level
                            - cli
                        - more
                        - progress-indicator
                            - delay
                            - shutdown
                            - type
                        - prompt
                            - context
                            - newline
                            - timestamp
                            - uncommitted-changes-indicator
                        - time-display
                        - time-format
            - configuration-mode
            - operations
                - global-timeouts
                    - asynchronous-execution
                    - asynchronous-retention
                    - synchronous-execution
            - remote-management
                - allow-unsecure-connection
                - client-tls-profile
                - connection-timeout
                - delay-on-boot
                - device-label
                - device-name
                - hello-interval
                - manager
                    - allow-unsecure-connection
                    - client-tls-profile
```

config system management-interface remote-management manager connection-timeout

```
                    - connection-timeout
                    - description
                    - device-label
                    - device-name
                    - manager-address
                    - manager-port
                    - router
                    - shutdown
                    - source-address
                    - source-port
                - router
                - shutdown
                - source-address
                - source-port
            - schema-path
            - yang-modules
                - nmda
                    - nmda-support
                - nokia-combined-modules
                - nokia-submodules
        - name
        - netconf
            - auto-config-save
            - call-home
                - device-labels
                    - advertise-operating-system
                    - advertise-software-version
                    - advertise-system-name
                    - device-label
                - netconf-client
                    - connection-type
                    - delay-on-boot
                    - description
                    - remote-port
                    - router-instance
                    - shutdown
                    - transport
            - capabilities
                - candidate
            - listen
                - delay-on-boot
                - port
                - shutdown
        - network-element-discovery
            - generate-traps
            - profile
                - neid
                - neip
                    - ipv4
                    - ipv6
                - platform-type
                - system-mac
                - vendor-id
        - ospf-dynamic-hostnames
        - persistence
            - ancp
                - description
                - location
            - dhcp-server
                - description
                - location
            - nat-port-forwarding
                - description
                - location
```

config system persistence options

```
            - options
                - dhcp-leasetime-threshold
        - rollback
            - local-max-checkpoints
            - remote-max-checkpoints
            - rescue-location
            - rollback-location
        - script-control
            - script
                - description
                - location
                - shutdown
            - script-policy
                - expire-time
                - lifetime
                - lock-override
                - max-completed
                - results
                - script
                - shutdown
        - security
            - cli-script
                - authorization
                    - cron
                        - cli-user
                    - event-handler
                        - cli-user
            - cli-session-group
                - combined-max-sessions
                - description
                - ssh-max-sessions
                - telnet-max-sessions
            - copy
            - dist-cpu-protection
                - policy
                    - description
                    - local-monitoring-policer
                        - description
                        - exceed-action
                        - log-events
                        - rate
                    - protocol
                        - dynamic-parameters
                            - detection-time
                            - exceed-action
                            - log-events
                            - rate
                        - enforcement
                    - static-policer
                        - description
                        - detection-time
                        - exceed-action
                        - log-events
                        - rate
            - ftp-server
            - keychain
                - description
                - direction
                    - bi
                        - entry
                            - begin-time
                            - option
                            - shutdown
                            - tolerance
```

config system security keychain direction uni

```
                    - uni
                        - receive
                            - entry
                                - begin-time
                                - end-time
                                - shutdown
                                - tolerance
                        - send
                            - entry
                                - begin-time
                                - shutdown
                - shutdown
                - tcp-option-number
                    - receive
                    - send
            - ldap
                - public-key-authentication
                - retry
                - route-preference
                - server
                    - address
                    - bind-authentication
                    - ldap-server
                    - search
                    - shutdown
                    - tls-profile
                - shutdown
                - timeout
                - use-default-template
            - management
                - allow-ftp
                - allow-grpc
                - allow-netconf
                - allow-ssh
                - allow-telnet
                - allow-telnet6
            - management-access-filter
                - ip-filter
                    - default-action
                    - entry
                        - action
                        - description
                        - dst-port
                        - l4-src-port
                        - log
                        - protocol
                        - router
                        - src-ip
                        - src-port
                    - renum
                    - shutdown
                - ipv6-filter
                    - default-action
                    - entry
                        - action
                        - description
                        - dst-port
                        - flow-label
                        - l4-src-port
                        - log
                        - next-header
                        - router
                        - src-ip
                        - src-port
```

config system security mgmt-access-filter ipv6-filter renum

```
                    - renum
                    - shutdown
                - mac-filter
                    - default-action
                    - entry
                        - action
                        - description
                        - log
                        - match
                            - dot1p
                            - dsap
                            - dst-mac
                            - etype
                            - snap-oui
                            - snap-pid
                            - src-mac
                            - ssap
                            - svc-id
                    - renum
                    - shutdown
            - management-interface
                - classic-cli
                    - read-algorithm
                    - write-algorithm
                - grpc
                    - hash-algorithm
                - md-cli
                    - command-accounting-during-load
                    - hash-algorithm
                - netconf
                    - hash-algorithm
                - output-authorization
                    - md-interfaces
                    - telemetry-data
                    - telemetry-default-user
            - password
                - admin-password
                - aging
                - attempts
                - authentication-order
                - complexity-rules
                    - allow-user-name
                    - credits
                    - disallow-sequence-keys
                    - minimum-classes
                    - minimum-length
                    - repeated-characters
                    - required
                - enable-admin-control
                    - tacplus-map-to-priv-lvl
                - hashing
                - health-check
                - history-size
                - minimum-age
                - minimum-change
            - per-peer-queuing
            - pki
                - ca-profile
                    - auto-crl-update
                        - crl-urls
                            - url-entry
                                - file-transmission-profile
                                - url
                        - periodic-update-interval
```

config system security pki ca-prof auto-crl-update pre-update-time

```
                        - pre-update-time
                        - retry-interval
                        - schedule-type
                        - shutdown
                    - cert-file
                    - cmpv2
                        - accept-unprotected-errormsg
                        - accept-unprotected-pkiconf
                        - always-set-sender-for-ir
                        - http-response-timeout
                        - http-response-timeout
                        - http-version
                        - key-list
                            - key
                        - recipient
                        - response-signing-cert
                        - same-recipnonce-for-pollreq
                        - url
                    - crl-file
                    - description
                    - ocsp
                        - responder-url
                        - service
                        - transmission-profile
                    - revocation-check
                    - shutdown
                - certificate-auto-update
                    - cert
                        - key
                        - profile
                - certificate-display-format
                - certificate-expiration-warning
                - certificate-update-profile
                    - hash-algorithm
                    - key-generation
                    - protocol
                    - retry-interval
                    - schedule
                        - time
                - common-name-list
                    - common-name
                - crl-expiration-warning
                - dynamic-ca
                - est-profile
                    - check-id-kp-cmcra-only
                    - client-tls-profile
                    - http-auth
                    - server
                    - transmission-profile
                - imported-format
                - maximum-cert-chain-depth
            - profile
                - cli-session-group
                - combined-max-sessions
                - default-action
                - entry
                    - action
                    - description
                    - match
                - grpc
                    - rpc-authorization
                        - gnmi-capabilities
                        - gnmi-get
                        - gnmi-set
```

config system security profile grpc rpc-authorization gnmi-subscribe

```
                        - gnmi-subscribe
                        - gnoi-cert-mgmt-cangenerate
                        - gnoi-cert-mgmt-getcert
                        - gnoi-cert-mgmt-install
                        - gnoi-cert-mgmt-revoke
                        - gnoi-cert-mgmt-rotate
                        - gnoi-file-get
                        - gnoi-file-put
                        - gnoi-file-remove
                        - gnoi-file-stat
                        - gnoi-file-transfertoremote
                        - gnoi-system-cancelreboot
                        - gnoi-system-ping
                        - gnoi-system-reboot
                        - gnoi-system-rebootstatus
                        - gnoi-system-setpackage
                        - gnoi-system-switchcontrolprocessor
                        - gnoi-system-time
                        - gnoi-system-traceroute
                        - md-cli-session
                    - netconf
                        - base-op-authorization
                            - action
                            - cancel-commit
                            - close-session
                            - commit
                            - copy-config
                            - create-subscription
                            - delete-config
                            - discard-changes
                            - edit-config
                            - get
                            - get-config
                            - get-data
                            - get-schema
                            - kill-session
                            - lock
                            - validate
                    - renum
                    - ssh-max-sessions
                    - telnet-max-sessions
                - radius
                    - access-algorithm
                    - accounting
                    - accounting-port
                    - authorization
                    - interactive-authentication
                    - port
                    - retry
                    - route-preference
                    - server
                    - shutdown
                    - timeout
                    - use-default-template
                - snmp
                    - access
                    - attempts
                    - community
                    - src-access-list
                        - src-host
                    - usm-community
                    - view
                        - mask
                - source-address
```

config system security source-address application

```
                - application
                - application6
            - ssh
                - authentication-method
                    - client
                        - public-key-only
                    - server
                        - public-key-only
                - client-cipher-list
                    - cipher
                - client-host-key-list
                    - host-key
                - client-kex-list
                    - kex
                - client-mac-list
                    - mac
                - key-re-exchange
                    - client
                        - mbytes
                        - minutes
                        - shutdown
                    - server
                        - mbytes
                        - minutes
                        - shutdown
                - listening-port
                - permit-empty-passwords
                - preserve-key
                - server-cipher-list
                    - cipher
                - server-host-key-list
                    - host-key
                - server-kex-list
                    - kex
                - server-mac-list
                    - mac
                - server-shutdown
            - tacplus
                - accounting
                - authorization
                - ignore-unknown-mandatory-vsas
                - interactive-authentication
                - priv-lvl-map
                    - priv-lvl
                - request-format
                    - access-operation-cmd
                - retry-timeout
                - route-preference
                - server
                - service-request
                    - nokia-grpc-rpc-authorization
                    - nokia-netconf-base-op-authorization
                    - nokia-user
                    - nokia-user-profile
                - shutdown
                - timeout
                - use-default-template
            - tech-support
                - ts-location
            - telnet
                - listening-port
            - telnet-server
            - telnet6-server
            - tls
```

config system security tls cert-profile

```
                    - cert-profile
                        - entry
                            - cert
                            - key
                            - send-chain
                                - ca-profile
                        - shutdown
                    - client-cipher-list
                        - cipher
                        - tls13-cipher
                    - client-group-list
                        - tls13-group
                    - client-signature-list
                        - tls13-signature
                    - client-tls-profile
                        - cert-profile
                        - cipher-list
                        - group-list
                        - protocol-version
                        - shutdown
                        - signature-list
                        - status-verify
                            - default-result
                            - ee-revocation
                        - trust-anchor-profile
                    - server-cipher-list
                        - cipher
                        - tls13-cipher
                    - server-group-list
                        - tls13-group
                    - server-signature-list
                        - tls13-signature
                    - server-tls-profile
                        - authenticate-client
                            - cn-authentication
                            - trust-anchor-profile
                        - cert-profile
                        - cipher-list
                        - group-list
                        - protocol-version
                        - shutdown
                        - signature-list
                        - status-verify
                            - default-result
                            - ee-revocation
                        - tls-re-negotiate-timer
                    - trust-anchor-profile
                        - trust-anchor
                - user
                    - access
                    - cli-engine
                    - console
                        - cannot-change-password
                        - login-exec
                        - member
                        - new-password-at-login
                    - home-directory
                    - password
                    - public-keys
                        - ecdsa
                            - ecdsa-key
                                - description
                                - key-value
                        - rsa
```

config system security user public-keys rsa rsa-key

```
                        - rsa-key
                            - description
                            - key-value
                - restricted-to-home
                - save-when-restricted
                - snmp
                    - authentication
                    - group
                - ssh-authentication-method
                    - client
                        - public-key-only
                    - server
                        - public-key-only
            - user-template
                - access
                - console
                    - login-exec
                - home-directory
                - profile
                - restricted-to-home
                - save-when-restricted
            - vprn-aaa-server
                - inband
                - outband
                - vprn
            - vprn-network-exceptions
        - snmp
            - engineID
            - general-port
            - max-bulk-duration
            - packet-size
            - shutdown
            - streaming
                - shutdown
            - transport
        - switchover-exec
        - telemetry
            - destination-group
                - allow-unsecure-connection
                - description
                - destination
                    - router-instance
                - tcp-keepalive
                    - idle-time
                    - interval
                    - retries
                    - shutdown
                - tls-client-profile
            - notification-bundling
                - max-msg-count
                - max-time-granularity
                - shutdown
            - persistent-subscriptions
                - delay-on-boot
                - subscription
                    - description
                    - destination-group
                    - encoding
                    - local-source-address
                    - mode
                    - originated-qos-marking
                    - sample-interval
                    - sensor-group
                    - shutdown
```

config system telemetry sensor-groups

```
            - sensor-groups
                - sensor-group
                    - description
                    - path
        - thresholds
            - cflash-cap-alarm
            - cflash-cap-alarm-pct
            - cflash-cap-warn
            - cflash-cap-warn-pct
            - kb-memory-use-alarm
            - kb-memory-use-warn
            - memory-use-alarm
            - memory-use-warn
            - rmon
                - alarm
                - event
        - time
            - dst-zone
                - end
                - offset
                - start
            - ntp
                - authentication-check
                - authentication-key
                - authentication-keychain
                - broadcast
                - broadcastclient
                - multicast
                - multicastclient
                - ntp-server
                - peer
                - server
                - shutdown
            - prefer-local-time
            - sntp
                - broadcast-client
                - server-address
                - shutdown
            - zone
        - usb
            - shutdown
```

3.4.22 configure test-oam Commands

```
        – test-oam
            – icmp
                – ipv6
                    – length-field
                    – maximum-original-datagram
            – mpls-echo-request-downstream-map
            – mpls-time-stamp-format
            – twamp
                – server
                    – allow-ipv6-udp-checksum-zero
                    – enforce-test-session-start-time
                    – inactivity-timeout
                    – max-conn-server
                    – max-sess-server
                    – prefix
                        – description
                        – max-conn-prefix
                        – max-sess-prefix
                    – shutdown
                – twamp-light
                    – inactivity-timeout
```

3.4.23 configure vrrp Commands

```
- vrrp
    - policy
        - delta-in-use-limit
        - description
        - priority-event
            - host-unreachable
                - drop-count
                - hold-clear
                - hold-set
                - interval
                - padding-size
                - priority
                - timeout
            - lag-port-down
                - hold-clear
                - hold-set
                - number-down
                    - priority
                - weight-down
                    - priority
            - mc-ipsec-non-forwarding
                - hold-clear
                - hold-set
                - priority
            - port-down
                - hold-clear
                - hold-set
                - priority
            - route-unknown
                - hold-clear
                - hold-set
                - less-specific
                - next-hop
                - priority
                - protocol
        - shutdown
```

3.5 debug Commands

```
- debug
    - certificate
        - auto-cert-update
        - auto-crl-update
            - ca-profile
        - cmpv2
            - ca-profile
        - est-profile
        - ocsp
            - ca-profile
    - ipsec
        - certificate
        - client-db
        - gateway
        - transport-mode
        - tunnel
    - lag
    - macsec
    - mirror-source
        - ip-filter
        - ipv6-filter
        - port
        - sap
        - shutdown
    - oam
        - lsp-ping-trace
    - pcap
        - capture
    - radius
    - router
        - autoconfigure
            - dhcp-client
                - events
                - packet
                    - detail-level
                    - mode
                - rtm
            - dhcp6-client
                - events
                - packet
                    - detail-level
                    - mode
                - rtm
        - bgp
            - events
            - graceful-restart
            - keepalive
            - notification
            - open
            - outbound-route-filtering
            - packets
            - route-refresh
            - rtm
            - socket
            - timers
            - update
        - igmp
            - group-interface
            - host
```

debug router igmp interface

```
                - interface
                - mcs
                - misc
                - packet
            - ip
                - arp
                - dhcp
                    - detail-level
                    - mode
                - dhcp6
                    - detail-level
                    - mode
                - event
                    - ipv6-error
                - icmp
                - icmp6
                - interface
                - neighbor
                - packet
                - route-table
                - tunnel-table
            - isis
                - adjacency
                - cspf
                - graceful-restart
                - interface
                - leak
                - lsdb
                - misc
                - packet
                - rtm
                - spf
                - summary
                - tunnel-endpoint
            - ldp
                - interface
                  - event
                      - messages
                  - packet
                      - hello
                - peer
                  - event
                      - bindings
                      - messages
                  - packet
                      - hello
                      - init
                      - keepalive
                      - label
            - local-dhcp-server
                - detail-level
                - mode
            - mld
                - group-interface
                - host
                - interface
                - ipsec-interface
                - mcs
                - misc
                - packet
            - mpls
                - event
                    - all
                    - frr
```

debug router mpls event iom

```
                - iom
                - lsp-setup
                - mbb
                - misc
                - pcc
                - te
                - xc
            - forwarding-policies
                - binding-label
                - endpoint
        - ospf
            - area
            - area-range
            - cspf
            - graceful-restart
            - interface
            - leak
            - lsdb
            - misc
            - neighbor
            - nssa-range
            - packet
            - rsvp-shortcut
            - rtm
            - sham-neighbor
            - spf
            - tunnel-endpoint
            - virtual-neighbor
        - ospf3
            - area
            - area-range
            - graceful-restart
            - interface
            - leak
            - lsdb
            - misc
            - neighbor
            - nssa-range
            - packet
            - rtm
            - spf
            - tunnel-endpoint
            - virtual-neighbor
        - pcep
            - pcc
                - all
                - connection
                    - all
                    - cspf-te
                    - db
                    - error
                    - msg
                    - packet
                    - red
                    - task
                - cspf-te
                - db
                - error
                - msg
                - packet
                - red
                - task
        - pim
            - adjacency
```

debug router pim all

```
            - all
            - assert
            - auto-rp
            - bgp
            - bsr
            - data
            - db
            - dynmldp
            - extranet
            - graft
            - interface
            - jp
            - mrib
            - msg
            - mvpn-rtcache
            - packet
            - red
            - register
            - rpfv
            - rtm
            - s-pmsi
            - tunnel-interface
        - radius
            - detail-level
            - packet-type
            - radius-attr
            - server-address
        - radius-proxy
            - server
                - client-address
                - detail-level
                - direction
                - dropped-only
                - packet-type
        - rip
            - auth
            - error
            - events
            - holddown
            - packets
            - requests
            - trigger
            - updates
        - ripng
            - error
            - events
            - holddown
            - packets
            - requests
            - trigger
            - updates
        - rpki-session
            - packet
                - all
                - cache-reset
                - cache-response
                - end-of-data
                - error-report
                - ipv4-prefix
                - ipv6-prefix
                - reset-query
                - serial-notify
                - serial-query
        - rsvp
```

debug router rsvp event

```
                - event
                    - all
                    - auth
                    - misc
                    - nbr
                    - path
                    - resv
                    - rr
                    - te-threshold-update
                - packet
                    - ack
                    - all
                    - bundle
                    - hello
                    - path
                    - patherr
                    - pathtear
                    - resv
                    - resverr
                    - resvtear
                    - srefresh
            - vrrp
                - events
                - packets
        - service
            - id
                - dhcp
                    - detail-level
                    - mac
                    - mode
                    - sap
                    - sdp
                - dhcp6
                    - detail-level
                    - mac
                    - mode
                    - sap
                - event-type
                - igmp-snooping
                    - detail-level
                    - evpn-mpls
                    - mac
                    - mode
                    - sap
                    - sdp
                - mld-snooping
                    - detail-level
                    - evpn-mpls
                    - mac
                    - mode
                    - sap
                    - sdp
                - proxy-arp
                - proxy-nd
                - sap
                    - event-type
                - sdp
                    - event-type
                - stp
                    - all-events
                    - bpdu
                    - bpdu
                    - core-connectivity
                    - exception
```

debug service id stp fsm-state-changes

```
                - fsm-state-changes
                - fsm-timers
                - port-role
                - port-state
                - sap
                - sdp
        - sdp
            - event-type
    - snmp
    - subscriber-mgmt
        - local-user-db
            - detail
    - system
        - grpc
            - client
            - type
        - grpc-tunnel
            - tunnel
        - http-connections
        - management-interface
            - remote-management
        - netconf
        - nsp-proxy
            - history
        - ntp
        - persistence
```

3.6 environment Commands

```
- environment
    - alias
    - create
    - kernel
    - more
    - reduced-prompt
    - saved-ind-prompt
    - shell
    - suggest-internal-objects
    - terminal
        - length
        - width
    - time-display
    - time-stamp
```

3.7 file Commands

```
- file
    - attrib
    - cd
    - checksum
    - copy
    - delete
    - dir
    - format
    - md
    - move
    - rd
    - repair
    - scp
    - shutdown
    - type
    - unzip
    - version
    - vi
```

3.8 Global Commands

```
        - back
        - echo
        - enable-admin
        - exec
        - exit
        - help
        - history
        - info
        - logout
        - password
        - ping
        - pwc
        - sleep
        - ssh
        - telnet
        - traceroute
        - tree
        - write
```

3.9 oam Commands

```
- oam
    - lsp-ping
        - bgp-label
        - ldp
        - prefix
        - rsvp-te
        - sr-isis
        - sr-ospf
        - sr-policy
        - sr-te
    - lsp-trace
        - bgp-label
        - ldp
        - prefix
        - rsvp-te
        - sr-isis
        - sr-ospf
        - sr-policy
        - sr-te
    - oam-pm
    - saa
    - sdp-mtu
    - sdp-ping
```

4 a Commands – Part I

4.1 aaa

aaa

Syntax
aaa

Context
[Tree] (config aaa)

Full Context
configure aaa

Description
Commands in this context configure authentication, authorization, and accounting.

Platforms
7705 SAR Gen 2

aaa

Syntax
aaa

Context
[Tree] (config>service>vprn aaa)

Full Context
configure service vprn aaa

Description
Commands in this context configure AAA on the VPRN.

Platforms
7705 SAR Gen 2

4.2 abort

abort

Syntax
abort

Context
[Tree] (config>router>bfd abort)

Full Context
configure router bfd abort

Description
This command discards the changes made to a BFD template during an active session.

Platforms
7705 SAR Gen 2

abort

Syntax
abort

Context
[Tree] (config>router>route-next-hop-policy abort)

Full Context
configure router route-next-hop-policy abort

Description
This command discards the changes made to route next-hop templates during an active session.

Platforms
7705 SAR Gen 2

abort

Syntax

abort

Context

[Tree] (config>router>policy-options abort)

Full Context

configure router policy-options abort

Description

This command is required to discard changes made to a route policy.

Platforms

7705 SAR Gen 2

4.3 accept-coa

accept-coa

Syntax

[**no**] **accept-coa**

Context

[Tree] (config>service>vprn>radius-server>server accept-coa)

[Tree] (config>router>radius-server>server accept-coa)

Full Context

configure service vprn radius-server server accept-coa

configure router radius-server server accept-coa

Description

This command configures this server for Change of Authorization messages. The system will process the CoA request from the external server if configured with this command; otherwise the CoA request is dropped.

The **no** form of this command disables the command.

Platforms

7705 SAR Gen 2

4.4 accept-from-ebgp

accept-from-ebgp

Syntax

accept-from-ebgp *family* [*family*]

no accept-from-ebgp

Context

[Tree] (config>service>vprn>bgp>group>neighbor>link-bandwidth accept-from-ebgp)

[Tree] (config>service>vprn>bgp>group>link-bandwidth accept-from-ebgp)

Full Context

configure service vprn bgp group neighbor link-bandwidth accept-from-ebgp

configure service vprn bgp group link-bandwidth accept-from-ebgp

Description

This command configures BGP to accept and use the link-bandwidth extended community attached to any route received from any EBGP peer in the scope of the command, as long as that route belongs to one of the listed address families.

The link-bandwidth extended community is encoded as a non-transitive type. This means that by default it should not be attached to any route advertised to an EBGP peer and it should be discarded when received in any route from an EBGP peer. This command overrides the standard behavior.

Up to three families may be configured.

The **no** form of this command restores the default behavior of discarding the link-bandwidth extended community in any route received from an EBGP peer.

Default

no accept-from-ebgp

Parameters

family

Specifies the address families for which receiving the link-bandwidth extended community from EBGP peers should be supported.

Values ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes.

label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes.

ipv6 — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes.

Platforms

7705 SAR Gen 2

accept-from-ebgp

Syntax

accept-from-ebgp *family* [*family*]

no accept-from-ebgp

Context

[Tree] (config>router>bgp>group>neighbor>link-bandwidth accept-from-ebgp)

[Tree] (config>router>bgp>group>link-bandwidth accept-from-ebgp)

Full Context

configure router bgp group neighbor link-bandwidth accept-from-ebgp

configure router bgp group link-bandwidth accept-from-ebgp

Description

This command configures BGP to accept and use the link-bandwidth extended community attached to any route received from any EBGP peer in the scope of the command, as long as that route belongs to one of the listed address families.

The link-bandwidth extended community is encoded as a non-transitive type. This means that by default it should not be attached to any route advertised to an EBGP peer and it should be discarded when received in any route from an EBGP peer. This command overrides the standard behavior.

Up to six families may be configured.

The **no** form of this command restores the default behavior of discarding the link-bandwidth extended community in any route received from an EBGP peer.

Default

no accept-from-ebgp

Parameters

family

Specifies the address families for which receiving the link-bandwidth extended community from EBGP peers should be supported.

Values ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes.

label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes.

vpn-ipv4 — Adds a link-bandwidth extended community to IPv4 VPN (SAFI 128) routes.

ipv6 — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes.

label-ipv6 — Adds a link-bandwidth extended community to labeled-unicast IPv6 routes.

vpn-ipv6 — Adds a link-bandwidth extended community to IPv6 VPN (SAFI 128) routes.

Platforms

7705 SAR Gen 2

4.5 accept-orf

accept-orf

Syntax

[**no**] **accept-orf**

Context

[Tree] (config>router>bgp>outbound-route-filtering>extended-community accept-orf)

[Tree] (config>router>bgp>group>outbound-route-filtering>extended-community accept-orf)

[Tree] (config>router>bgp>group>neighbor>outbound-route-filtering>extended-community accept-orf)

Full Context

configure router bgp outbound-route-filtering extended-community accept-orf

configure router bgp group outbound-route-filtering extended-community accept-orf

configure router bgp group neighbor outbound-route-filtering extended-community accept-orf

Description

This command instructs the router to negotiate the receive capability in the BGP ORF negotiation with a peer, and accept filters that the peer wants to send.

The **no** form of this command causes the router to remove the accept capability in the BGP ORF negotiation with a peer, and to clear any existing ORF filters that are currently in place.

Default

no accept-orf

Platforms

7705 SAR Gen 2

4.6 accept-unprotected-errormsg

accept-unprotected-errormsg

Syntax

[**no**] **accept-unprotected-errormsg**

Context

[Tree] (config>system>security>pki>ca-profile>cmpv2 accept-unprotected-errormsg)

Full Context

configure system security pki ca-profile cmpv2 accept-unprotected-errormsg

Description

This command enables the system to accept both protected and unprotected CMPv2 error message. Without this command, system will only accept protected error messages.

The **no** form of this command causes the system to only accept protected PKI confirmation message.

Default

no accept-unprotected-errormsg

Platforms

7705 SAR Gen 2

4.7 accept-unprotected-pkiconf

accept-unprotected-pkiconf

Syntax

[**no**] **accept-unprotected-pkiconf**

Context

[Tree] (config>system>security>pki>ca-profile>cmpv2 accept-unprotected-pkiconf)

Full Context

configure system security pki ca-profile cmpv2 accept-unprotected-pkiconf

Description

This command enables the system to accept both protected and unprotected CMPv2 PKI confirmation messages. Without this command, the system will only accept protected PKI confirmation message.

The **no** form of this command causes the system to only accept protected PKI confirmation message.

Default

no accept-unprotected-pkiconf

Platforms

7705 SAR Gen 2

4.8 access

access

Syntax

access

Context

[Tree] (config>port>ethernet access)

Full Context

configure port ethernet access

Description

This command configures Ethernet access port parameters.

Platforms

7705 SAR Gen 2

access

Syntax

access

Context

[Tree] (config>port access)

[Tree] (config>card>mda access)

Full Context

configure port access

configure card mda access

Description

This command enables the access context to configure egress and ingress pool policy parameters.

On the MDA level, access egress and ingress pools are only allocated on channelized MDAs.

Platforms

7705 SAR Gen 2

access

Syntax

access

Context

[Tree] (config>card>fp>ingress access)

Full Context

configure card fp ingress access

Description

This CLI node contains the access forwarding-plane parameters.

Platforms

7705 SAR Gen 2

access

Syntax

access

Context

[Tree] (config>lag access)

Full Context

configure lag access

Description

Commands in this context configure access parameters.

Platforms

7705 SAR Gen 2

access

Syntax

[**no**] **access**

Context

[Tree] (config>service>vprn>snmp access)

Full Context

configure service vprn snmp access

Description

This command enables SNMP access using VPRN interface addresses. This command allows SNMP messages destined to the VPRN interface IP addresses for this VPRN (including VPRN interfaces that are bound to R-VPLS services) to be processed by the SNMP agent on the router. SNMP messages that arrive on VPRN interfaces but are destined to IP addresses in the Base routing context that can be accessed in the VPRN (for example, the router system address via grt leaking) do not require **snmp access** to be enabled but do require **allow-local-management** to be enabled.

Using an SNMP community defined inside the VPRN context (**configure service vprn snmp community**) allows access to a subset of the full SNMP data model. This subset can be seen in the output of **show system security view "vprn-view"**.

Using an SNMP community defined in the system context (**configure system security snmp community**) allows access to the full SNMP data model (unless otherwise restricted used SNMP views).

Alternatively, grt leaking and a Base routing IP address can be used (along with an SNMP community defined at the system context) to get access to the entire SNMP data model (see the **allow-local-management** command).

The Nokia NSP cannot discover or fully manage an SR OS router using an SNMP community defined inside the VPRN context. Full SNMP access requires using one of the approaches described above.

See the *7705 SAR Gen 2 System Management Guide* for detailed information about SNMP.

Platforms

7705 SAR Gen 2

access

Syntax

[**no**] **access** [**ftp**] [**snmp**] [**console**] [**li**] [**netconf**] [**grpc**] [**scp-sftp**] [**console-port-cli**] [**ssh-cli**] [**telnet-cli**] [**bluetooth**]

Context

[Tree] (config>system>security>user-template access)

[Tree] (config>system>security>user access)

Full Context

configure system security user-template access

configure system security user access

Description

This command configures user permissions for router management access methods.

To deny an existing access method, enter the **no** form of this command followed by the method to be denied; for example, **no access ftp** denies FTP access.

The **no** form of this command removes the user permission for all management access methods.

Default

no access

Parameters

ftp

Specifies FTP access.

snmp

Specifies SNMP access. This keyword is only configurable in the **configure system security user** context.

console

Specifies Bluetooth, console port CLI, SCP/SFTP, SSH CLI, and Telnet CLI access.

li

Specifies Lawful Intercept (LI) command access.

netconf

Specifies NETCONF access.

grpc

Specifies gRPC access.

scp-sftp

Specifies SCP/SFTP access.

console-port-cli

Specifies console port CLI access.

ssh-cli

Specifies SSH CLI access.

telnet-cli

Specifies Telnet CLI access.

bluetooth

Specifies Bluetooth access.

Platforms

7705 SAR Gen 2

access

Syntax

[**no**] **access group** *group-name* **security-model** *security-model* **security-level** *security-level* [**context** *context-name* [**prefix -match**]] [**read** *view-name-1*] [**write** *view-name-2*] [**notify** *view-name-3*]

Context

[Tree] (config>system>security>snmp access)

Full Context

configure system security snmp access

Description

This command creates an association between a user group, a security model, and the views that the user group can access. Access parameters must be configured unless security is limited to the preconfigured access groups and views for SNMPv1 and SNMPv2. An access group is defined by a unique combination of the group name, security model and security level.

Access groups are used by the usm-community command.

Access must be configured unless security is limited to SNMPv1/SNMPv2c with community strings. See the **community** command.

Default access group configurations cannot be modified or deleted.

To remove the user group with associated, security model(s), and security level(s), use:

no access group *group-name*

To remove a security model and security level combination from a group, use:

no access group *group-name* **security-model** {**snmpv1** | **snmpv2c** | **usm**} **security-level** {**no-auth-no-privacy** | **auth-no-privacy** | **privacy**}

Parameters

group-name

Specify a unique group name up to 32 characters.

security-model {snmpv1 | snmpv2c | usm}

Specifies the security model required to access the views configured in this node. A group can have multiple security models. For example, one view may only require SNMPv1/ SNMPv2c access while another view may require USM (SNMPv3) access rights.

security-level {no-auth-no-priv | auth-no-priv | privacy}

Specifies the required authentication and privacy levels to access the views configured in this node.

security-level no-auth-no-privacy

Specifies that no authentication and no privacy (encryption) is required. When configuring the user's authentication, select the **none** option.

security-level auth-no-privacy

Specifies that authentication is required but privacy (encryption) is not required. When this option is configured, both the **group** and the **user** must be configured for authentication.

security-level privacy

Specifies that both authentication and privacy (encryption) is required. When this option is configured, both the **group** and the user must be configured for **authentication**. The user must also be configured for **privacy**.

context-name

Specifies a set of SNMP objects that are associated with the context-name.

The *context-name* is treated as either a full context-name string or a context name prefix depending on the keyword specified (**exact** or **prefix**).

prefix-match

Specifies the context name **prefix-match** keywords, **exact** or **prefix**.

The VPRN context names begin with a **vprn** prefix. The numerical value is associated with the service ID that the VPRN was created with and identifies the service in the service domain. For example, when a new VPRN service is created such as **config>service>vprn 2345 customer 1**, a VPRN with context name **vprn2345** is created.

The **exact** keyword specifies that an exact match between the context name and the prefix value is required. For example, when context **vprn2345 exact** is entered, matches for only **vprn2345** are considered.

The **prefix** keyword specifies that only a match between the prefix and the starting portion of context name is required. If only the **prefix** keyword is specified, simple wildcard processing is used. For example, when context **vprn prefix** is entered, all **vprn** contexts are matched.

Default **exact**

view-name-1

Specifies the SNMP view used to control which MIB objects can be accessed using a read (get) operation.

view-name-2

Specifies the SNMP view used to control which MIB objects can be accessed using a write (set) operation.

view-name-3

Specifies the SNMP view used to control which MIB objects can be accessed for notifications.

Values none

Platforms

7705 SAR Gen 2

4.9 access-algorithm

access-algorithm

Syntax

access-algorithm {**direct** | **round-robin** | **hash-based**}

no access-algorithm

Context

[Tree] (config>aaa>radius-srv-plcy>servers access-algorithm)

Full Context

configure aaa radius-server-policy servers access-algorithm

Description

This command configures the algorithm used to select a RADIUS server from the pool of configured RADIUS servers.

Default

access-algorithm direct

Parameters

direct

Specifies that the first server is used as primary server for all requests, the second as secondary and so on.

round-robin

Specifies that the first server is used as primary server for the first request, the second server as primary for the second request, and so on. If the router gets to the end of the list, it starts again with the first server.

hash-based

Select a RADIUS server based on the calculated hash result of the configured **load-balance-key** under the **radius-proxy server** hierarchy. This parameter is only applicable for radius-proxy server scenarios and results in an unpredictable RADIUS server selection if used in other scenarios.

Platforms

7705 SAR Gen 2

access-algorithm

Syntax

access-algorithm {**direct** | **round-robin**}

no access-algorithm

Context

[Tree] (config>service>vprn>aaa>rmt-srv>radius access-algorithm)

[Tree] (config>system>security>radius access-algorithm)

Full Context

configure service vprn aaa remote-servers radius access-algorithm

configure system security radius access-algorithm

Description

This command indicates the algorithm used to access the set of RADIUS servers.

Default

access-algorithm direct

Parameters

direct

Specifies that the first server is used as primary server for all requests, the second as secondary and so on.

round-robin

Specifies that the first server is used as primary server for the first request, the second server as primary for the second request, and so on. If the router gets to the end of the list, it starts again with the first server.

Platforms

7705 SAR Gen 2

4.10 access-operation-cmd

access-operation-cmd

Syntax

[**no**] **access-operation-cmd** *access-operation*

Context

[Tree] (config>service>vprn>aaa>rmt-srv>tacplus>req access-operation-cmd)

[Tree] (config>system>security>tacplus>request-format access-operation-cmd)

Full Context

configure service vprn aaa remote-servers tacplus request-format access-operation-cmd

configure system security tacplus request-format access-operation-cmd

Description

This command sends an operation argument in authorization requests.

In model-driven interfaces, this command configures the system to send the operation in the cmd argument, and the path in the cmd-args argument, in TACACS+ authorization requests. This command does not apply to authorization requests in classic interfaces.

The **no** form of this command removes the operation from the configuration.

Default

no access-operation-cmd

Parameters

access-operation

Specifies that an operation in the authorization request is sent.

Values delete — Keyword that sends the operation "cmd=delete" and "cmd-args=path".

Platforms

7705 SAR Gen 2

4.11 accounting

accounting

Syntax

[**no**] **accounting**

Context

[Tree] (config>service>vprn>aaa>rmt-srv>radius accounting)

[Tree] (config>system>security>radius accounting)

Full Context

configure service vprn aaa remote-servers radius accounting

configure system security radius accounting

Description

This command enables RADIUS accounting.

The **no** form of this command disables RADIUS accounting.

Default

no accounting

Platforms

7705 SAR Gen 2

accounting

Syntax

accounting [**record-type** { **start-stop** | **stop-only**}]

no accounting

Context

[Tree] (config>service>vprn>aaa>rmt-srv>tacplus accounting)

[Tree] (config>system>security>tacplus accounting)

Full Context

configure service vprn aaa remote-servers tacplus accounting

configure system security tacplus accounting

Description

This command configures the type of accounting record packet that is to be sent to the TACACS+ server. The **record-type** parameter indicates whether TACACS+ accounting start and stop packets be sent or just stop packets be sent.

Default

no accounting

Parameters

record-type start-stop

Specifies that a TACACS+ start packet is sent whenever the user executes a command and a TACACS+ stop packet when command execution is complete.

record-type stop-only

Specifies that only a TACACS+ stop packet is sent whenever the command execution is complete.

Platforms

7705 SAR Gen 2

4.12 accounting-files-total-size

accounting-files-total-size

Syntax

accounting-files-total-size *megabytes*

Context

[Tree] (config>log>storage accounting-files-total-size)

Full Context

configure log file-storage-control accounting-files-total-size

Description

This command configures the limit for the total space that all accounting files can occupy on each storage device on the active CPM.

When this threshold is reached, new accounting files are no longer created in the `\act-collect` directory of the storage device until SR OS removes older accounting files from the `\act` directory and the occupancy is below the limit. Currently open, in-progress accounting files in the `\act-collect` directory are not affected by this limit and are completed.

When unconfigured, there is no specific limit for the total size of all accounting files.

Only accounting files in the `\act` directory with system generated names (including no file extension) are applicable toward the total size limit.

If a user manually adds or deletes accounting files from the `\act` directory, the size of the files is not taken into account for up to 1 hour.

The configured total size limit is not validated against the actual size of the installed storage devices. If the configured limit is larger than the installed compact flash (CF) device, the limit is never reached.

The **no** form of this command removes the total size limit for accounting files.

Default

no accounting-files-total-size

Parameters

megabytes

Specifies the total size limit for accounting files, in MB.

Values 50 to 4,194,304 MBytes (4 TBytes, 2^{22} MB)

Default 0

Platforms

7705 SAR Gen 2

4.13 accounting-policy

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy

Context

[Tree] (config>service>vpls>sap accounting-policy)

[Tree] (config>service>vprn>if>sap accounting-policy)

[Tree] (config>service>vpls>spoke-sdp accounting-policy)

[Tree] (config>service>vpls>mesh-sdp accounting-policy)

[Tree] (config>service>vprn>if>spoke-sdp accounting-policy)

[Tree] (config>service>ies>if>sap accounting-policy)

Full Context

configure service vpls sap accounting-policy

configure service vprn interface sap accounting-policy

configure service vpls spoke-sdp accounting-policy

configure service vpls mesh-sdp accounting-policy

configure service vprn interface spoke-sdp accounting-policy

configure service ies interface sap accounting-policy

Description

This command creates the accounting policy context that can be applied to an interface SAP or interface SAP spoke SDP.

An accounting policy must be defined before it can be associated with a SAP or SDP.

If the *policy-id* does not exist, an error message is generated.

A maximum of one accounting policy can be associated with a SAP or SDP at one time. Accounting policies are configured in the **config>log** context.

The **no** form of this command removes the accounting policy association from the SAP or SDP, and the accounting policy reverts to the default.

Default

no accounting policy

Parameters

acct-policy-id

Specifies the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 to 99

Platforms

7705 SAR Gen 2

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy

Context

[Tree] (config>card>fp>ingress>access>queue-group accounting-policy)

[Tree] (config>card>fp>ingress>network>queue-group accounting-policy)

Full Context

configure card fp ingress access queue-group accounting-policy

configure card fp ingress network queue-group accounting-policy

Description

This command configures an accounting policy that can apply to a queue-group on the forwarding plane.

An accounting policy must be configured before it can be associated to an interface. If the accounting *policy-id* does not exist, an error is returned.

Accounting policies associated with service billing can only be applied to SAPs. The accounting policy can be associated with an interface at a time.

The **no** form of this command removes the accounting policy association from the queue-group.

Default

No accounting policies are specified by default. You must explicitly specify a policy. If configured, the accounting policy configured as the default is used.

Parameters

acct-policy-id

Specifies the name of the accounting policy to use for the queue-group.

Values 1 to 99

Platforms

7705 SAR Gen 2

accounting-policy

Syntax

accounting-policy *policy-id*

no accounting-policy

Context

[Tree] (config>port>ethernet>network>egr>qgrp accounting-policy)

[Tree] (config>port>ethernet>network accounting-policy)

[Tree] (config>port>ethernet>access>ing>qgrp accounting-policy)

[Tree] (config>port>ethernet accounting-policy)

[Tree] (config>port>ethernet>access>egr>qgrp accounting-policy)

Full Context

configure port ethernet network egress queue-group accounting-policy

configure port ethernet network accounting-policy

configure port ethernet access ingress queue-group accounting-policy

configure port ethernet accounting-policy

configure port ethernet access egress queue-group accounting-policy

Description

This command configures an accounting policy that can apply to an interface.

An accounting policy must be configured before it can be associated to an interface. If the accounting *policy-id* does not exist, an error is returned.

Accounting policies associated with service billing can only be applied to SAPs. Accounting policies associated with network ports can only be associated with interfaces. Only one accounting policy can be associated with an interface at a time.

The **no** form of this command removes the accounting policy association from the network interface, and the accounting policy reverts to the default.

Default

No accounting policies are specified by default. You must explicitly specify a policy. If configured, the accounting policy configured as the default is used.

Parameters

policy-id

The accounting *policy-id* of an existing policy. Accounting policies record either service (access) or network information. A network accounting policy can only be associated with the network port configurations. Accounting policies are configured in the **config>log>accounting-policy** context.

Values 1 to 99

Platforms

7705 SAR Gen 2

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy [*acct-policy-id*]

Context

[Tree] (config>service>epipe>sap accounting-policy)

[Tree] (config>service>epipe>spoke-sdp accounting-policy)

Full Context

configure service epipe sap accounting-policy

configure service epipe spoke-sdp accounting-policy

Description

This command creates the accounting policy context that can be applied to a SAP.

An accounting policy must be defined before it can be associated with a SAP. If the *policy-id* does not exist, an error message is generated.

A maximum of one accounting policy can be associated with a SAP at one time. Accounting policies are configured in the **config>log** context.

The **no** form of this command removes the accounting policy association from the SAP, and the accounting policy reverts to the default.

Default

no accounting policy

Parameters

acct-policy-id

Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 to 99

Platforms

7705 SAR Gen 2

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy

Context

[Tree] (config>service>ies>if>spoke-sdp accounting-policy)

Full Context

configure service ies interface spoke-sdp accounting-policy

Description

This command configures an accounting-policy.

Parameters

acct-policy-id

Specifies an accounting policy ID.

Values 1 to 99

Platforms

7705 SAR Gen 2

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy

Context

[Tree] (config>saa>test accounting-policy)

Full Context

configure saa test accounting-policy

Description

This command associates an accounting policy to the SAA test. The accounting policy must already be defined before it can be associated otherwise an error message is generated.

A notification (trap) is issued whenever a test is completed or terminates.

The **no** form of this command removes the accounting policy association.

Parameters

acct-policy-id

Specifies the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 to 99

Platforms

7705 SAR Gen 2

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy

Context

[Tree] (config>oam-pm>session>meas-interval accounting-policy)

Full Context

configure oam-pm session meas-interval accounting-policy

Description

This optional command allows the operator to assign an accounting policy and the policy-id (configured under the **config>log>accounting-policy**) with a record-type of complete-pm. This runs the data collection process for completed measurement intervals in memory, file storage, and maintenance functions moving data from memory to flash. A single accounting policy can be applied to a measurement interval.

The **no** form of this command removes the accounting policy.

Parameters

acct-policy-id

Specifies the accounting policy to be applied to the measurement interval.

Values 1 to 99

Platforms

7705 SAR Gen 2

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy

Context

[Tree] (config>service>sdp accounting-policy)

[Tree] (config>service>pw-template accounting-policy)

Full Context

configure service sdp accounting-policy

configure service pw-template accounting-policy

Description

This command creates the accounting policy context that can be applied to an SDP. An accounting policy must be defined before it can be associated with a SDP. If the *acct-policy-id* does not exist, an error message is generated.

A maximum of one accounting policy can be associated with a SDP at one time. Accounting policies are configured in the **config>log** context.

The **no** form of this command removes the accounting policy association from the SDP, and the accounting policy reverts to the default.

Default

no accounting-policy

Parameters

acct-policy-id

Specifies the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 to 99

Platforms

7705 SAR Gen 2

accounting-policy

Syntax

accounting-policy *policy-id* [**interval** *minutes*]

no accounting-policy *policy-id*

Context

[Tree] (config>log accounting-policy)

Full Context

configure log accounting-policy

Description

This command creates an access or network accounting policy. An accounting policy defines the accounting records that are created.

Access accounting policies are policies that can be applied to one or more SAPs. Changes made to an existing policy, using any of the sub-commands, are applied immediately to all SAPs where this policy is applied.

If an accounting policy is not specified on a SAP, then accounting records are produced in accordance with the access policy designated as the **default**. If a default access policy is not specified, then no accounting records are collected other than the records for the accounting policies that are explicitly configured.

Only one policy can be regarded as the default access policy. If a policy is configured as the default policy, then a **no default** command must be used to allow the data that is currently being collected to be written before a new access default policy can be configured.

Network accounting policies are policies that can be applied to one or more network ports or SONET/SDH channels. Any changes made to an existing policy, using any of the sub-commands, will be applied immediately to all network ports or SONET/SDH channels where this policy is applied.

If no accounting policy is defined on a network port, accounting records will be produced in accordance with the default network policy as designated with the **default** command. If no network default policy is created, then no accounting records will be collected other than the records for the accounting policies explicitly configured. Default accounting policies cannot be explicitly applied. For example, for **accounting-policy 10**, if default is set, then that policy cannot be used:

```
A:node-2>config>service>vpls>spoke-sdp# accounting-policy 10
```

Only one policy can be regarded as the default network policy. If a policy is configured as the default policy, then a **no default** command must be used to allow the data that is currently being collected to be written before a new network default policy can be configured.

The **no** form of this command deletes the policy from the configuration. The accounting policy cannot be removed unless it is removed from all the SAPs, network ports or channels where the policy is applied.

Parameters

policy-id

Specifies the policy ID that uniquely identifies the accounting policy, expressed as a decimal integer.

Values 1 to 99

Platforms

7705 SAR Gen 2

4.14 accounting-port

accounting-port

Syntax

accounting-port *port*

no accounting-port

Context

[Tree] (config>system>security>radius accounting-port)

[Tree] (config>service>vprn>aaa>rmt-srv>radius accounting-port)

Full Context

configure system security radius accounting-port

configure service vprn aaa remote-servers radius accounting-port

Description

This command specifies a UDP port number on which to contact the RADIUS server for accounting requests.

Default

accounting-port 1813

Parameters

port

Specifies the UDP port number.

Values 1 to 65535

Platforms

7705 SAR Gen 2

4.15 acct-interim

acct-interim

Syntax

acct-interim min *min-val* **max** *max-val* **lifetime** *lifetime*

no acct-interim

Context

[Tree] (config>aaa>radius-srv-plcy>servers>buffering acct-interim)

Full Context

configure aaa radius-server-policy servers buffering acct-interim

Description

This command enables RADIUS accounting interim update message buffering.

1. The message is stored in the buffer, a lifetime timer is started and the message is sent to the RADIUS server
2. If after retry*timeout seconds no RADIUS accounting response is received for the interim update then a new attempt to send the message is started after minimum[(min-val*2n), max-val] seconds.
3. Repeat step 2 until for one of the following:
   a. a RADIUS accounting response is received.
   b. the lifetime of the buffered message expires.
   c. a new RADIUS accounting interim-update or a RADIUS accounting stop for the same accounting session-id and radius-server-policy is stored in the buffer.
   d. the message is manually purged from the message buffer via a clear command.
4. The message is purged from the buffer.

The **no** form of this command disables RADIUS accounting interim update message buffering.

Parameters

min-val

Specifies the minimum interval in seconds between attempts to resend the RADIUS accounting interim update.

Values 1 to 3600

max-val

Specifies the maximum interval in seconds between attempts to resend the RADIUS accounting interim update.

Values 1 to 3600

lifetime

Specifies the lifetime in hours.

Values 1 to 25

Platforms

7705 SAR Gen 2

4.16 acct-on-off

acct-on-off

Syntax

acct-on-off

acct-on-off monitor-group *group-name*

acct-on-off oper-state-change [**group** *group-name*]

Context

[Tree] (config>aaa>radius-srv-plcy acct-on-off)

Full Context

configure aaa radius-server-policy acct-on-off

Description

This command controls the sending of Accounting-On and Accounting-Off messages and the acct-on-off oper-state of the radius-server-policy:

acct-on-off: enables the sending of Accounting-On and Accounting-Off messages for this radius-server-policy. The acct-on-off oper-state is always not blocked.

acct-on-off oper-state-change [group *group-name*]: enables the sending of Accounting-On and Accounting-Off messages for this radius-server-policy. The acct-on-off oper-state is function of the Accounting-response received for the Accounting-On and Accounting-Off. Optionally, sets the acct-on-off oper-state of the acct-on-off-group.

acct-on-off monitor-group *group-name*: no Accounting-On and Accounting-Off messages are sent for this radius-server-policy. The acct-on-off oper-state is inherited from the acct-on-off-group.

The **no** form of this command disables the sending of Accounting-On and Accounting-Off messages.

Parameters

group-name

Specifies the name of an acct-on-off group up to 32 characters.

Platforms

7705 SAR Gen 2

4.17 acct-port

acct-port

Syntax

acct-port *port*

no acct-port

Context

[Tree] (config>service>vprn>radius-server>server acct-port)

[Tree] (config>router>radius-server>server acct-port)

Full Context

configure service vprn radius-server server acct-port

configure router radius-server server acct-port

Description

This command specifies the UDP listening port for RADIUS accounting requests.

The **no** form of this command resets the UDP port to its default value (1813).

Default

acct-port 1813

Parameters

port

Specifies the UDP listening port for accounting requests of the external RADIUS server.

Values 1 to 65535

Platforms

7705 SAR Gen 2

4.18 acct-stats

acct-stats

Syntax

[**no**] **acct-stats**

Context

[Tree] (config>ipsec>rad-acct-plcy>include acct-stats)

Full Context

configure ipsec radius-accounting-policy include-radius-attribute acct-stats

Description

This command enables the system to include accounting attributes in RADIUS acct-stop and interim-update packets.

The **no** form of this command disables the system from including accounting attributes in RADIUS acct-stop and interim-update packets.

Platforms

7705 SAR Gen 2

4.19 acct-stop

acct-stop

Syntax

acct-stop min *min-val* **max** *max-val* **lifetime** *lifetime*

no acct-stop

Context

[Tree] (config>aaa>radius-srv-plcy>servers>buffering acct-stop)

Full Context

configure aaa radius-server-policy servers buffering acct-stop

Description

This command enables RADIUS accounting stop message buffering.

1. The message is stored in the buffer, a lifetime timer is started and the message is sent to the RADIUS server
2. If after retry*timeout seconds no RADIUS accounting response is received for the accounting stop, then a new attempt to send the message is started after minimum[(min-val*2n), max-val] seconds.
3. Repeat step 2 until one of the following events occurs:
   a. A RADIUS accounting response is received.
   b. The lifetime of the buffered message expires.
   c. The message is manually purged from the message buffer via a clear command.
4. The message is purged from the buffer.

The **no** form of this command disables RADIUS accounting stop message buffering.

Parameters

min-val

Specifies the minimum interval in seconds between attempts to resend the RADIUS accounting stop.

Values 1 to 3600

max-val

Specifies the maximum interval in seconds between attempts to resend the RADIUS accounting stop.

Values 1 to 3600

lifetime

Specifies the lifetime in hours.

Values 1 – 25

Platforms

7705 SAR Gen 2

4.20 ack

ack

Syntax

ack [**detail**]

no ack

Context

[Tree] (debug>router>rsvp>packet ack)

Full Context

debug router rsvp packet ack

Description

This command debugs ack events.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about ack events.

Platforms

7705 SAR Gen 2

4.21 acknowledgment

acknowledgment

Syntax

[**no**] **acknowledgment**

Context

[Tree] (config>service>vpls>spoke-sdp>control-channel-status acknowledgment)

[Tree] (config>service>epipe>spoke-sdp>control-channel-status acknowledgment)

Full Context

configure service vpls spoke-sdp control-channel-status acknowledgment

configure service epipe spoke-sdp control-channel-status acknowledgment

Description

This command enables the acknowledgment of control channel status messages. By default, no acknowledgment packets are sent.

Platforms

7705 SAR Gen 2

acknowledgment

Syntax

[**no**] **acknowledgment**

Context

[Tree] (config>service>ies>if>spoke-sdp>control-channel-status acknowledgment)

Full Context

configure service ies interface spoke-sdp control-channel-status acknowledgment

Description

This command enables the acknowledgment of control channel status messages. By default, no acknowledgment packets are sent.

Default

no acknowledgment

Platforms

7705 SAR Gen 2

acknowledgment

Syntax

[**no**] **acknowledgment**

Context

[Tree] (config>service>vprn>if>spoke-sdp>control-channel-status acknowledgment)

Full Context

configure service vprn interface spoke-sdp control-channel-status acknowledgment

Description

This command enables the acknowledgment of control channel status messages. By default, no acknowledgment packets are sent.

Platforms

7705 SAR Gen 2

4.22 action

action

Syntax

action bypass-host-creation

action drop

no action

Context

[Tree] (config>filter>dhcp-filter>entry action)

Full Context

configure filter dhcp-filter entry action

Description

This command specifies the action to take on DHCP host creation when the filter entry matches.

The **no** form of this command reverts to the default wherein the host creation proceeds as normal.

Parameters

bypass-host-creation

Specifies that the host creation is bypassed.

drop

Specifies that the DHCP message is dropped.

Platforms

7705 SAR Gen 2

action

Syntax

action bypass-host-creation [**na**] [**pd**]

action drop

no action

Context

[Tree] (config>filter>dhcp6-filter>entry action)

Full Context

configure filter dhcp6-filter entry action

Description

This command specifies the action to take on DHCP6 host creation when the filter entry matches.

The **no** form of this command reverts to the default wherein the host creation proceeds as normal.

Parameters

bypass-host-creation

Specifies that the host creation is bypassed.

Values **na** — Bypasses the DHCP6 NA hosts creation.

pd — Bypasses the DHCP6 PD hosts creation.

drop

Specifies that the DHCP6 message is dropped.

Platforms

7705 SAR Gen 2

action

Syntax

action {**accept** | **next-entry** | **next-policy** | **drop** | **reject**}

no action

Context

[Tree] (config>router>policy-options>policy-statement>entry action)

Full Context

configure router policy-options policy-statement entry action

Description

This command creates the context to configure actions to take for routes matching a route policy statement entry.

This command is required and must be entered for the entry to be active.

Any route policy entry without the **action** command will be considered incomplete and will be inactive.

The **no** form of this command deletes the action context from the entry.

Default

no action

Parameters

accept

Specifies that routes matching the entry match criteria will be accepted and propagated.

next-entry

Specifies that the actions specified would be made to the route attributes and then policy evaluation would continue with next policy entry (if any others are specified).

next-policy

Specifies that the actions specified would be made to the route attributes and then policy evaluation would continue with next route policy (if any others are specified).

drop

Specifies that routes matching the entry match criteria should be rejected. This parameter provides a context for modifying route properties.

reject

Specifies that routes matching the entry match criteria should be rejected. This parameter does not provide a context for modifying route properties.

Platforms

7705 SAR Gen 2

action

Syntax

action *dhcp-action*

no action

Context

[Tree] (config>service>vpls>sap>dhcp>option action)

[Tree] (config>service>ies>if>dhcp>option action)

[Tree] (config>service>vprn>if>dhcp>option action)

Full Context

configure service vpls sap dhcp option action

configure service ies interface dhcp option action

configure service vprn interface dhcp option action

Description

This command configures the processing required when the SR OS receives a DHCP request that already has a Relay Agent Information Option (Option 82) field in the packet.

The **no** form of this command returns the system to the default value.

Default

action keep — Per RFC 3046, *DHCP Relay Agent Information Option*, section 2.1.1, *Reforwarded DHCP requests*. The default is to keep the existing information intact. The exception to this is if the giaddr of the received packet is the same as the ingress address on the router. In that case the packet is dropped and an error is logged.

Parameters

replace

In the upstream direction (from the user), the existing Option 82 field is replaced with the Option 82 field from the router. In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop

Specifies that the packet is dropped, and an error is logged.

keep

Specifies that the existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on towards the client.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router inserts its own VSO into the Option 82 field. This is only done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO is added to the message.

Platforms

7705 SAR Gen 2

action

Syntax

action {**drop** | **forward**}

no action

Context

[Tree] (config>service>vprn>log>filter>entry action)

[Tree] (config>log>filter>entry action)

Full Context

configure service vprn log filter entry action

configure log filter entry action

Description

This command specifies a drop or forward action associated with the filter entry. If neither **drop** nor **forward** is specified, the default-action will be used for traffic that conforms to the match criteria. This could be considered a No-Op filter entry used to explicitly exit a set of filter entries without modifying previous actions.

Multiple action statements entered will overwrite previous actions.

The **no** form of this command removes the specified action statement.

Default

Action specified by the **default-action** command will apply.

Parameters

drop

Specifies packets matching the entry criteria will be dropped.

forward

Specifies packets matching the entry criteria will be forwarded.

Platforms

7705 SAR Gen 2

action

Syntax

action {**drop** | **forward**}

no action

Context

[Tree] (config>log>filter>entry action)

Full Context

configure log filter entry action

Description

This command specifies a drop or forward action associated with the filter entry. If neither **drop** nor **forward** is specified, the **default-action** will be used for traffic that conforms to the match criteria. This could be considered a No-Op filter entry used to explicitly exit a set of filter entries without modifying previous actions.

Multiple action statements entered will overwrite previous actions.

The **no** form of this command removes the specified **action** statement.

Default

no action

Parameters

drop

Specifies packets matching the entry criteria will be dropped.

forward

Specifies packets matching the entry criteria will be forwarded.

Platforms

7705 SAR Gen 2

action

Syntax

action {**log-only** | **reset-mda** | **fail-mda**}

no action

Context

[Tree] (config>card>mda>event action)

Full Context

configure card mda event action

Description

This command defines the action to be taken when a specific hardware error event is raised against the target mda.

Only one action can be enabled at a time. Entering a new action will override a previously defined action.

The **no** form of this command sets the action to the default value.

Default

action log-only

Parameters

log-only

Specifies to pass the log event to log management. No other action is taken.

reset-mda

Specifies to reset the mda.

fail-mda

Specifies to set the operational state of the mda to Failed. This Failed state will persist until the clear mda command is issued (reset) or the mda is removed and re-inserted (re-seat).

Platforms

7705 SAR Gen 2

action

Syntax

[**no**] **action**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization action)

Full Context

configure system security profile netconf base-op-authorization action

Description

This command enables the NETCONF <action> RPC.

The **no** form of this command disables the RPC.

Default

no action

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

action

Syntax

[**no**] **action** [**secondary**]

Context

[Tree] (config>filter>ip-filter>entry action)

[Tree] (config>filter>ipv6-filter>entry action)

Full Context

configure filter ip-filter entry action

configure filter ipv6-filter entry action

Description

Commands in this context configure a primary (no option specified) or secondary (**secondary** option specified) action to be performed on packets matching this filter entry. An ACL filter entry remains inactive (is not programmed in hardware) until a specific action is configured for that entry.

A primary action supports any filter entry action, a secondary action is used for redundancy and defines a redundant Layer 3 PBR action for an Layer 3 PBR primary action or a redundant L2 PBF action for a Layer 2 PBF primary action.

The **no** form of this command removes the specific action configured in the context of the action command. The primary action cannot be removed if a secondary action exists.

Default

no action

Parameters

secondary

Specifies a secondary action to be performed on packets matching this filter entry. A secondary action can only be configured if a primary action is configured.

Platforms

7705 SAR Gen 2

action

Syntax

action [**fc** *fc-name*] [**priority** {**high** | **low**}] [**policer** *policer-id*]

no action

Context

[Tree] (config>qos>sap-ingress>ipv6-criteria>entry action)

[Tree] (config>qos>sap-ingress>ip-criteria>entry action)

[Tree] (config>qos>sap-ingress>mac-criteria>entry action)

Full Context

configure qos sap-ingress ipv6-criteria entry action

configure qos sap-ingress ip-criteria entry action

configure qos sap-ingress mac-criteria entry action

Description

This mandatory command associates the forwarding class or enqueuing priority with specific IP, IPv6, or MAC criteria entry ID. The action command supports setting the forwarding class parameter to a subclass. Packets that meet all match criteria within the entry have their forwarding class and enqueuing priority overridden based on the parameters included in the **action** parameters. When the forwarding class is not specified in the **action** command syntax, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueuing priority is not specified in the action, a matching packet preserves (or inherits) the existing enqueuing priority derived from earlier matches in the classification hierarchy.

When a policer is specified in the action, a matching packet is directed to the configured policer instead of the policer/queue assigned to the forwarding class of the packet.

The **action** command must be executed for the match criteria to be added to the active list of entries. If the entry is designed to prevent more explicit (higher entry ID) entries from matching certain packets, the **fc** *fc-name* and **match** *protocol* fields should not be defined when executing action. This allows packets matching the entry to preserve the forwarding class and enqueuing priority derived from previous classification rules.

Each time action is executed on a specific entry ID, the previously entered values for **fc** *fc-name* and **priority** are overridden with the newly defined parameters or inherit previous matches when a parameter is omitted.

The **no** form of this command removes the entry from the active entry list. Removing an entry on a policy immediately removes the entry from all SAPs using the policy. All previous parameters for the action is lost.

If no action is specified, the action specified by the **default-fc** command will be used.

Parameters

fc *fc-name*

The value given for **fc** *fc-name* must be one of the predefined forwarding classes in the system. Specifying the **fc** *fc-name* is required. When a packet matches the rule, the forwarding class is only overridden when the **fc** *fc-name* parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

The *subclass-name* parameter is optional and used with the *fc-name* parameter to define a pre-existing subclass. The *fc-name* and *subclass-name* parameters must be separated by a period (.). If *subclass-name* does not exist in the context of *fc-name*, an error will occur. If *subclass-name* is removed using the **no fc** *fc-name.subclass-name* **force** command, the **default-fc** command will automatically drop the *subclass-name* and only use *fc-name* (the parent forwarding class for the subclass) as the forwarding class.

Values

fc: *class*[.*subclass*]

class: be, l2, af, l1, h2, ef, h1, nc

subclass: 29 characters max

Default Inherit (When **fc** *fc-name* is not defined, the rule preserves the previous forwarding class of the packet.)

priority

The **priority** parameter overrides the default enqueuing priority for all packets received on a SAP using this policy that match this rule. Specifying the priority (**high** or **low**) is optional. When a packet matches the rule, the enqueuing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and priority is not explicitly defined in the rule, the enqueuing priority is inherited based on previous rule matches.

Default Inherit (When the **priority** (**high** or **low**) is not defined, the rule preserves the previous enqueuing priority of the packet)

high

The **high** parameter is used in conjunction with the **priority** parameter. Setting the **priority** enqueuing parameter to **high** for a packet increases the likelihood to enqueue the packet when the queue is congested. The enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the queue, the significance of the enqueuing priority is lost.

low

The **low** parameter is used in conjunction with the **priority** parameter. Setting the **priority** enqueuing parameter to **low** for a packet decreases the likelihood to enqueue the packet when the queue is congested. The enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

Default Inherit

policer-id

A valid policer-id must be specified. The parameter policer-id references a policer-id that has already been created within the sap-ingress QoS policy.

Values 1 to 63

Platforms

7705 SAR Gen 2

action

Syntax

action [**fc** *fc-name*] [**profile** {**in** | **out** | **exceed** | **inplus**}] [**policer** *policer-id*] [**port-redirect-group-queue**] [**queue** *queue-id*] [**use-fc-mapped-queue**]

no action

Context

[Tree] (config>qos>sap-egress>ip-criteria>entry action)

[Tree] (config>qos>sap-egress>ipv6-criteria>entry action)

Full Context

configure qos sap-egress ip-criteria entry action

configure qos sap-egress ipv6-criteria entry action

Description

This command defines the reclassification actions that should be performed on any packet matching the defined IP flow criteria within the entries match node. When defined under the **ip-criteria** context, the reclassification only applies to IPv4 packets. When defined under the **ipv6-criteria** context, the reclassification only applies to IPv6 packets.

If an egress packet on the SAP matches the specified IP flow entry, the forwarding class, or profile or egress queue accounting behavior may be overridden. By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions. Matching an IP flow reclassification entry will override all IP precedence- or DSCP-based reclassification rule actions when an explicit reclassification action is defined for the entry.

It is also possible to redirect the egress packet to a configured policer. The forwarding class or profile can also be optionally specified.

When an IP flow entry is first created, the entry will have no explicit behavior defined as the reclassification actions to be performed. In **show** and **info** commands, the entry will display no action as the specified reclassification action for the entry. When the entry is defined with no action, the entry will not be populated in the IP flow reclassification list used to evaluate packets egressing a SAP with the SAP egress policy defined. An IP flow reclassification entry is only added to the evaluation list when the action command for the entry is executed either with explicit reclassification entries or without any actions defined. Specifying action without any trailing reclassification actions allows packets matching the entry to exit the evaluation list without matching entries lower in the list. Executing no action on an entry removes the entry from the evaluation list and also removes any explicitly defined reclassification actions associated with the entry.

The **fc** keyword is optional. When specified, the egress classification rule will overwrite the forwarding class derived from ingress. The new forwarding class is used for egress remarking and queue mapping decisions.

The **profile** keyword is optional. When specified, the egress classification rule will overwrite the profile of the packet derived from ingress. The new profile value is used for egress remarking and queue congestion behavior.

The **policer** keyword is optional. When specified, the egress packet will be redirected to the configured policer. Optional parameters allow the user to control how the forwarded policed traffic exits the egress port. By default, the policed forwarded traffic will use a queue in the egress port's policer-output-queue queue group; alternatively, a queue in an instance of a user-configured queue group can be used or a local SAP egress queue.

The **no** form of this command removes all reclassification actions from the IP flow reclassification entry and also removes the entry from the evaluation list. An entry removed from the evaluation list will not be matched to any packets egress a SAP associated with the SAP egress QoS policy.

Parameters

fc *fc-name*

The fc reclassification action is optional. When specified, packets matching the IP flow reclassification entry will be explicitly reclassified to the forwarding class specified as fc-name regardless of the ingress classification decision. The fc-name defined must be one of the eight forwarding classes supported by the system. To remove the forwarding class reclassification action for the IP flow entry, the action command must be re-executed without the fc reclassification action defined.

Values

fc	class	
	class	be, l2, af, l1, h2, ef, h1, nc

profile {in | out | exceed | inplus}

The profile reclassification action is optional. When specified, packets matching the IP flow reclassification entry will be explicitly reclassified to the configured profile regardless of the ingress profiling decision. To remove the profile reclassification action for the IP flow reclassification entry, the action command must be re-executed without the profile reclassification action defined.

in

The **in** parameter is mutually exclusive to the **exceed**, **inplus**, and **out** parameters following the profile reclassification action keyword. **In, exceed**, **inplus**, or **out** must be

specified when the profile keyword is present. When **in** is specified, any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.

out

The **out** parameter is mutually exclusive to the **exceed**, **inplus**, and **in** parameters following the profile reclassification action keyword. **In, exceed**, **inplus**, or **out** must be specified when the profile keyword is present. When **out** is specified, any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.

exceed

The **exceed** parameter is mutually exclusive to the **out**, **inplus**, and **in** parameters following the profile reclassification action keyword. **In, exceed**, **inplus**, or **out** must be specified when the profile keyword is present. When **exceed** is specified, any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.

inplus

The **inplus** parameter is mutually exclusive to the **out**, **exceed**, and **in** parameters following the profile reclassification action keyword. **In, exceed**, **inplus**, or **out** must be specified when the profile keyword is present. When **inplus** is specified, any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

policer *policer-id*

When the action policer command is executed, a valid policer ID must be specified. The parameter policer ID references a policer ID that has already been created within the SAP egress QoS policy.

Values 1 to 63

port-redirect-group-queue queue *queue-id*

Used to override the forwarding class default egress queue destination to an egress port queue group. The specific egress queue group instance to use is specified at the time the QoS policy is applied to the SAP. Therefore, this parameter is only valid if SAP-based redirection is required. The queue parameter overrides the policer's default egress queue destination to a specified queue-id in the egress port queue group instance.

Values 1 to 8

queue *queue-id*

This parameter overrides the policer's default egress queue destination to a specified local SAP queue of that queue-id. A queue of ID queue-id must exist within the egress QoS policy.

Values 1 to 8

use-fc-mapped-queue

This parameter overrides the policer's default egress queue destination to the queue mapped by the traffic's forwarding class.

Platforms

7705 SAR Gen 2

action

Syntax

action [**fc** *fc-name* **profile** {**in** | **out** | **exceed** | **inplus**}] [**port-redirect-group** {**queue** *queue-id* | **policer** *policer-id* [**queue** *queue-id*]}]

Context

[Tree] (config>qos>network>egress>ip-criteria>entry action)

[Tree] (config>qos>network>egress>ipv6-criteria>entry action)

Full Context

configure qos network egress ip-criteria entry action

configure qos network egress ipv6-criteria entry action

Description

This command defines the reclassification actions that are performed on any packet matching the defined IP flow criteria within the entry's matched node. When defined under the **ip-criteria** context, the reclassification only applies to IPv4 packets. When defined under the **ipv6-criteria** context, the reclassification only applies to IPv6 packets.

If an egress packet matches the specified IP flow entry, the forwarding class and profile may be overridden. By default, the forwarding class and profile of the packet are derived from ingress classification and profiling functions. Matching an IP flow reclassification entry will override all IP precedence-based or DSCP-based reclassification rule actions when an explicit reclassification action is defined for the entry.

When an IP flow entry is first created, the entry will have no explicit behavior defined as the reclassification actions to be performed. When the entry is defined with no action, the entry will not be populated in the IP flow reclassification list used to evaluate egress packets. An IP flow reclassification entry is only added to the evaluation list when the action command for the entry is executed.

The **fc** and **profile** keywords are optional. When specified, the egress classification rule will overwrite the forwarding class and profile derived from ingress. The new forwarding class and profile are used for egress remarking, queue mapping decisions, and queue congestion behavior.

The **port-redirect-group** keyword is optional. When specified, the egress packet will be redirected to the configured queue or policer in the specified egress network queue group. By default, the policed forwarded traffic will use the regular network queue to which the packet's forwarding class is mapped. Alternatively, a queue in the network egress queue group instance can be used for post-policed traffic by specifying a queue after the **policer** parameter. The **port-redirect-group** keyword requires that the network egress queue group instance is specified when this network QoS policy is applied to a network interface.

The **no** form of this command removes all reclassification actions from the IP flow reclassification entry and also removes the entry from the evaluation list. An entry removed from the evaluation list will not be matched to any egress packets.

Default

no action

Parameters

fc *fc-name*

The fc reclassification action is optional. When specified, packets matching the IP flow reclassification entry will be explicitly reclassified to the forwarding class specified as *fc-name* regardless of the ingress classification decision. The *fc-name* defined must be one of the eight forwarding classes supported by the system. The profile reclassification action is mandatory when an **fc** is specified. To remove the forwarding class reclassification action for the IP flow entry, the action command must be re-executed without the fc reclassification action defined.

Values be, l2, af, l1, h2, ef, h1, nc

profile {in | out | exceed | inplus}

The profile reclassification action is mandatory when an fc is specified, otherwise it is optional. When specified, packets matching the IP flow reclassification entry will be explicitly reclassified to the configured profile regardless of the ingress profiling decision. **In**, **exceed**, **inplus**, or **out** must be specified when the profile keyword is present. To remove the profile reclassification action for the IP flow reclassification entry, the action command must be re-executed without the profile reclassification action defined.

in

When specified, any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.

out

When specified, any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.

exceed

When specified, any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.

inplus

When specified, any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

queue *queue-id*

Used to override the forwarding class default egress queue destination to the specified network egress queue group instance queue. The specific egress queue group instance to use is specified at the time the QoS policy is applied to the network interface.

Values 1 to 8

policer *policer-id*

Specifies a valid policer ID that has already been created within the network egress queue group instance.

Values 1 to 16

queue *queue-id*

The queue following the configured policer overrides the default policed traffic egress queue destination to a specified queue in the network egress queue group instance.

Values 1 to 8

Platforms

7705 SAR Gen 2

action

Syntax

action fc *fc-name* **profile** {**in** | **out**}

no action

Context

[Tree] (config>qos>network>ingress>ipv6-criteria>entry action)

[Tree] (config>qos>network>ingress>ip-criteria>entry action)

Full Context

configure qos network ingress ipv6-criteria entry action

configure qos network ingress ip-criteria entry action

Description

This command defines the reclassification actions that are performed on any packet matching the defined IP flow criteria within the entry's matched node. When defined under the **ip-criteria** context, the reclassification only applies to IPv4 packets. When defined under the **ipv6-criteria** context, the reclassification only applies to IPv6 packets.

If an ingress packet matches the specified IP flow entry, the forwarding class and profile may be overridden. By default, the forwarding class and profile of the packet are derived from ingress classification and profiling functions. Matching an IP flow reclassification entry will override all non-criteria reclassification rule actions when an explicit reclassification action is defined for the entry.

When an IP flow entry is first created, the entry will have no explicit behavior defined as the reclassification actions to be performed. When the entry is defined with no action, the entry will not be populated in the IP flow reclassification list used to evaluate ingress packets. An IP flow reclassification entry is only added to the evaluation list when the action command for the entry is executed.

The **no** form of this command removes all reclassification actions from the IP flow reclassification entry and also removes the entry from the evaluation list. An entry removed from the evaluation list will not be matched to any ingress packets.

Default

no action

Parameters

fc *fc-name*

The fc reclassification action is optional. When specified, packets matching the IP flow reclassification entry will be explicitly reclassified to the forwarding class specified as *fc-*

name regardless of the ingress classification decision. The *fc-name* defined must be one of the eight forwarding classes supported by the system. The profile reclassification action is mandatory when an **fc** is specified. To remove the forwarding class reclassification action for the IP flow entry, the action command must be re-executed without the fc reclassification action defined.

Values be, l2, af, l1, h2, ef, h1, nc

profile {in | out}

The profile reclassification action is mandatory. Packets matching the IP flow reclassification entry will be explicitly reclassified to the configured profile regardless of other ingress profiling decisions. **In** or **out** must be specified when the profile keyword is present. To remove the profile reclassification action for the IP flow reclassification entry, the action command must be re-executed without the profile reclassification action defined.

in

When specified, any packets matching the reclassification rule will be treated as in-profile by the ingress forwarding plane.

out

When specified, any packets matching the reclassification rule will be treated as out-of-profile by the ingress forwarding plane.

Platforms

7705 SAR Gen 2

action

Syntax

action {**replace** | **drop** | **keep**}

no action

Context

[Tree] (config>router>if>dhcp>option action)

Full Context

configure router interface dhcp option action

Description

This command configures the processing required when the router receives a DHCP request that already has a Relay Agent Information Option (Option 82) field in the packet.

The **no** form of this command returns the system to the default value.

Default

Per RFC 3046, *DHCP Relay Agent Information Option*, section 2.1.1, *Reforwarded DHCP requests*, the default is to keep the existing information intact. The exception to this is if the GI address of the received

packet is the same as the ingress address on the router. In that case the packet is dropped and an error is logged.

Parameters

replace

In the upstream direction (from the user), the existing Option 82 field is replaced with the Option 82 field from the router. In the downstream direction (toward the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop

The packet is dropped, and an error is logged.

keep

The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on toward the client.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.

Platforms

7705 SAR Gen 2

action

Syntax

action {**permit** | **deny** | **deny-host-unreachable**}

no action

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry action)

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter>entry action)

[Tree] (config>system>security>mgmt-access-filter>ip-filter>entry action)

Full Context

configure system security management-access-filter mac-filter entry action

configure system security management-access-filter ipv6-filter entry action

configure system security management-access-filter ip-filter entry action

Description

This command creates the action associated with the management access filter match criteria entry.

The **action** keyword is required. If no **action** is defined, the filter is ignored. If multiple action statements are configured, the last one overwrites previous configured actions.

If the packet does not meet any of the match criteria the configured **default action** is applied.

Parameters

permit

Specifies that packets matching the configured criteria will be permitted.

deny

Specifies that packets matching the configured selection criteria will be denied and that a ICMP host unreachable message will not be issued.

deny-host-unreachable

Specifies that packets matching the configured selection criteria will be denied and that a host unreachable message will not be issued.

The **deny-host-unreachable** parameter only applies to ip-filter and ipv6-filter.

Platforms

7705 SAR Gen 2

action

Syntax

action {**deny** | **permit** | **read-only**}

Context

[Tree] (config>system>security>profile>entry action)

Full Context

configure system security profile entry action

Description

This command configures the action associated with the profile entry.

Parameters

deny

Specifies that commands matching the entry command match criteria are to be denied.

permit

Specifies that commands matching the entry command match criteria is permitted.

read-only

Specifies the commands matching the entry command match criteria is available with read-only access.

Platforms
7705 SAR Gen 2

4.23 action-list

action-list

Syntax
action-list

Context
[Tree] (config>log>event-handling>handler action-list)

Full Context
configure log event-handling handler action-list

Description
Commands in this context configure the EHS handler action list.

Platforms
7705 SAR Gen 2

4.24 activate

activate

Syntax
activate [*file-url*] [**now**]

Context
[Tree] (admin>system>license activate)

Full Context
admin system license activate

Description
This command performs an activation on the license file pointed to by the command line argument. The file is first validated as described in the **admin>system>license>validate** command and upon success, replaces the existing license attributes in the system with the information in the new license file.

The license attributes that are active on a system can be viewed with the **show>licensing>entitlements** command.

Note:
If the CLM tool is being used for license management, it shall perform the validation and activation and there is no need to enter these commands manually.

Parameters

file-url

Specifies the file URL location to read the license file.

Values local-url, remote-url

now

If the **now** keyword is not present, the operator is prompted to confirm the activation. With the **now** keyword the license file is activated without the additional prompt.

Platforms

7705 SAR Gen 2

activate

Syntax

activate card *cpm-slot* **serial-number** *cpm-serial-number* **confirmation-code** *code*

Context

[Tree] (admin>system>security>secure-boot activate)

Full Context

admin system security secure-boot activate

Description

This command activates Secure Boot to enforce digital signature verification of the software on every boot.

Once Secure Boot is activated on a CPM, the capability is permanently enabled and cannot be disabled.

Parameters

cpm-slot

Specifies the CPM slot.

Values A,B

cpm-serial-number

Specifies the CPM serial number, up to 256 characters.

code

Specifies the secure boot confirmation code, up to 32 characters.

Platforms

7705 SAR Gen 2

4.25 active-hold-delay

active-hold-delay

Syntax

active-hold-delay *active-hold-delay*

no active-hold-delay

Context

[Tree] (config>service>epipe>endpoint active-hold-delay)

Full Context

configure service epipe endpoint active-hold-delay

Description

This command specifies that the node will delay sending the change in the T-LDP status bits for the VLL endpoint when the MC-LAG transitions the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby** or when any object in the endpoint. For example, SAP, ICB, or regular spoke SDP, transitions from up to down operational state.

By default, when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby**, the node sends immediately new T-LDP status bits indicating the new value of "standby" over the spoke SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

There is no delay applied to the VLL endpoint status bit advertisement when the MC-LAG transitions the LAG subgroup which hosts the SAP from **standby** to **active** or when any object in the endpoint transitions to an operationally up state.

Default

active-hold-delay 0

Parameters

active-hold-delay

Specifies the active hold delay in 100s of milliseconds.

A value of zero means that when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby,** the node sends immediately new T-LDP status bits indicating the new value of **standby** over the spoke SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

Values 0 to 60

Platforms

7705 SAR Gen 2

4.26 active-mda-limit

active-mda-limit

Syntax

active-mda-limit *number*

no active-mda-limit

Context

[Tree] (config>isa>nat-group active-mda-limit)

Full Context

configure isa nat-group active-mda-limit

Description

This command configures the number of active ESA-VM or ISA members in a NAT group.

The system automatically selects which ESA-VMs or ISAs are active. In active/standby (A/S) redundancy mode, the correlation between ESA-VM or ISA members is direct, meaning each ESA-VM or ISA equates to one member. In active/active (A/A) redundancy mode, an individual ESA-VM or ISA may be associated with multiple members.

For A/S redundancy, any surplus ESA-VMs or ISAs beyond the configured active threshold automatically transition to standby. These standby units remain idle until an active unit fails, at which point a standby unit takes over, handling traffic from only one failed active unit. This setup allows for the configuration of multiple standby units to provide resilience against several concurrent failures.

In A/A redundancy, the combination of this command and the **failed-mda-limit** command guides the distribution of resources among ESA-VMs or ISAs, essentially defining how the members are structured.

In both A/S and A/A modes, the system strives to maintain the configured number of active members as outlined by the active MDA limit, drawing from the pool of available spare resources to compensate for any failures. If the actual number of active members drops below this limit because of a lack of available spares, the NAT group status changes to degraded. In this state, traffic intended for the missing ESA-VM or ISA members (up to the active MDA limit) is blackholed. In Layer 2-aware NAT this condition can be circumvented where traffic can bypass NAT altogether and be directly routed within the internal network that may have an alternate path to a backup NAT system. For additional details, see "L2-Aware bypass" in the *7705 SAR Gen 2 Multiservice ISA and ESA Guide*

The **no** form of this command removes the active MDA limit configuration.

Default

no active-mda-limit

Parameters

number

Specifies the active MDA limit.

Values 1 to 28

Platforms

7705 SAR Gen 2

4.27 active-mda-number

active-mda-number

Syntax

active-mda-number *number*

no active-mda-number

Context

[Tree] (config>isa>tunnel-grp active-mda-number)

Full Context

configure isa tunnel-group active-mda-number

Description

This command specifies the number of active MS-ISA within all configured MS-ISA in the tunnel-group with multi-active enabled. IPsec traffic will be load balanced across all active MS-ISAs. If the number of configured MS-ISA is greater than the active-mda-number then the delta number of MS-ISA will be backup.

Default

active-mda-number 1

Parameters

number

Specifies the number of active MDAs.

Values 1 to 16

Platforms

7705 SAR Gen 2

4.28 active-outbound-sa

active-outbound-sa

Syntax

active-outbound-sa *spi*

no active-outbound-sa

Context

[Tree] (config>grp-encryp>encryp-keygrp active-outbound-sa)

Full Context

configure group-encryption encryption-keygroup active-outbound-sa

Description

This command specifies the Security Association, referenced by the Security Parameter Index (SPI), to use when performing encryption and authentication on NGE packets egressing the node for all services configured using this key group.

The **no** form of the command returns the parameter to its default value and is the same as removing this key group from all outbound direction key groups in all services configured with this key group (that is, all packets of services using this key group will egress the node in without being encrypted).

Parameters

spi

Specifies the SPI to use for packets of services using this key group when egressing the node.

Values 1 to 127

Platforms

7705 SAR Gen 2

4.29 active-psk

active-psk

Syntax

active-psk *active-pre-shared-key*

no active-psk

Context

[Tree] (config>macsec>conn-assoc>static-cak active-psk)

Full Context

configure macsec connectivity-association static-cak active-psk

Description

This command specifies the active transmitting pre-shared-key. If two pre-shared-keys are configured, the arriving MACsec MKA can be decrypted via CAKs of both pre-shared keys; however, only the active-psk will be used for TX encryption of MKA PDUs.

Default

active-psk 1

Parameters

active-pre-shared-key

Specifies the value of the pre-shared-key.

Values 1 or 2

Platforms

7705 SAR Gen 2

4.30 ad-per-evi-routes

ad-per-evi-routes

Syntax

ad-per-evi-routes

Context

[Tree] (config>service>system>bgp-evpn ad-per-evi-routes)

Full Context

configure service system bgp-evpn ad-per-evi-routes

Description

Commands in this context configure how Ethernet AD per-EVI routes are generated.

Default

ad-per-evi-routes

Platforms

7705 SAR Gen 2

4.31 ad-validation

ad-validation

Syntax

ad-validation {**fall-through** | **drop**}

no ad-validation

Context

[Tree] (config>system>dns>dnssec ad-validation)

Full Context

configure system dns dnssec ad-validation

Description

This command enables validation of the presence of the AD-bit in responses from the DNS servers, and reports a warning to the SECURITY log if DNSSEC validation was not possible.

This command requires either the fall-through or drop parameters be configured. When the fall-through parameter is supplied, the system will allow DNS responses that do not pass DNSSEC validation to be accepted and logged. When the drop parameter is specified, the system will reject and log DNS responses that do not pass DNSSEC validation and the resolution will appear to fail.

Default

no ad-validation

Parameters

fall-through

Specifies that the DNSSEC validator should allow non-DNSSEC responses to fall-through to permit resolution in case of validation failure.

drop

Specifies that the DNSSEC validator should drop non-DNSSEC responses in case of validation failure.

Platforms

7705 SAR Gen 2

4.32 adapt-qos

adapt-qos

Syntax

adapt-qos {**link** | **port-fair** | **distribute** [**include-egr-hash-cfg**]}

Context

[Tree] (config>lag>access adapt-qos)

Full Context

configure lag access adapt-qos

Description

This command specifies how the LAG SAP queue and virtual scheduler buffering and rate parameters are adapted over multiple active XMAs/MDAs. This command applies only to access LAGs.

Default

adapt-qos distribute

Parameters

link

Specifies that the LAG will create the SAP queues and virtual schedulers with the actual parameters on each LAG member port.

port-fair

Places the LAG instance into a mode that enforces QoS bandwidth constraints in the following manner:

- all egress QoS objects associated with the LAG instance are created on a per port basis
- bandwidth is distributed over these per port objects based on the proportion of the port's bandwidth relative to the total of all active ports bandwidth within the LAG
- the **include-egr-hash-cfg** behavior is automatically enabled allowing the system to detect objects that hash to a single egress link in the lag and enabling full bandwidth for that object on the appropriate port

distribute

Creates an additional internal virtual scheduler per IOM/XCM as parent of the configured SAP queues and virtual schedulers per LAG member port on that IOM/XCM. This internal virtual scheduler limits the total amount of egress bandwidth for all member ports on the IOM/XCM to the bandwidth specified in the egress qos policy.

include-egr-hash-cfg

Specifies whether explicitly configured hashing should factor into the egress buffering and rate distribution.

When this parameter is configured, all SAPs on this LAG which have explicit hashing configured, the egress HQoS and HPol (including queues, policers, schedulers and arbiters) will receive 100% of the configured bandwidth (essentially operating in adapt-qos link mode). For any Multi-Service-Sites assigned to such a LAG, bandwidth will continue to be divided according to adapt-qos distribute mode.

A LAG instance that is currently in adapt-qos link mode may be placed at any time in port-fair mode. Similarly, a LAG instance that is currently in adapt-qos port-fair mode may be placed at any time in link mode. However, a LAG instance in adapt-qos distribute mode may not be placed into port-fair (or link) mode while QoS objects are associated with the LAG instance. To move from distribute to port-fair mode it is necessary to remove all QoS objects from the LAG instance.

Platforms

7705 SAR Gen 2

4.33 adaptation-rule

adaptation-rule

Syntax

adaptation-rule [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]

Context

[Tree] (config>service>ies>if>sap>egress>queue-override>queue adaptation-rule)

[Tree] (config>service>vpls>sap>egress>queue-override>queue adaptation-rule)

[Tree] (config>service>vpls>sap>ingress>queue-override>queue adaptation-rule)

[Tree] (config>service>ies>if>sap>ingress>queue-override>queue adaptation-rule)

Full Context

configure service ies interface sap egress queue-override queue adaptation-rule

configure service vpls sap egress queue-override queue adaptation-rule

configure service vpls sap ingress queue-override queue adaptation-rule

configure service ies interface sap ingress queue-override queue adaptation-rule

Description

This command overrides specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the adaptation rule is performed under the **hs-wrr-group** within the SAP egress QoS policy.

The **no** form of this command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default

no adaptation-rule

Parameters

pir

Specifies the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir

Specifies the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id* **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule

Specifies the CIR and PIR adaptation rules.

Values max — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue is equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue is equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue is the rate closest to the rate specified using the **rate** command.

Platforms

7705 SAR Gen 2

adaptation-rule

Syntax

adaptation-rule [**pir** *adaptation-rule*] [**cir** {**max** | **min** | **closest**}]

no adaptation-rule

Context

[Tree] (config>port>ethernet>network>egr>qgrp>qover>q adaptation-rule)

[Tree] (config>port>ethernet>access>ing>qgrp>qover>q adaptation-rule)

[Tree] (config>port>ethernet>access>egr>qgrp>qover>q adaptation-rule)

Full Context

configure port ethernet network egress queue-group queue-overrides queue adaptation-rule

configure port ethernet access ingress queue-group queue-overrides queue adaptation-rule

configure port ethernet access egress queue-group queue-overrides queue adaptation-rule

Description

This command specifies the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case the configuration of the adaptation rule is performed under the **hs-wrr-group** within the egress queue group template.

The **no** form of this command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default

adaptation-rule pir closest cir closest

Parameters

pir

Defines the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir

Defines the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id* **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule

Specifies the adaptation rule to be used while computing the operational CIR or PIR value.

Values **max** — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

Platforms

7705 SAR Gen 2

adaptation-rule

Syntax

adaptation-rule [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]

no adaption-rule

Context

[Tree] (config>service>epipe>sap>ingress>queue-override>queue adaptation-rule)

[Tree] (config>service>epipe>sap>egress>queue-override>queue adaptation-rule)

Full Context

configure service epipe sap ingress queue-override queue adaptation-rule

configure service epipe sap egress queue-override queue adaptation-rule

Description

This command can be used to override specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case the configuration of the adaptation rule is performed under the hs-wrr-group within the SAP egress QoS policy.

The **no** form of this command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default

no adaptation-rule

Parameters

pir

The **pir** parameter defines the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir

The **cir** parameter defines the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id* **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule

Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.

Values **max** — The **max** (maximum) keyword is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) keyword is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

Platforms

7705 SAR Gen 2

adaptation-rule

Syntax

adaptation-rule [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]

no adaptation-rule

Context

[Tree] (config>service>vprn>if>sap>ingress>queue-override>queue adaptation-rule)

[Tree] (config>service>vprn>if>sap>egress>queue-override>queue adaptation-rule)

Full Context

configure service vprn interface sap ingress queue-override queue adaptation-rule

configure service vprn interface sap egress queue-override queue adaptation-rule

Description

This command can be used to override specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the adaptation rule is performed under the **hs-wrr-group** within the SAP egress QoS policy.

The **no** form of this command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default

no adaptation-rule

Parameters

pir

The **pir** parameter defines the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir

The **cir** parameter defines the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id* **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule

Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.

Values **max** — The **max** (maximum) keyword is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) keyword is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

Platforms

7705 SAR Gen 2

4.34 adaptive

adaptive

Syntax

[**no**] **adaptive**

Context

[Tree] (config>router>mpls>lsp adaptive)

[Tree] (config>router>mpls>lsp>secondary adaptive)

[Tree] (config>router>mpls>lsp>primary adaptive)

[Tree] (config>router>mpls>lsp-template adaptive)

Full Context

configure router mpls lsp adaptive

configure router mpls lsp secondary adaptive

configure router mpls lsp primary adaptive

configure router mpls lsp-template adaptive

Description

This command enables the make-before-break functionality for an LSP or LSP path. When enabled for the LSP, make-before-break will be performed for primary path and all the secondary paths of the LSP.

Default

adaptive

Platforms

7705 SAR Gen 2

4.35 add-paths

add-paths

Syntax

[**no**] **add-paths**

Context

[Tree] (config>router>bgp add-paths)

[Tree] (config>router>bgp>group add-paths)

[Tree] (config>router>bgp>group>neighbor add-paths)

Full Context

configure router bgp add-paths

configure router bgp group add-paths

configure router bgp group neighbor add-paths

Description

This command allows the add-paths node to be the configured for one or more families of the BGP instance, a group or a neighbor. The BGP add-paths capability allows the router to send and/or receive multiple paths per prefix to/from a peer. The add-paths command without additional parameters is equivalent to removing Add-Paths support for all address families, which causes sessions that previously negotiated the add-paths capability for one or more address families to go down and come back up without the add-paths capability.

The no form of this command (no add-paths) removes add-paths from the configuration of BGP, the group or the neighbor, causing sessions established using add-paths to go down and come back up without the add-paths capability.

Default

no add-paths

Platforms

7705 SAR Gen 2

4.36 add-paths-send-limit

add-paths-send-limit

Syntax

add-paths-send-limit *send-limit*

no add-paths-send-limit

Context

[Tree] (config>router>policy-options>policy-statement>default-action add-paths-send-limit)

Full Context

configure router policy-options policy-statement default-action add-paths-send-limit

Description

This command sets the *send-limit* to a specific value for all routes matched by the policy entry or default action. Add-paths allows a BGP router to send multiple paths for the same NLRI/prefix to a peer

advertising the add-paths receive capability. The *send-limit* dictates the maximum number of paths that can be advertised.

The default *send-limit* is controlled by the instance, group or neighbor level configuration and applies to all prefixes in a particular address family. Using route policies allows the default send-limit to be overridden to use a larger or smaller maximum value on a per-prefix basis. For example, if, for most prefixes advertised to a peer, at most 1 path should be advertised but for a few exceptional prefixes up to 4 paths should be advertised, then the neighbor-level *send-limit* can be set to a value of 1 and the **add-paths-send-limit** in the policy entry that matches the exceptional routes can be set to a value of 4.

Default

no add-paths-send-limit

Parameters

send-limit

Specifies the maximum number of paths to advertise for matched routes to an Add-Paths peer. If the value is **multipaths**, then BGP advertises all of the used BGP multipaths for each matched route that is the best path for its prefix (NLRI). Add paths can be advertised only if the peer has signaled support for receiving multiple add paths.

Values 1 to 16, none, multipaths

Platforms

7705 SAR Gen 2

4.37 add-to-received-bgp

add-to-received-bgp

Syntax

add-to-received-bgp *weight*

no add-to-received-bgp

Context

[Tree] (config>service>vprn>bgp>group>evpn-link-bandwidth add-to-received-bgp)

[Tree] (config>service>vprn>bgp>group>neighbor>evpn-link-bandwidth add-to-received-bgp)

Full Context

configure service vprn bgp group evpn-link-bandwidth add-to-received-bgp

configure service vprn bgp group neighbor evpn-link-bandwidth add-to-received-bgp

Description

This command configures the weight value added to all BGP PE-CE routes for the purpose of weighted ECMP if EVPN-IFL and BGP PE-CE routes are combined into the same ECMP set.

For the load-balancing between EVPN-IFL and BGP PE-CE routes the **configure service vprn bgp eibgp-loadbalance** command must already be configured on the system.

The **no** form of this command disables the weight value added to all BGP PE-CE routes.

Default

no add-to-received-bgp

Parameters

weight

Specifies the weight value added to all BGP PE-CE routes.

Values 1 to 128

Platforms

7705 SAR Gen 2

4.38 add-to-received-ebgp

add-to-received-ebgp

Syntax

add-to-received-ebgp *family* [*family*]

no add-to-received-ebgp

Context

[Tree] (config>service>vprn>bgp>group>link-bandwidth add-to-received-ebgp)

[Tree] (config>service>vprn>bgp>group>neighbor>link-bandwidth add-to-received-ebgp)

Full Context

configure service vprn bgp group link-bandwidth add-to-received-ebgp

configure service vprn bgp group neighbor link-bandwidth add-to-received-ebgp

Description

This command configures BGP to automatically add a link-bandwidth extended community to every route received from a directly connected (single-hop) EBGP peer within the scope of the command, as long as that route belongs to one of the listed address families.

The link-bandwidth extended community added by this command encodes the local-AS number of receiving BGP instance and the bandwidth of the interface to the directly connected EBGP peer.

Up to three families may be configured.

The **no** form of this command removes the link-bandwidth extended community added to received BGP routes.

Default

no add-to-received-ebgp

Parameters

family

Specifies the address families for which receiving the link-bandwidth extended community from EBGP peers should be supported.

Values ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes.

label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes.

ipv6 — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes.

Platforms

7705 SAR Gen 2

add-to-received-ebgp

Syntax

add-to-received-ebgp *family* [*family*]

no add-to-received-ebgp

Context

[Tree] (config>router>bgp>group>link-bandwidth add-to-received-ebgp)

[Tree] (config>router>bgp>group>neighbor>link-bandwidth add-to-received-ebgp)

Full Context

configure router bgp group link-bandwidth add-to-received-ebgp

configure router bgp group neighbor link-bandwidth add-to-received-ebgp

Description

This command configures BGP to automatically add a link-bandwidth extended community to every route received from a directly connected (single-hop) EBGP peer within the scope of the command, as long as that route belongs to one of the listed address families.

The link-bandwidth extended community added by this command encodes the local-AS number of receiving BGP instance and the bandwidth of the interface to the directly connected EBGP peer.

Up to six families may be configured.

The **no** form of this command removes the link-bandwidth extended community added to received BGP routes.

Default

no add-to-received-ebgp

Parameters

family

Specifies the address families for which receiving the link-bandwidth extended community from EBGP peers should be supported.

Values ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes.

label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes.

vpn-ipv4 — Adds a link-bandwidth extended community to IPv4 VPN (SAFI 128) routes.

ipv6 — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes.

label-ipv6 — Adds a link-bandwidth extended community to labeled-unicast IPv6 routes.

vpn-ipv6 — Adds a link-bandwidth extended community to IPv6 VPN (SAFI 128) routes.

Platforms

7705 SAR Gen 2

4.39 address

address

Syntax

address gi-address [**scope** *scope*]

address *ip-address*[*/prefix-length*]

address pool *pool-name* [**secondary-pool** *sec-pool-name*] [**delimiter** *delimiter*]

address use-pool-from-client [**delimiter** *delimiter*]

no address

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host address)

Full Context

configure subscriber-mgmt local-user-db ipoe host address

Description

This command configures how the IP address is defined for this host.

When the user database is used from a local DHCP server, then this command defines how to define the IP address the server offers to the DHCP-client.

When the user-db is used for PPPoE authentication, the **gi-address** parameter cannot be used. A fixed IP address causes PPPoE to use this IP address. If no IP address is specified, the PPPoE looks for IP address by other means (DHCP). If a pool name is given, this pool is sent in the DHCP request so it can be used in by the DHCP server to determine which address to give to the host.

The **no** form of this command causes no IP address to be assigned to this host. In a user database referred to from a local DHCP server, creating a host without address information causes the matching client never to get an IP address.

The **no** form of this command reverts to the default.

Parameters

gi-address

When specified, the gi-address of the DHCP message is taken to look for a subnet in the local DHCP server. The first available free address of the subnet is taken and "offered" to the host. When **local-user-db** is used for PPPoE authentication, this has the same result as **no address**.

ip-address

Specifies the fixed IP address to use for this host.

Values a.b.c.d

pool-name/sec-pool-name

Specifies the primary (and secondary) pool (in the local DHCP server), up to 32 characters, to look for an available address. The first available IP address from any subnet in the pool is used. When the local user database is used for PPPoE authentication, this causes the specified pool name to be sent to the DHCP server in a vendor-specific sub-option under Option 82.

use-pool-from-client

Use the pool-name in the Option 82 vendor-specific sub-option.

delimiter

Specifies a single ASCII character specifies the delimiter of separating primary and secondary pool names in option82 VSO.

Platforms

7705 SAR Gen 2

address

Syntax

address *ipv6-address/prefix-length* [**eui-64**] [**track-srrp** *srrp-instance*] [**modifier** *cga-modifier*] [**dad-disable**] [**primary-preference** *primary-preference*]

no address *ipv6-address/prefix-length*

Context

[Tree] (config>service>ies>if>ipv6 address)

[Tree] (config>service>vprn>if>ipv6 address)

Full Context

configure service ies interface ipv6 address

configure service vprn interface ipv6 address

Description

This command assigns an IPv6 address/subnet to the interface.

Up to 16 total primary and secondary IPv4 and IPv6 addresses can be assigned to network interfaces, and up to 256 to access interfaces.

Caution:
Configurations must not exceed 16 secondary IP addresses when IPsec, GRE, L2TPv3, or IP in IP protocols are active on an access interface.

The **no** form of this command removes the IPv6 address from the interface.

Parameters

ipv6-address/prefix-length

Specifies the IPv6 address on the interface.

Values	ipv6-address/prefix:	ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
			x:x:x:x:x:x:d.d.d.d
			x [0 to FFFF]H
			d [0 to 255]D
	prefix-length		1 to 128

eui-64

When the **eui-64** keyword is specified, a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from MAC address on Ethernet interfaces. For interfaces without a MAC address, for example ATM interfaces, the Base MAC address of the chassis is used.

srrp-instance

Specifies the SRRP instance ID that this interface route needs to track.

Values 1 to 4294967295

cga-modifier

Specifies the modifier in 32 hexadecimal nibbles.

Values 0x0–0xFFFFFFFF

dad-disable

Disables Duplicate Address Detection (DAD) and sets the address to preferred, even if there is a duplicated address.

primary-preference

Specifies a *primary-preference* index to an IPv6 address of the interface to enforce the order in which the address is used by control plane protocols and applications which require a fixed address of the interface. These include LDP and Segment Routing.

When originating packets from this interface, the source IPv6 address follows the selection rules in RFC 6724 except for the specific cases where a fixed address is required. In the latter case, the IPv6 address with the lowest primary-preference index is selected. If the selected address is removed, the system selects the IPv6 address with the next lowest *primary-preference* index.

The system assigns the next available index value to any IPv6 address of the interface when configured without the *primary-preference* index value specified. The address index space is unique across all addresses of a given interface.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

address

Syntax

address {*ip-address/mask* | *ip-address netmask*} [**broadcast {all-ones** | **host-ones}**] [**track-srrp** *srrp-instance*]

no address [*ip-address/mask* | *ip-address netmask*]

Context

[Tree] (config>service>ies>if address)

[Tree] (config>service>vprn>nw-if address)

[Tree] (config>service>vprn>if address)

Full Context

configure service ies interface address

configure service vprn network-interface address

configure service vprn interface address

Description

This command assigns an IP address, IP subnet, and broadcast address format to an IES IP router interface. Only one IP address can be associated with an IP interface. Use the **secondary** command to assign multiple addresses.

An IP address must be assigned to each IES or VPRN IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the router.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. The show commands display CIDR notation and is stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

Table 14: Address Admin and Operational States

Address	Admin State	Oper State
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an administratively up state and an address is assigned, it becomes operationally up and the protocol interfaces and the MPLS LSPs associated with that IP interface are reinitialized.

The **no** form of this command removes the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Parameters

ip-address

Specifies the IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

/

The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "*/*" and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.

mask-length

Specifies the subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address.

Note:
A mask length of 32 is reserved for loopback addresses (includes system addresses).

Default 0 to 31

mask

The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that is used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.254.

Note:
A mask of 255.255.255.255 is reserved for system IP addresses.

broadcast

Overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed.

This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) is received by the IP interface.

Default host-ones

all-ones

Specifies the broadcast address used by the IP interface for this IP address is 255.255.255.255, also known as the local broadcast.

host-ones

Specifies that the broadcast address used by the IP interface for this IP address is the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

srrp-instance

Tracks the specified SRRP instance state on the IPv6 address.

Platforms

7705 SAR Gen 2

address

Syntax

address *ip-address* [/*mask*] [*netmask*]

no address

Context

[Tree] (config>service>vpls>interface address)

Full Context

configure service vpls interface address

Description

This command assigns an IP address, IP subnet, and broadcast address format to an IES IP router interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each IES IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the router.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. The show commands display CIDR notation and is stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

Use the **no** form of this command to remove the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Table 15: Address Admin and Operational States

Address	Admin State	Oper State
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an administratively up state and an address is assigned, it becomes operationally up and the protocol interfaces and the MPLS LSPs associated with that IP interface will be reinitialized.

Parameters

ip-address

The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP netmask

The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 to 255.255.255.254. A mask of 255.255.255.255 is reserved for system IP addresses.

Platforms

7705 SAR Gen 2

address

Syntax

address *ip-address*

no address

Context

[Tree] (config>service>vprn>log>syslog address)

Full Context

configure service vprn log syslog address

Description

This command adds the syslog target host IP address to/from a syslog ID.

The *ip-address* parameter is mandatory. If no **address** is configured, syslog data cannot be forwarded to the syslog target host.

Only one address can be associated with a *syslog-id*. If multiple addresses are entered, the last address entered overwrites the previous address.

The same syslog target host can be used by multiple log IDs.

The **no** form of this command removes the syslog target host IP address.

Default

no address

Parameters

ip-address

Specifies the IP address of the syslog target host in dotted decimal notation.

Values		
	ipv4-address	a.b.c.d
	ipv6-address	x:x:x:x:x:x:x:x[-interface]
		x:x:x:x:x:x:d.d.d.d[-interface]
		x: [0 to FFFF]H
		d: [0 to 255]D
		interface: 32 characters maximum, mandatory for link local addresses

Platforms

7705 SAR Gen 2

address

Syntax

[**no**] **address** *ip-address*

Context

[Tree] (config>service>vprn>pim>rp>rp-candidate address)

[Tree] (config>service>vprn>pim>rp>bsr-candidate address)

Full Context

configure service vprn pim rp rp-candidate address

configure service vprn pim rp bsr-candidate address

Description

This command configures a static bootstrap or rendezvous point (RP) as long as the source is not directly attached to this router.

Use the **no** form of this command to remove the static RP from the configuration.

Default

No IP address is specified.

Parameters

ip-address

The static IP address of the RP. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values 1.0.0.0 to 223.255.255.255

Platforms

7705 SAR Gen 2

address

Syntax

[**no**] **address** *ipv6-address*

Context

[Tree] (config>service>vprn>pim>rp>ipv6>bsr-candidate address)

[Tree] (config>service>vprn>pim>rp>ipv6>rp-candidate address)

Full Context

configure service vprn pim rp ipv6 bsr-candidate address

configure service vprn pim rp ipv6 rp-candidate address

Description

This command configures a static bootstrap or rendezvous point (RP) as long as the source is not directly attached to this router.

Use the **no** form of this command to remove the static RP from the configuration.

Default

No IP address is specified.

Parameters

ipv6-address

The static IP address of the RP. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values

ipv6-address	: x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x [0 to FFFF]H

d [0 to 255]D

Platforms

7705 SAR Gen 2

address

Syntax

[**no**] **address** *ip-address*

Context

[Tree] (config>service>vprn>pim>rp>static address)

Full Context

configure service vprn pim rp static address

Description

This command configures the static rendezvous point (RP) address.

The **no** form of this command removes the static RP entry from the configuration.

Platforms

7705 SAR Gen 2

address

Syntax

address prefix *ip-prefix/ip-prefix-len*

address from *begin-ip-address* **to** *end-ip-address*

no address

Context

[Tree] (config>ipsec>ts-list>remote>entry address)

[Tree] (config>ipsec>ts-list>local>entry address)

Full Context

configure ipsec ts-list remote entry address

configure ipsec ts-list local entry address

Description

This command specifies the address range in the IKEv2 traffic selector.

Default

no address

Parameters

ip-prefix/ip-prefix-len

Specifies the IP prefix and subnet mask.

begin-ip-address

Specifies the beginning address of the range for this entry.

end-ip-address

Specifies the ending address of the range for this entry.

Platforms

7705 SAR Gen 2

address

Syntax

address *ip-address*

no address

Context

[Tree] (config>router>pim>rp>bsr-candidate address)

Full Context

configure router pim rp bsr-candidate address

Description

This command configures the candidate BSR IP address. This address is for Bootstrap router election.

The **no** form of this command removes the IP address from the BSR candidate configuration.

Default

no address

Parameters

ip-address

Specifies the IP host address used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values 1.0.0.0 – 223.255.255.255

Platforms

7705 SAR Gen 2

address

Syntax

address *ipv6-address*

no address

Context

[Tree] (config>router>pim>rp>ipv6>bsr-candidate address)

Full Context

configure router pim rp ipv6 bsr-candidate address

Description

This command configures the candidate BSR IPv6 address. This address is for Bootstrap router election.

The **no** form of this command removes the IPv6 address from the BSR candidate configuration.

Default

no address

Parameters

ipv6-address

Specifies the IPv6 host address used by the interface within the subnet.

Values

x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

Platforms

7705 SAR Gen 2

address

Syntax

address *ipv6-address*

no address

Context

[Tree] (config>router>pim>rp>ipv6>rp-candidate address)

Full Context

configure router pim rp ipv6 rp-candidate address

Description

This command configures the local IPv6 RP address. This address is sent in the RP candidate advertisements to the bootstrap router.

The **no** form of this command removes the IPv6 address from the RP candidate configuration.

Default

no address

Parameters

ipv6-address

Specifies the IPv6 RP address.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

prefix-length: 16 to 128

Platforms

7705 SAR Gen 2

address

Syntax

address *ip-address*

no address

Context

[Tree] (config>router>pim>rp>rp-candidate address)

Full Context

configure router pim rp rp-candidate address

Description

This command configures the local RP address. This address is sent in the RP candidate advertisements to the bootstrap router.

The **no** form of this command removes the IP address from the RP candidate configuration.

Default

no address

Parameters

ip-address

Specifies the *ip-address.*

Values 1.0.0.0 – 223.255.255.255

Platforms

7705 SAR Gen 2

address

Syntax

address *ip-address*

no address

Context

[Tree] (config>router>pim>rp>static address)

[Tree] (config>router>pim>rp>ipv6>static address)

Full Context

configure router pim rp static address

configure router pim rp ipv6 static address

Description

This command configures the Rendezvous Point (RP) address that should be used by the router for the range of multicast groups configured by the range command.

The **no** form of this command removes the IP address from the static configuration.

Parameters

ip-address

Specifies the static IP address of the RP. The *ip-address* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values 1.0.0.0 – 223.255.255.255

Platforms

7705 SAR Gen 2

address

Syntax

address {*ip-address*/*mask* | *ip-address netmask*} [**broadcast** {**all-ones** | **host-ones**}] [**track-srrp** *srrp-instance*] [**gre-termination**]

no address

Context

[Tree] (config>router>if address)

Full Context

configure router interface address

Description

This command assigns an IP address, IP subnet, and broadcast address format to an IP interface. Only one IP address can be associated with an IP interface. Use the **secondary** command to assign additional addresses.

An IP address must be assigned to each IP interface. An IP address and a mask combine to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the router.

From Release 19.10, The overlap restriction is not applicable for host-addresses configured on loopback interfaces. For example, a loopback interface addresses configured with mask of 32 or netmask of 255.255.255.255 can overlap with other prefixes on other IP interfaces in the same routing context within the router.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. **Show** commands display CIDR notation and are stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

The **no** form of this command removes the IP address assignment from the IP interface. Interface specific configurations for MPLS are also removed. This will operationally stop any MPLS LSPs that explicitly reference that IP address. When a new IP address is configured, interface specific configurations for MPLS need to be added. IEEE 1588 port based timestamping configured with **ptp-hw-assist** is also disabled.

Default

no address

Parameters

ip-address

Specifies the IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values 1.0.0.0 to 223.255.255.255

/

The forward slash is a parameter delimiter that separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "*/*" and the *mask* parameter. If a forward slash does not immediately follow the *ip-address*, a dotted decimal mask must follow the prefix.

mask

Specifies the subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask* parameter. The *mask* parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 1— 32. A mask length of 32 is reserved for system IP addresses.

Values 1 to 32

netmask

Specifies the subnet mask in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

broadcast

The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones,** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed.

This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) will be received by the IP interface.

Default host-ones

Values all-ones, host-ones

all-ones

The **all-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.

host-ones

Specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet

described by the *ip-address* and the *netmask* with all the host bits set to binary 1. This is the default broadcast address used by an IP interface.

srrp-instance

Specifies the SRRP instance ID that this interface route needs to track.

Values 1 to 4294967295

gre-termination

The optional **gre-termination** keyword allows GRE SDP tunnel packets to terminate on the router interface using the /31 value of the configured IP address. Refer to the *7705 SAR Gen 2 Services Overview Guide* for information about using **gre-termination**.

Platforms

7705 SAR Gen 2

address

Syntax

address *ipv6-address/prefix-length* [**eui-64**] [**track-srrp** *srrp-instance*] [**modifier** *cga-modifier*] [**dad-disable**] [**primary-preference** *primary-preference*]

no address *ipv6-address/prefix-length*

Context

[Tree] (config>router>if>ipv6 address)

Full Context

configure router interface ipv6 address

Description

This command assigns an IPv6 address to the interface. Up to 16 total primary and secondary IPv4 and IPv6 addresses can be assigned to network interfaces, and up to 256 to access interfaces.

Caution:
Configurations must not exceed 16 IPv6 addresses when IPsec, GRE, L2TPv3, or IP in IP protocols are active on an access interface.

A global IPv6 address together with the *prefix-length* create a locally configured interface IPv6 prefix and subnet. The defined global IP prefix must be unique within the context of a routing instance. It cannot overlap with any other existing global IP prefix defined on another IP interface within the same routing context in the router.

This overlap restriction is not applicable for IPv6 host addresses configured on loopback interfaces. For example, an IPv6 loopback host address configured upon a loopback interface may overlap with another prefix subnet configured on another IP interface within the same routing context.

Parameters

ipv6-address/prefix-length

Specifies the IPv6 address on the interface.

Values

ipv6-address/prefix-length:	ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x [0 to FFFF]H
		d [0 to 255]D
	prefix-length	1 to 128

eui-64

When the **eui-64** keyword is specified, a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from MAC address on Ethernet interfaces. For interfaces without a MAC address, for example POS interfaces, the Base MAC address of the chassis should be used.

srrp-instance

Indicates the unique identifier of the tracked SRRP instance.

Values 1 to 4294967295

cga-modifier

Sets the modifier for cryptographically-assigned addresses.

Values 0x0..0xFFFFFFFF...(32 hex nibbles)

dad-disable

Disables Duplicate Address Detection (DAD) and sets the address to preferred, even if there is a duplicated address.

primary-preference

Specifies a *primary-preference* index to an IPv6 address of the interface to enforce the order in which the address is used by control plane protocols and applications which require a fixed address of the interface. These include LDP and Segment Routing.

When originating packets from this interface, the source IPv6 address follows the selection rules in RFC 6724 except for the specific cases where a fixed address is required. In the latter case, the IPv6 address with the lowest primary-preference index is selected. If the selected address is removed, the system selects the IPv6 address with the next lowest *primary-preference* index.

The system assigns the next available index value to any IPv6 address of the interface when configured without the *primary-preference* index value specified. The address index space is unique across all addresses of a given interface.

Values 1 to 4294967295

srrp

Tracks the specified SRRP instance state on the IPv6 address.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

address

Syntax

[**no**] **address** *ip-prefix/ip-prefix-length* [**active** | **standby** | **standby/A** | **standby/B** | **standby/C** | **standby/D**]

Context

[Tree] (bof address)

Full Context

bof address

Description

This command assigns an IP address to the management Ethernet port on a CPM. The IP addresses are applied by the boot loader and the running image. The active and standby IP addresses must be on the same subnet.

An address must be assigned with the **active** keyword and for systems with a redundant CPM an additional address may be assigned with the **standby** keyword. The active address is used by the active CPM whether its CPM A or CPM B and the standby address, if specified, is used by the standby CPM whether its CPM B or CPM A.

Deleting a BOF address entry is not allowed from a remote session.

Note that changing the active and standby addresses without reboot standby CPM may cause a boot-env sync to fail.

The **no** form of this command deletes the IP address from the CPM Ethernet port.

Parameters

ip-prefix/ip-prefix-length

Specifies the destination address of the aggregate route in dotted decimal notation.

Values

ipv4-prefix	*a.b.c.d* (host bits must be 0)
ipv4-prefix-length	0 to 32
ipv6-prefix	*x:x:x:x:x:x:x:x* (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d

	x:	[0 to FFFF]H
	d:	[0 to 255]D
ipv6-prefix-length	0 to 128	

active | standby | standby/A | standby/B | standby/C | standby/D
specifies which CPM Ethernet address is being configured

Default active

Platforms
7705 SAR Gen 2

address

Syntax
address *ip-address*

no address

Context
[Tree] (config>log>syslog address)

Full Context
configure log syslog address

Description
This command adds the syslog target host IP address to/from a syslog ID.

This parameter is mandatory. If no **address** is configured, syslog data cannot be forwarded to the syslog target host.

Only one address can be associated with a *syslog-id*. If multiple addresses are entered, the last address entered overwrites the previous address.

The same syslog target host can be used by multiple log IDs.

The **no** form of this command removes the syslog target host IP address.

Default
no address

Parameters
ip-address
Specifies the IP address of the syslog target host in dotted decimal notation.

Values ipv4-address a.b.c.d

	ipv6-address	x:x:x:x:x:x:x:x[-interface] x:x:x:x:x:x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface: 32 characters maximum, mandatory for link local addressesipv6-address x:x:x:x:x:x:x:x[-interface] x:x:x:x:x:x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface: 32 characters maximum, mandatory for link local addresses

Platforms

7705 SAR Gen 2

address

Syntax

address *ip-address* [**port** *port*]

no address

Context

[Tree] (config>system>security>ldap>server address)

Full Context

configure system security ldap server address

Description

This command configures the IPv4 or IPv6 address for the LDAP server.

The **no** version of this command removes the server address.

Parameters

ip-address

The IP address of the LDAP server.

Values	ipv4-address	a.b.c.d (host bits must be 0)
	ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0..FFFF]H

d: [0..255]D

port

Specifies the port ID. The port is the LDAP server listening port; by default it is 389 but if the listening port on LDAP server is changed, this command needs to be configured accordingly.

Values 1 to 65535

Default 389

Platforms

7705 SAR Gen 2

4.40 address-map

address-map

Syntax

address-map *start-inside-ip-address* **to** *end-inside-ip-address* **subscriber-type** *nat-sub-type* **nat-policy** *nat-policy-name* [**create**]

no address-map *start-inside-ip-address* **to** *end-inside-ip-address* **subscriber-type** *nat-sub-type* **nat-policy** *nat-policy-name*

Context

[Tree] (config>router>nat>inside>deterministic address-map)

[Tree] (config>service>vprn>nat>inside>deterministic address-map)

Full Context

configure router nat inside deterministic address-map

configure service vprn nat inside deterministic address-map

Description

This command configures the mapping of the inside IP addresses of deterministic NAT44 subscribers to the outside IP addresses in a NAT pool. This mapping is applicable is applicable only to deterministic NAT44 with a single ESA-VM in a NAT-group. The number of subscribers per outside IP address is flexible and not restricted to a discrete range governed by the 2^n rule.

When configured, the **classic-lsn-max-subscriber-limit** command must be set to 1.

The **no** form of this command removes the configuration.

Parameters

start-inside-ip-address

Specifies the first IP address in the inside IP address range.

Values

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 32
ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6-prefix-length	0 to 128

end-inside-ip-address

Specifies the last IP address in the inside IP address range.

Values

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 32
ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6-prefix-length	0 to 128

nat-sub-type

Specifies the NAT subscriber type.

Values classic-lsn-sub, dslite-lsn-sub

nat-policy-name

Specifies the NAT policy name, up to 32 characters, that is referencing a NAT pool.

Values classic-lsn-sub, dslite-lsn-sub

create

Keyword used to create the address mapping.

Platforms

7705 SAR Gen 2

4.41 address-pref

address-pref

Syntax

address-pref {**ipv4-only** | **ipv6-first**}

no address-pref

Context

[Tree] (config>system>dns address-pref)

Full Context

configure system dns address-pref

Description

This command configures the DNS address resolving order preference. By default, DNS names are queried for A-records only (address-preference is IPv4-only).

If the address-preference is set to IPv6-first, the DNS server will be queried for AAAA-records (IPv6) first and if a successful replied is not received, then the DNS server is queried for A-records.

Default

address-pref ipv4-only

Platforms

7705 SAR Gen 2

4.42 address-range

address-range

Syntax

no address-range *start-ip-address end-ip-address* [**failover** {**local** | **remote** | **access-driven**}]

no address-range *start-ip-address end-ip-address*

Context

[Tree] (config>router>dhcp>server>pool>subnet address-range)

Full Context

configure router dhcp local-dhcp-server pool subnet address-range

Description

This command configures a range of IP addresses to be served from the pool. All IP addresses between the start and end IP addresses are included (other than specific excluded addresses).

The **no** form of this command removes the address-range parameters from the configuration.

Parameters

start-ip-address

Specifies the start address of this range to include. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

Values a.b.c.d

end-ip-address

Specifies the end address of this range to include. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

Values a.b.c.d

local

Specifies that the local DHCP server has the ownership of this dress range in a redundant setup under normal operation.

remote

Specifies that the remote DHCP server has the ownership of this address range in a redundant setup under normal operation.

access-driven

Specifies that the DHCP server failover system is in control by the access protection mechanisms (SRRP or MC-LAG).

Platforms

7705 SAR Gen 2

address-range

Syntax

address-range *start-ip-address end-ip-address* [**create**]

no address-range *start-ip-address end-ip-address*

Context

[Tree] (config>service>vprn>nat>outside>pool address-range)

[Tree] (config>router>nat>outside>pool address-range)

Full Context

configure service vprn nat outside pool address-range

configure router nat outside pool address-range

Description

This command configures a NAT address range.

Parameters

start-ip-address

Specifies the beginning IP address in a.b.c.d form.

end-ip-address

Specifies the ending IP address in a.b.c.d. form.

create

This parameter must be specified to create the address range instance

Platforms

7705 SAR Gen 2

4.43 address-source

address-source

Syntax

address-source router *router-instance* **dhcp-server** *local-dhcp4-svr-name* **pool** *dhcp4-server-pool* [**secondary-pool** *secondary-pool-name*]

address-source service-name *service-name* **dhcp-server** *local-dhcp4-svr-name* **pool** *dhcp4-server-pool* [**secondary-pool** *secondary-pool-name*]

address-source router *router-instance* **dhcp-server** *local-dhcp6-svr-name* **pool** *dhcp6-server-pool*

address-source service-name *service-name* **dhcp-server** *local-dhcp6-svr-name* **pool** *dhcp6-server-pool*

no address-source

Context

[Tree] (config>service>ies>if>sap>ipsec-gw>lcl-addr-assign>ipv4 address-source)

[Tree] (config>service>vprn>if>sap>ipsec-gw>lcl-addr-assign>ipv4 address-source)

[Tree] (config>service>vprn>if>sap>ipsec-gw>lcl-addr-assign>ipv6 address-source)

[Tree] (config>service>ies>if>sap>ipsec-gw>lcl-addr-assign>ipv6 address-source)

Full Context

configure service ies interface sap ipsec-gw local-address-assignment ipv4 address-source

configure service vprn interface sap ipsec-gw local-address-assignment ipv4 address-source

configure service vprn interface sap ipsec-gw local-address-assignment ipv6 address-source

configure service ies interface sap ipsec-gw local-address-assignment ipv6 address-source

Description

This command specifies the IPv4 or IPv6 source of the local address assignment for the IPsec gateway, which is a pool of a local DHCPv4 or DHCPv6 server. The system will assign an internal address to an IKEv2 remote-access client from the specified pool.

Beside the IP address, netmask and DNS server can also be returned. For IPv4, the netmask and DNS server address can be returned from the specified pool, as well as the IP address. The netmask returned to the IPsec client is derived from the subnet length from the **subnet** *x.x.x.x/m* **create** configuration, not the **subnet-mask** configuration in the subnet context. For IPv6, the DNS server address can be returned from the specified pool, as well as the IP address.

For IPv4, a secondary pool can be optionally specified. The secondary pool is used if the system is unable to assign addresses from the primary pool.

Default

no address-source

Parameters

router-instance

Specifies the router instance ID where the local DHCPv4 or DHCPv6 server is defined, up to 32 characters.

This variant of this command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **address-source service-name** *service-name* variant can be used in all configuration modes.

service-name

Specifies the name of the service where the local DHCPv4 or DHCPv6 server is defined, up to 64 characters.

local-dhcp4-svr-name

Specifies the name of the local DHCPv4 server, up to 32 characters.

local-dhcp6-svr-name

Specifies the name of the local DHCv6 server, up to 32 characters.

dhcp4-server-pool

The name of the pool defined in the specified DHCPv4 server, up to 32 characters.

dhcp6-server-pool

The name of the pool defined in the specified DHCPv6 server, up to 32 characters.

secondary-pool-name

The name of the secondary pool defined in the specified server, up to 32 characters.

Platforms

7705 SAR Gen 2

4.44 adj-sid-hold

adj-sid-hold

Syntax

adj-sid-hold *seconds*

no adj-sid-hold

Context

[Tree] (config>router>isis>segm-rtng adj-sid-hold)

Full Context

configure router isis segment-routing adj-sid-hold

Description

This command configures a timer to hold the ILM or LTN of an adjacency SID following a failure of the adjacency.

When an adjacency to a neighbor fails, the following procedures are followed for both the LFA protected SID and the LFA unprotected SID of this adjacency in SR-MPLS. An adjacency can have both types of SIDs assigned by configuration. An LFA protected adjacency SID is eligible for LFA protection, but the following procedures apply even if a LFA backup was not programmed at the time of the failure. An LFA unprotected adjacency SID is not eligible for LFA protection.

- IGP withdraws the advertisement of the link TLV as well as its adjacency SID sub-TLV.
- The adjacency SID hold timer starts.
- The LTN and ILM records of the adjacency are kept in the datapath for as long as the adjacency SID hold time is running. This allows packets to flow over the LFA backup path, when the adjacency is protected, and allows the ingress LER or PCE time to compute a new path of the SR-TE LSP after IGP converges.
- If the adjacency is restored while the adjacency SID hold timer is running, the timer is aborted, and the adjacency SID remains programmed in the datapath with the retained SID values. However, the backup NHLFE may change if a new LFA SPF runs while the adjacency SID hold timer running. An update to the backup NHLFE is performed immediately following the LFA SPF. In all cases, the adjacency keeps its assigned SID label value.
- If the adjacency SID hold timer expires before the adjacency is restored, the SID is deprogrammed from the datapath and the label returned into the common pool where it was drawn from. Users of the adjacency (for example, SR policy and SR-TE LSP) are also informed.

  When the adjacency is subsequently restored, it gets assigned its allocated static-label value or a new dynamic-label value.
- A new PG-ID is assigned each time an adjacency comes back up. This PG-ID is used by the ILM and LTN of the adjacency SID and of all downstream node SIDs that resolve to a next hop over this adjacency.

The **no** form of this command reverts to the default value.

Default

adj-sid-hold 15

Parameters

seconds

Specifies the adjacency SID hold time, in seconds.

Values 1 to 1800

Platforms

7705 SAR Gen 2

adj-sid-hold

Syntax

adj-sid-hold *seconds*

no adj-sid-hold

Context

[Tree] (config>router>ospf>segm-rtng adj-sid-hold)

Full Context

configure router ospf segment-routing adj-sid-hold

Description

This command configures a timer to hold the ILM or LTN of an adjacency SID following a failure of the adjacency.

When an adjacency to a neighbor fails, the following procedures are followed for both the LFA protected SID and the LFA unprotected SID of this adjacency in SR-MPLS. An adjacency can have both types of SIDs assigned by configuration. An LFA protected adjacency SID is eligible for LFA protection, but the following procedures apply even if a LFA backup was not programmed at the time of the failure. An LFA unprotected adjacency SID is not eligible for LFA protection.

- IGP withdraws the advertisement of the link TLV as well as its adjacency SID sub-TLV.
- The adjacency SID hold timer starts.
- The LTN and ILM records of the adjacency are kept in the datapath for as long as the adjacency SID hold time is running. This allows packets to flow over the LFA backup path, when the adjacency is protected, and allows the ingress LER or PCE time to compute a new path of the SR-TE LSP after IGP converges.
- If the adjacency is restored while the adjacency SID hold timer is running, the timer is aborted, and the adjacency SID remains programmed in the datapath with the retained SID values. However, the backup NHLFE may change when a new LFA SPF is run while the adjacency SID hold timer running. An update to the backup NHLFE is performed immediately following the LFA SPF. In all cases, the adjacency keeps its assigned SID label value.

- If the adjacency SID hold timer expires before the adjacency is restored, the SID is deprogrammed from the datapath and the label returned into the common pool where it was drawn from. Users of the adjacency (for example, SR policy and SR-TE LSP) are also informed.

  When the adjacency is subsequently restored, it gets assigned its allocated static label value or a new dynamic label value.
- A new PG-ID is assigned each time an adjacency comes back up. This PG-ID is used by the ILM and LTN of the adjacency SID and of all downstream node SIDs that resolve to a next hop over this adjacency.

The **no** form of this command reverts to the default value.

Default

adj-sid-hold 15

Parameters

seconds

Specifies the adjacency SID hold time, in seconds.

Values 1 to 1800

Platforms

7705 SAR Gen 2

4.45 adjacency

adjacency

Syntax

[**no**] **adjacency**

Context

[Tree] (debug>router>pim adjacency)

Full Context

debug router pim adjacency

Description

This command enables debugging for PIM adjacencies.

The **no** form of this command disables debugging for PIM adjacencies.

Platforms

7705 SAR Gen 2

adjacency

Syntax

[**no**] **adjacency** [*ip-int-name* | *ip-address* | *nbr-system-id*]

Context

[Tree] (debug>router>isis adjacency)

Full Context

debug router isis adjacency

Description

This command enables debugging for IS-IS adjacency.

The **no** form of the command disables debugging.

Parameters

ip-address

When specified, only adjacencies with the specified interface address are debugged.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

ip-int-name

When specified, only adjacencies with the specified interface name are debugged.

nbr-system-id

When specified, only the adjacency with the specified ID is debugged.

Platforms

7705 SAR Gen 2

4.46 adjacency-set

adjacency-set

Syntax

[**no**] **adjacency-set** *id*

Context

[Tree] (config>router>ospf>segm-rtng adjacency-set)

[Tree] (config>router>isis>segm-rtng adjacency-set)

Full Context

configure router ospf segment-routing adjacency-set

configure router isis segment-routing adjacency-set

Description

This command creates an adjacency set. An adjacency set consists of one or more adjacency SIDs originating on this node. The constituent adjacencies may terminate on different nodes.

The **no** form of this command removes the specified adjacency set.

Parameters

id

Specifies an unsigned integer representing the identifier of the adjacency set.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

adjacency-set

Syntax

[**no**] **adjacency-set** *id*

Context

[Tree] (config>router>ospf>area>interface adjacency-set)

[Tree] (config>router>isis>interface adjacency-set)

Full Context

configure router ospf area interface adjacency-set

configure router isis interface adjacency-set

Description

This command associates an interface with an adjacency set. The adjacency set must have been defined under the IS-IS or OSPF segment-routing context.

The **no** form of this command removes the association.

Parameters

id

Specifies an unsigned integer representing the identifier of the adjacency set.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

4.47 adjacency-sid

adjacency-sid

Syntax

adjacency-sid label *value*

no adjacency-sid

Context

[Tree] (config>router>ospf>area>interface adjacency-sid)

Full Context

configure router ospf area interface adjacency-sid

Description

This command allows a static value to be assigned to an adjacency SID in OSPF segment routing.

The **label** option specifies that the value is assigned to an MPLS label.

The **no** form of this command removes the adjacency SID.

Parameters

label value

Specifies the value of adjacency SID label.

Values 18432 to 52428 | 1048575 (FP4 or FP5 only)

Platforms
7705 SAR Gen 2

adjacency-sid

Syntax
adjacency-sid

Context
[Tree] (config>router>isis>segm-rtng adjacency-sid)
[Tree] (config>router>ospf>segm-rtng adjacency-sid)

Full Context
configure router isis segment-routing adjacency-sid
configure router ospf segment-routing adjacency-sid

Description
Commands in this context configure two SR-MPLS adjacency SIDs per interface.

Platforms
7705 SAR Gen 2

4.48 admin

admin

Syntax
admin

Context
[Tree] (admin)

Full Context
admin

Description
Commands in this context configure administrative system parameters. Only authorized users can execute the commands in the **admin** context.

Platforms

7705 SAR Gen 2

4.49 admin-group

admin-group

Syntax

[**no**] **admin-group** *group-name* [*group-name*]

no admin-group

Context

[Tree] (config>router>if>if-attribute admin-group)

[Tree] (config>service>ies>if>if-attribute admin-group)

[Tree] (config>router>mpls>interface admin-group)

[Tree] (config>service>vprn>if>if-attribute admin-group)

Full Context

configure router interface if-attribute admin-group

configure service ies interface if-attribute admin-group

configure router mpls interface admin-group

configure service vprn interface if-attribute admin-group

Description

This command configures the admin group membership of an interface. The user can apply admin groups to an IES, VPRN, network IP, or MPLS interface.

Each single operation of the **admin-group** command allows a maximum of five (5) groups to be specified at a time. However, a maximum of 32 groups can be added to a given interface through multiple operations. Once an admin group is bound to one or more interface, its value cannot be changed until all bindings are removed.

The configured **admin-group** membership will be applied in all levels or areas the interface is participating in. The same interface cannot have different memberships in different levels or areas.

Only the admin groups bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The **no** form of this command deletes one or more of the admin-group memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.

Default

no admin-group

Parameters

group-name

Specifies up to five groups, each up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain. Each single operation of the **admin-group** command allows a maximum of 5 groups to be specified. However, a maximum of 32 groups can be added to a given interface through multiple operations.

Platforms

7705 SAR Gen 2

admin-group

Syntax

admin-group *group-name* **value** *group-value*

no admin-group *group-name*

Context

[Tree] (config>router>if-attribute admin-group)

Full Context

configure router if-attribute admin-group

Description

This command defines an Administrative Group (AG) that can be associated with an IP or MPLS interface.

AGs, also known as affinity, are used to tag IP and MPLS interfaces that share a specific characteristic with the same identifier. For example, an AG identifier can represent:

- all links that connect to core routers
- all links that have a bandwidth higher than 10 Gb
- all links that are dedicated to a specific service

First configure locally on each router the name and identifier of each AG. A maximum of 32 AGs can be configured per system.

After configuring the router name and identifier, configure the AG membership of an interface. You can apply AGs to a IES, VPRN, network IP, or MPLS interface.

When applied to MPLS interfaces, the interfaces can be included or excluded in the LSP path definition by inferring the AG name. CSPF computes a path that satisfies the AG include and exclude constraints.

When applied to IES, VPRN, or network IP interfaces, the interfaces can be included or excluded in the route next-hop selection by inferring the AG name in a route next-hop policy template applied to an interface or a set of prefixes.

The following provisioning rules apply to the AG configuration. The system rejects the creation of an AG:

- if the name of the AG is the same as that of an existing group, even if the new AG group value is different from the existing group value
- if the AG reuses the same group value but with a different name from an existing group

Only the AGs bound to an MPLS interface are advertised area wide in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

Parameters

group-name

Specifies the name of the group, up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain

group-value

Specifies the integer value associated with the group. The association of group name and value should be unique within an IP or MPLS domain.

Values 0 to 31 – Specifies the value range to use with link LFA next-hop policies or is used as a link color (AG or EAG) with Segment Routing Flex-Algorithms.

32 to 255 – Specifies the value range to use when the EAG is used as a link color with Segment Routing Flex-Algorithms. This higher range fails if used for other applications, such as LFA next-hop policies.

Platforms

7705 SAR Gen 2

admin-group

Syntax

admin-group *admin-group*

no admin-group *admin-group*

Context

[Tree] (config>router>fad>flex-algo>exclude admin-group)

Full Context

configure router flexible-algorithm-definitions flex-algo exclude admin-group

Description

This command configures an administrative group link that will be excluded from the topology graph of the flexible algorithm. If multiple administrative groups are configured, they are all excluded from the topology graph.

Administrative groups are attributes associated with a link. Frequently these administrative groups are described as link colors.

The **no** form of this command removes the admin-group from being excluded from the topology graph.

Default

no admin-group

Parameters

admin-group

Configures an administrative group link to exclude from the topology graph of the configured FAD.

Platforms

7705 SAR Gen 2

admin-group

Syntax

admin-group *admin-group*

no admin-group *admin-group*

Context

[Tree] (config>router>fad>flex-algo>include-all admin-group)

Full Context

configure router flexible-algorithm-definitions flex-algo include-all admin-group

Description

This command configures an administrative group link that will be included in the topology graph of the defined FAD. If multiple administrative groups are configured, groups must be present in a link before the link is included in the flexible algorithm topology graph.

The **no** form of this command removes the specified *admin-group* from being included in the topology graph.

Default

no admin-group

Parameters

admin-group

Configures an administrative group to include in topology graph of the configured FAD.

Platforms

7705 SAR Gen 2

admin-group

Syntax

admin-group *admin-group*

no admin-group *admin-group*

Context

[Tree] (config>router>fad>flex-algo>include-any admin-group)

Full Context

configure router flexible-algorithm-definitions flex-algo include-any admin-group

Description

This command configures an administrative group link that will be included in the topology graph of the configured FAD. If multiple administrative groups are configured, at least one of the administrative groups must be present in a link before the link is included into the flexible algorithm topology graph.

The **no** form of this command removes the *admin-group* from being included in the topology graph.

Default

no admin-group

Parameters

admin-group

Configures an administrative group to include in the topology graph of the configured FAD.

Platforms

7705 SAR Gen 2

4.50 admin-group-frr

admin-group-frr

Syntax

[**no**] **admin-group-frr**

Context

[Tree] (config>router>mpls admin-group-frr)

Full Context

configure router mpls admin-group-frr

Description

This command enables the use of the admin-group constraints in the association of a manual or dynamic bypass LSP with the primary LSP path at a Point-of-Local Repair (PLR) node.

When this command is enabled, each PLR node reads the admin-group constraints in the FAST_REROUTE object in the Path message of the LSP primary path. If the FAST_REROUTE object is not included in the Path message, then the PLR will read the admin-group constraints from the Session Attribute object in the Path message.

If the PLR is also the ingress LER for the LSP primary path, then it just uses the admin-group constraint from the LSP and/or path level configurations.

The PLR node then uses the admin-group constraints along with other constraints, such as hop-limit and SRLG, to select a manual or dynamic bypass among those that are already in use.

If none of the manual or dynamic bypass LSP satisfies the admin-group constraints, and/or the other constraints, the PLR node will request CSPF for a path that merges the closest to the protected link or node and that includes or excludes the specified admin-group IDs.

If the user changes the configuration of the above command, it will not have any effect on existing bypass associations. The change will only apply to new attempts to find a valid bypass.

The **no** form of this command disables the use of administrative group constraints on a FRR backup LSP at a PLR node.

Default

no frr-admin-group

Platforms

7705 SAR Gen 2

4.51 admin-password

admin-password

Syntax

admin-password *password* [**hash** | **hash2**]

no admin-password

Context

[Tree] (config>system>security>password admin-password)

Full Context

configure system security password admin-password

Description

This command allows a user (with admin permissions) to configure a password that enables a user to become an administrator.

This password is valid only for one session. When enabled, no authorization to TACACS+ or RADIUS is performed and the user is locally regarded as an admin user.

This functionality can be enabled in two contexts:

config>system>security>password>admin-password

<global> enable-admin

If the admin-password is configured in the **config>system>security>password** context, then any user can enter the special mode by entering the **enable-admin** command.

enable-admin is in the default profile. By default, all users are given access to this command.

After the **enable-admin** command is entered, the user is prompted for a password. If the password matches, user is given unrestricted access to all the commands.

The minimum length of the password is determined by the **minimum-length** command. The complexity requirements for the password are determined by the **complexity** command.

Note:
The *password* argument of this command is not sent to the servers. This is consistent with other commands that configure secrets.

The usernames and passwords in the FTP and TFTP URLs will not be sent to the authorization or accounting servers when the **file>copy** *source-url dest-url* command is executed.

For example:

file copy ftp://test:secret@10.20.31.79/test/srcfile cf1:\destfile

In this example, the username 'test' and password 'secret' will not be sent to the AAA servers (or to any logs). They will be replaced with "*****".

The **no** form of this command removes the admin password from the configuration.

Note:
This command applies to a local user, in addition to users on RADIUS, TACACS, and LDAP.

Default

no admin-password

Parameters

password

Configures the password that enables a user to become a system administrator. The maximum length can be up to 56 characters if unhashed, 60 characters if hashed with bcrypt, from 87 to 92 characters if hashed with sha2-pbkdf2, 32 characters if the hash keyword is specified, or 54 characters if the hash2 keyword is specified. The unhashed cleartext password form should meet all the requirements that are defined by the **complexity** command.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form or hashed with bcrypt or PBKDF2. For security, all keys are stored in the configuration file in hashed form (using bcrypt or PBKDF2, depending on the hashing configuration parameter) or, for backward compatibility, can be stored in encrypted form with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form or hashed with bcrypt or PBKDF2. For security, all keys are stored in the configuration file in hashed form (using bcrypt or PBKDF2, depending

on the hashing configuration parameter) or, for backward compatibility, can be stored in encrypted form with the **hash** or **hash2** parameter specified.

Platforms

7705 SAR Gen 2

4.52 admin-state

admin-state

Syntax

[no] admin-state

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>info-output>always-display admin-state)

Full Context

configure system management-interface cli md-cli environment info-output always-display admin-state

Description

This command configures that the values of the **admin-state** elements in the **info** output (without the **detail** option) are always displayed, even if they are the default values.

The **no** form of this command excludes the values of the **admin-state** elements from the **info** output display.

Default

no admin-state

Platforms

7705 SAR Gen 2

4.53 admin-status

admin-status

Syntax

admin-status {**rx** | **tx** | **tx-rx** | **disabled**}

Context

[Tree] (config>port>ethernet>lldp>dstmac admin-status)

Full Context

configure port ethernet lldp dest-mac admin-status

Description

This command configures LLDP transmission/reception frame handling.

Default

admin-status disabled

Parameters

rx

Specifies the LLDP agent will receive, but will not transmit LLDP frames on this port.

tx

Specifies that the LLDP agent will transmit LLDP frames on this port and will not store any information about the remote systems connected.

tx-rx

Specifies that the LLDP agent transmits and receives LLDP frames on this port.

disabled

Specifies that the LLDP agent does not transmit or receive LLDP frames on this port. If there is remote systems information which is received on this port and stored in other tables, before the port's admin status becomes disabled, then the information will naturally age out.

Platforms

7705 SAR Gen 2

admin-status

Syntax

admin-status {**rx** | **tx** | **tx-rx** | **disabled**}

Context

[Tree] (config>lag>lldp-member-template>dstmac admin-status)

Full Context

configure lag lldp-member-template dest-mac admin-status

Description

This command configures the LLDP transmission and reception frame handling.

Default

admin-status disabled

Parameters

rx

Keyword to specify that the LLDP agent receives, but does not transmit LLDP frames on this port.

tx

Keyword to specify that the LLDP agent transmits LLDP frames on this port and does not store any information about the remote systems connected.

tx-rx

Keyword to specify that the LLDP agent transmits and receives LLDP frames on this port.

disabled

Keyword to specify that the LLDP agent does not transmit or receive LLDP frames on this port. If remote system information is received on this port and stored in other tables before the administrative status of the port becomes disabled, the information naturally ages out.

Platforms

7705 SAR Gen 2

4.54 admin-tag

admin-tag

Syntax

[**no**] **admin-tag** *tag-value*

Context

[Tree] (config>router>mpls>lsp-template admin-tag)

[Tree] (config>router>mpls>lsp admin-tag)

Full Context

configure router mpls lsp-template admin-tag

configure router mpls lsp admin-tag

Description

This assigns an administrative tag to an LSP. The administrative tag can be used to enable routes with certain administrative tags to resolve using LSPs of matching administrative tags.

Up to four tags can be assigned to an LSP.

The administrative tag must exist under **config>router>admin-tags**.

The **no** form of this command removes the administrative tag.

Parameters

tag-value

The value of the admin-tag, up to 32 characters.

Platforms

7705 SAR Gen 2

admin-tag

Syntax

[**no**] **admin-tag** *tag*

Context

[Tree] (config>router>admin-tags admin-tag)

Full Context

configure router admin-tags admin-tag

Description

This command configures an admin tag value in the nodal LSP administrative tag database.

Up to 256 admin tags can be configured.

The **no** form of this command removes the admin tag.

Parameters

tag

The value of the administrative tag, up to 32 characters.

Platforms

7705 SAR Gen 2

4.55 admin-tag-policy

admin-tag-policy

Syntax

admin-tag-policy *policy-name*

no admin-tag-policy

Context

[Tree] (config>router>policy-options>policy-statement>default-action admin-tag-policy)

[Tree] (config>router>policy-options>policy-statement>entry>action admin-tag-policy)

Full Context

configure router policy-options policy-statement default-action admin-tag-policy

configure router policy-options policy-statement entry action admin-tag-policy

Description

This command assigns a route admin tag policy as an action in a route policy.

The admin tag policy must exist under **config>router>admin-tags**.

The **no** form of this command removes the admin tag policy.

Parameters

policy-name

Specifies the name of the admin tag policy, up to 64 characters.

Platforms

7705 SAR Gen 2

4.56 admin-tags

admin-tags

Syntax

admin-tags

Context

[Tree] (config>router admin-tags)

Full Context

configure router admin-tags

Description

Commands in this context configure admin tags and router admin tag policy templates used for route resolution to LSPs.

Platforms

7705 SAR Gen 2

4.57 adspec

adspec

Syntax

[**no**] **adspec**

Context

[Tree] (config>router>mpls>lsp-template adspec)

[Tree] (config>router>mpls>lsp adspec)

Full Context

configure router mpls lsp-template adspec

configure router mpls lsp adspec

Description

When enabled, the ADSPEC object will be included in RSVP messages for this LSP. The ADSPEC object is used by the ingress LER to discover the minimum value of the MTU for links in the path of the LSP. By default, the ingress LER derives the LSP MTU from that of the outgoing interface of the LSP path.

A bypass LSP always signals the ADSPEC object since it protects both primary paths which signal the ADSPEC object and primary paths which do not. This means that MTU of LSP at ingress LER may change to a different value from that derived from the outgoing interface even if the primary path has ADSPEC disabled.

Default

no adspec — No ADSPEC objects are included in RSVP messages.

Platforms

7705 SAR Gen 2

4.58 adv-adj-addr-only

adv-adj-addr-only

Syntax

[**no**] **adv-adj-addr-only**

Context

[Tree] (config>router>ldp>session-params>peer adv-adj-addr-only)

Full Context

configure router ldp session-parameters peer adv-adj-addr-only

Description

This command provides a means for an LDP router to advertise only the local IPv4 or IPv6 interfaces it uses to establish hello adjacencies with an LDP peer. By default, when a router establishes an LDP session with a peer, it advertises in an LDP Address message the addresses of all local interfaces to allow the peer to resolve LDP FECs distributed by this router. Similarly, a router sends a Withdraw Address message to of all its peers to withdraw a local address if the corresponding interface went down or was deleted.

This new option reduces CPU processing when a large number of LDP neighbors come up or go down. The new CLI option is strongly recommended in mobile backhaul networks where the number of LDP peers can be very large.

The **no** form of this command reverts LDP to the default behavior of advertising all local interfaces.

Platforms

7705 SAR Gen 2

4.59 adv-local-lsr-id

adv-local-lsr-id

Syntax

[**no**] **adv-local-lsr-id**

Context

[Tree] (config>router>ldp>targeted-session>peer-template adv-local-lsr-id)

[Tree] (config>router>ldp>session-params>peer adv-local-lsr-id)

Full Context

configure router ldp targeted-session peer-template adv-local-lsr-id

configure router ldp session-parameters peer adv-local-lsr-id

Description

This command advertises a local LSR ID over a specified LDP session.

Advertisement of a local LSR ID over a given LDP session is configured using the **adv-local-lsr-id** command in the peer session-parameters. If a user disables the **adv-local-lsr-id** command, then the system will withdraw the FEC for the local LSR ID.

The SR OS router uses the following rules when advertising a local LSR ID:

- If the session parameters have the default configuration and the targeted peer template has the default configuration, the local LSR ID is not advertised.

- If the session parameters have the default configuration but the targeted peer template has an explicit configuration for advertisement of the local LSR ID, the targeted peer template configuration is used.
- If the session parameters have an explicit configuration for advertisement of the local LSR ID but the targeted peer template has the default configuration, the session parameter configuration is used.
- If both the session parameters and the targeted peer template have an explicit configuration for advertisement of the local LSR ID, then the session parameter configuration is used.

The **no** form of this command withdraws the FEC for the local LSR ID.

Default

no adv-local-lsr-id

Platforms

7705 SAR Gen 2

4.60 adv-mtu-override

adv-mtu-override

Syntax

[**no**] **adv-mtu-override**

Context

[Tree] (config>service>sdp adv-mtu-override)

Full Context

configure service sdp adv-mtu-override

Description

This command overrides the advertised VC-type MTU of all spoke-sdps of L2 services using this SDP-ID. When enabled, the router signals a VC MTU equal to the service MTU, which includes the Layer 2 header. It also allows this router to accept an MTU advertised by the far-end PE which value matches either its advertised MTU or its advertised MTU minus the L2 headers.

By default, the router advertises a VC-MTU equal to the L2 service MTU minus the Layer 2 header and always matches its advertised MTU to that signaled by the far-end PE router, otherwise the spoke-sdp goes operationally down.

When this command is enabled on the SDP, it has no effect on a spoke-sdp of an IES/VPRN spoke interface using this SDP-ID. The router continues to signal a VC MTU equal to the net IP interface MTU, which is min{ip-mtu, sdp operational path mtu - L2 headers}. The router also continues to make sure that the advertised MTU values of both PE routers match or the spoke-sdp goes operationally down.

The **no** form of the command disables the VC-type MTU override and returns to the default behavior.

Default

no adv-mtu-override

Platforms

7705 SAR Gen 2

4.61 adv-noaddrs-global

adv-noaddrs-global

Syntax

adv-noaddrs-global [**esm-proxy**] [**esm-relay**] [**relay**] [**server**]

no adv-noaddrs-global

Context

[Tree] (config>system>dhcp6 adv-noaddrs-global)

Full Context

configure system dhcp6 adv-noaddrs-global

Description

This command configures the different DHCPv6 applications to send the NoAddrsAvail Status-Code in DHCPv6 Advertise messages at the global DHCP message level.

By default, all applications send the NoAddrsAvail Status-Code in DHCPv6 Advertise messages at the IA_NA Option level.

Different applications for which NoAddrsAvail Status-Code in DHCPv6 Advertise messages can be configured at the global DHCP message level.

The only valid combination in current SR OS is **adv-noaddrs-global esm-relay server**.

The **no** form of this command reverts to the default.

Default

no adv-noaddrs-global. All applications send the NoAddrsAvail Status-Code in DHCPv6 Advertise messages at the IA_NA Option level.

Parameters

esm-proxy

Specifies the DHCPv6 proxy server on subscriber group-interfaces. Not supported in current SR OS.

esm-relay

Specifies the DHCPv6 relay on subscriber group-interfaces. Must be enabled together with the DHCPv6 server (server) application.

relay

Specifies the DHCPv6 relay on regular IES or VPRN interfaces. Not supported in current SR OS.

server

Specifies the DHCPv6 server. Must be enabled together with the DHCPv6 relay on subscriber interfaces (esm-relay) application.

Platforms

7705 SAR Gen 2

4.62 adv-service-mtu

adv-service-mtu

Syntax

adv-service-mtu *octets*

no adv-service-mtu

Context

[Tree] (config>service>epipe>spoke-sdp adv-service-mtu)

[Tree] (config>service>vpls>spoke-sdp adv-service-mtu)

Full Context

configure service epipe spoke-sdp adv-service-mtu

configure service vpls spoke-sdp adv-service-mtu

Description

This command configures the MTU value signaled in the targeted LDP for the spoke-SDP and is used to validate the value signaled by the far-end PE. If configured, this value is used instead of the service MTU. However, the configuration does not affect the locally enforced value, which is still based on the service MTU. This command cannot be configured on a spoke-SDP that is bound to an SDP with the **adv-mtu-override** command.

When configured, an adjusted service MTU is used. See the **service-mtu** command for more information.

The **no** form of this command removes the configuration.

Default

no adv-service-mtu

Parameters

octets

The size of the MTU in octets, expressed as a decimal integer.

Values 0 to 9782

Platforms

7705 SAR Gen 2

adv-service-mtu

Syntax

adv-service-mtu *number*

no adv-service-mtu

Context

[Tree] (config>service>epipe>bgp adv-service-mtu)

[Tree] (config>service>vpls>bgp adv-service-mtu)

Full Context

configure service epipe bgp adv-service-mtu

configure service vpls bgp adv-service-mtu

Description

This command configures the Layer 2 MTU value (advertised for BGP signaling) or the MTU interface parameter (advertised for LDP signaling) for the service. The configured MTU information is used to validate the value signaled by the far-end PE. However, this configuration does not affect the locally enforced value, which is still based on the service MTU.

The **no** form of this command reverts to the default Layer 2 MTU value for BGP signaling or to the default MTU interface parameter for LDP signaling for the service, which uses an adjusted **service-mtu** value. See the **service-mtu** command for more information.

Default

no adv-service-mtu

Parameters

number

Specifies the size, in octets, of the Layer 2 MTU value to advertise for BGP signaling for the service.

Values 0 to 9782

Platforms

7705 SAR Gen 2

4.63 advertise

advertise

Syntax

advertise {**static** | **dynamic**} [**route-tag** [*1..255*]]

no advertise {**static** | **dynamic**}

Context

[Tree] (config>service>ies>if>vpls>evpn>arp advertise)

[Tree] (config>service>ies>if>vpls>evpn>nd advertise)

Full Context

configure service ies interface vpls evpn arp advertise

configure service ies interface vpls evpn nd advertise

Description

This command enables the advertisement of static and dynamic ARP and ND entries that are installed in the ARP and ND cache into EVPN MAC/IP routes. This command must be used along with **no learn-dynamic**.

Default

no advertise

Parameters

static

Enables ARP/ND host routes to be created in the route table from EVPN ARP/ND entries

dynamic

Enables ARP/ND host routes to be created in the route table out of dynamic ARP/ND entries (learned from ARP/ND messages received from the hosts).

route-tag

Specifies the route tag that is added in the route table for ARP/ND host routes of type **dynamic**, or **static**. This tag can be matched on BGP VRF export and BGP peer export policies.

Values 1 to 255

Platforms

7705 SAR Gen 2

advertise

Syntax

advertise {**static** | **dynamic**} [**route-tag** [*1..255*]] **interface-less-routing** [**bgp-evpn-instance** [*1..1*]]

advertise {**static** | **dynamic**} [**route-tag** [*1..255*]]

no advertise {**static** | **dynamic**}

Context

[Tree] (config>service>vprn>if>vpls>evpn>arp advertise)

[Tree] (config>service>vprn>if>vpls>evpn>nd advertise)

Full Context

configure service vprn interface vpls evpn arp advertise

configure service vprn interface vpls evpn nd advertise

Description

This command enables the advertisement of static and dynamic ARP and ND entries that are installed in the ARP and ND cache into EVPN MAC/IP routes. This command must be used along with the **no learn-dynamic** command.

Default

no advertise

Parameters

static

Enables ARP or ND host routes to be created in the route table from EVPN ARP or ND entries

dynamic

Enables ARP or ND host routes to be created in the route table out of dynamic ARP or ND entries (learned from ARP or ND messages received from the hosts).

route-tag

Keyword to specify the route tag is added in the route table for ARP or ND host routes of type **dynamic**, or **static**. This tag can be matched on BGP VRF export and BGP peer export policies.

Values 1 to 255

interface-less-routing

Keyword to specify that the advertisement in EVPN MAC/IP advertisement routes include the label1 and route target of the R-VPLS EVPN service and the label2 value and route target of the EVPN interface-less instance in the linked VPRN.

bgp-evpn-instance

Keyword to specify the EVPN interface-less BGP instance from which the label and route target are taken when advertising the ARP or ND entry in an EVPN MAC/IP advertisement route.

Values 1 to 1

Platforms

7705 SAR Gen 2

advertise

Syntax

advertise *fad-name*

no advertise

Context

[Tree] (config>router>isis>flex-algos>flex-algo advertise)

[Tree] (config>router>ospf>flex-algos>flex-algo advertise)

Full Context

configure router isis flexible-algorithms flex-algo advertise

configure router ospf flexible-algorithms flex-algo advertise

Description

This command enables the advertisement of a locally configured Flexible Algorithm Definition (FAD).

A locally defined FAD is only advertised if it is administratively enabled. A router can advertise only a single locally defined FAD by using the *fad-name* as reference anchor.

The winning FAD used by a router must be consistent with the winning FAD on all other routers. This avoids routing loops and traffic blackholing. The winning FAD is selected using a tie-breaker algorithm that first selects the highest advertised FAD priority and next the highest system Id.

The **no** form of this command removes the advertisement of a flexible algorithm definition.

Default

no advertise

Parameters

fad-name

Configures the FAD name, up to 32 characters. By default, no locally configured FAD is advertised.

Platforms

7705 SAR Gen 2

advertise

Syntax

[**no**] **advertise**

advertise weight dynamic [**max-dynamic-weight** *max-dynamic-weight*]

advertise weight *weight*

Context

[Tree] (configure>service>vpls>bgp-evpn>ip-route-link-bw advertise)

[Tree] (configure>service>vprn>bgp-evpn>mpls>evpn-link-bw advertise)

Full Context

configure service vpls bgp-evpn ip-route-link-bandwidth advertise

configure service vprn bgp-evpn mpls evpn-link-bandwidth advertise

Description

This command enables the advertisement of the EVPN link bandwidth extended community along with the IP Prefix routes.

The **no** form of this command disables the advertisement of the EVPN link bandwidth extended community.

Default

no advertise

Parameters

weight

Specifies the weight advertised in the EVPN link bandwidth extended community for the advertised EVPN IP prefix routes for the service.

Values 1 to 128

weight dynamic

Keyword to specify that the weight is dynamically set based on the number of BGP PE-CE paths for the IP-Prefix that is advertised in an EVPN IP-Prefix route.

max-dynamic-weight

Specifies the maximum weight advertised in the EVPN link bandwidth extended community for the advertised EVPN IP-Prefix routes for the service. If **weight dynamic** is configured, the actual advertised weight is the minimum of the number of BGP PE-CE paths for the prefix and the configured maximum weight.

Values 1 to 128

Platforms

7705 SAR Gen 2

advertise

Syntax

advertise [**holdtime** *seconds*]

no advertise

Context

[Tree] (configure>router>bgp>group>bfd-strict-mode advertise)

[Tree] (configure>router>bgp>group>neighbor>bfd-strict-mode advertise)

[Tree] (configure>service>vprn>bgp>group>bfd-strict-mode advertise)

[Tree] (configure>router>bgp>bfd-strict-mode advertise)

[Tree] (configure>service>vprn>bgp>bfd-strict-mode advertise)

[Tree] (configure>service>vprn>bgp>group>neighbor>bfd-strict-mode advertise)

Full Context

configure router bgp group bfd-strict-mode advertise

configure router bgp group neighbor bfd-strict-mode advertise

configure service vprn bgp group bfd-strict-mode advertise

configure router bgp bfd-strict-mode advertise

configure service vprn bgp bfd-strict-mode advertise

configure service vprn bgp group neighbor bfd-strict-mode advertise

Description

This command configures BGP to advertise the Strict-BFD capability to peers that are within scope of this command and meet the following requirements:

- The **bfd-enable** command that applies to the peer is enabled (through either configuration or inheritance).
- The interface associated with the peer has a valid BFD configuration.

When the preceding conditions are satisfied and two peers attempting to form a session both advertise the Strict-BFD capability, the BGP finite state machine in each router transitions the session state to established after the BFD session with the peer enters the up state.

The **no** form of this command prevents BGP from advertising the Strict-BFD capability to peers.

Default

no advertise

Parameters

seconds

Specifies the maximum time (in seconds) BGP waits for the BFD session to come up, provided that the Strict-BFD procedures apply to a session, and the negotiated BGP hold

time is zero (no keepalives). If the negotiated BGP hold time is greater than zero, the **holdtime** parameter is not considered.

Values 1 to 65535

Default 30

Platforms

7705 SAR Gen 2

4.64 advertise-admin-group

advertise-admin-group

Syntax

advertise-admin-group {prefer-ag | eag-only | ag-eag}

no advertise-admin-group

Context

[Tree] (config>router>ospf>flex-algos advertise-admin-group)

[Tree] (config>router>isis>flex-algos advertise-admin-group)

Full Context

configure router ospf flexible-algorithms advertise-admin-group

configure router isis flexible-algorithms advertise-admin-group

Description

This command configures the type of Aministrative Group (AG) or Extended Administrative Group (EAG) TLVs the router advertises as the Interior Gateway Protocol (IGP) link attribute. This command is configured for this IGP instance.

The **no** form of this command removes the configuration.

Default

prefer-ag

Parameters

prefer-ag

Keyword to specify that the router advertises the Administrative Group (AG) TLV as the IGP link attribute if the affinity bits in the **configure router if-attribute admin-group** *value* command are configured between 0 to 31. If no EAG (32 to 255) affinity bits are configured, only the AG TLV is advertised as the IGP link attribute.

If the affinity bits are configured in both the AG (0 to 31) and EAG (32 to 255) range, the router advertises both the AG and the EAG TLVs as the IGP link attributes.

eag-only

Keyword to specify that the router advertises only the EAG TLV as the IGP link attribute. No AG TLV is advertised if this keyword is configured.

ag-eag

Keyword to specify that the router can advertise both the AG and the EAG TLVs as the IGP link attributes, even without the affinity bit in the EAG range configured in the **configure router if-attribute admin-group** *value* command. If no affinity bit is configured in the AG range (0 to 31), the router prunes the AG TLV. Configuring this keyword allows for backward compatibility for vendor implementations that support only AG, while still supporting EAG.

Platforms

7705 SAR Gen 2

4.65 advertise-bgp

advertise-bgp

Syntax

advertise-bgp route-distinguisher *rd* [**community** *community*]

no advertise-bgp route-distinguisher *rd*

Context

[Tree] (config>service>pw-routing>local-prefix advertise-bgp)

Full Context

configure service pw-routing local-prefix advertise-bgp

Description

This command enables a given prefix to be advertised in MP-BGP for dynamic MS-PW routing.

The **no** form of this command will explicitly withdraw a route if it has been previously advertised.

Default

no advertise-bgp

Parameters

rd

Specifies an 8-octet route distinguisher associated with the prefix. Up to 4 unique route distinguishers can be configured and advertised for a given prefix though multiple instances of the advertise-bgp command. This parameter is mandatory.

Values (6 bytes, other 2 Bytes of type will be automatically generated) asn:number1 (RD Type 0): 2bytes ASN and 4 bytes locally administered number ip-address:number2 (RD Type 1): 4bytes IPv4 and 2 bytes locally administered number;

community

An optional BGP communities attribute associated with the advertisement. To delete a previously advertised community, advertise-bgp route-distinguisher must be run again with the same value for the RD but excluding the community attribute.

Values

community	{2-byte-as-number:comm-va1}
2-byte-asnumber	0 to 65535
comm.-val	0 to 65535

Platforms

7705 SAR Gen 2

4.66 advertise-delay

advertise-delay

Syntax

[**no**] **advertise-delay**

Context

[Tree] (config>router>ospf>te-opts advertise-delay)

Full Context

configure router ospf traffic-engineering-options advertise-delay

Description

This command configures the advertisement of link delay in the IGP LSDB within the OSPF-TE TLV attribute or when the Application Specific Link Attribute (ASLA) is enabled within the SR-TE ASLA.

When the router is configured with the **configure router ospf traffic-engineering-options sr-te application-specific-link-attributes** command to generate SR-TE ASLA attributes, link delay is advertised as a legacy RFC 3630 TE TLV when RSVP-TE is enabled and as an ASLA RFC 8920 TLV for SR-TE when MPLS is enabled for an interface.

SR OS accepts and handles both legacy RSVP-TE TLVs and ASLAs for the RSVP application. However, SR OS only advertises RFC 3630 legacy RSVP-TE TLVs (as recommended by RFC 8920) to avoid compatibility issues.

The **no** form of this command disables link delay advertisement.

Default

no advertise-delay

Platforms

7705 SAR Gen 2

advertise-delay

Syntax

[**no**] **advertise-delay**

Context

[Tree] (config>router>isis>te advertise-delay)

Full Context

configure router isis traffic-engineering-options advertise-delay

Description

This command enables the advertisement of link delay in the IGP LSDB within legacy Traffic Engineering (TE) attributes in IS-IS or within the Application Specific Link Attribute (ASLA) when ASLA is enabled for SR-TE or RSVP-TE applications.

When **application-link-attributes legacy** command is configured for SR-TE or RSVP-TE, link delay is advertised as a legacy TE TLV with the ASLA legacy bit set.

The **no** form of this command disables link delay advertisement.

Default

no advertise-delay

Platforms

7705 SAR Gen 2

4.67 advertise-external

advertise-external

Syntax

[**no**] **advertise-external** [**ipv4**] [**ipv6**] [**label-ipv4**] [**label-ipv6**]

Context

[Tree] (config>router>bgp advertise-external)

Full Context

configure router bgp advertise-external

Description

This command allows BGP to advertise its best external route to a destination even when its best overall route is an internal route. Entering the command (or its **no** form) with no address family parameters is equivalent to specifying all supported address families.

The **no** form of this command disables Advertise Best External for the BGP family.

Default

no advertise-external

Parameters

ipv4

Enables the best-external advertisement for unlabeled unicast IPv4 routes.

ipv6

Enables the best-external advertisement for unlabeled unicast IPv6 routes.

label-ipv4

Enables the best-external advertisement for labeled-unicast IPv4 routes.

label-ipv6

Enables the best-external advertisement for labeled-unicast IPv6 routes.

Platforms

7705 SAR Gen 2

4.68 advertise-inactive

advertise-inactive

Syntax

[**no**] **advertise-inactive**

Context

[Tree] (config>service>vprn>bgp>group advertise-inactive)

[Tree] (config>service>vprn>bgp>group>neighbor advertise-inactive)

[Tree] (config>service>vprn>bgp advertise-inactive)

Full Context

configure service vprn bgp group advertise-inactive

configure service vprn bgp group neighbor advertise-inactive

configure service vprn bgp advertise-inactive

Description

This command enables or disables the advertising of inactive BGP routers to other BGP peers.

By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a given destination.

When the BGP **advertise-inactive** command is configured so that it applies to a BGP session it has the following effect on the IPv4, IPv6, mcast-ipv4, mcast-ipv6, label-IPv4 and label-IPv6 routes advertised to that peer:

- If the active route for the IP prefix is a BGP route then that route is advertised.
- If the active route for the IP prefix is a non-BGP route and there is at least one valid but inactive BGP route for the same destination then the best of the inactive and valid BGP routes is advertised unless the non-BGP active route is matched and accepted by an export policy applied to the session.
- If the active route for the IP prefix is a non-BGP route and there are no (valid) BGP routes for the same destination then no route is advertised for the prefix unless the non-BGP active route is matched and accepted by an export policy applied to the session.

Default

no advertise-inactive

Platforms

7705 SAR Gen 2

advertise-inactive

Syntax

[**no**] **advertise-inactive**

Context

[Tree] (config>router>bgp>group>neighbor advertise-inactive)

[Tree] (config>router>bgp>group advertise-inactive)

[Tree] (config>router>bgp advertise-inactive)

Full Context

configure router bgp group neighbor advertise-inactive

configure router bgp group advertise-inactive

configure router bgp advertise-inactive

Description

This command enables the advertising of inactive BGP routes to other BGP peers. By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the used route within the system for a given destination.

The **no** form of this command disables the advertising of inactive BGP routers to other BGP peers.

Default

no advertise-inactive

Platforms

7705 SAR Gen 2

4.69 advertise-ipv6-next-hops

advertise-ipv6-next-hops

Syntax

advertise-ipv6-next-hops [**ipv4**]

no advertise-ipv6-next-hops

Context

[Tree] (config>service>vprn>bgp>group>neighbor advertise-ipv6-next-hops)

[Tree] (config>service>vprn>bgp>group advertise-ipv6-next-hops)

[Tree] (config>service>vprn>bgp advertise-ipv6-next-hops)

Full Context

configure service vprn bgp group neighbor advertise-ipv6-next-hops

configure service vprn bgp group advertise-ipv6-next-hops

configure service vprn bgp advertise-ipv6-next-hops

Description

When this command is configured, with the IPv4 option, so that it applies to a BGP session established on top of IPv6 transport, IPv4 BGP routes can be advertised with a true IPv6 address when originated or when **next-hop-self** (configured or automatic) is applied.

If an IPv4 route must originate or be advertised with a **next-hop-self** and the corresponding **advertise-ipv6-next-hops** command option does not apply to the session or if an appropriate **extended-nh-encoding** capability was not received from the remote peer, then the route is advertised with the IPv4 system address as the BGP next-hop.

If an IPv4 route is matched by a BGP export policy entry that tries to change the next hop to an IPv6 address and the corresponding **advertise-ipv6-next-hops** command option does not apply to the session or if an appropriate **extended-nh-encoding** capability was not received from the remote peer, then the route is handled as though it was rejected by the policy entry.

This command has no effect on sessions established over IPv4 transport.

The **no** form of this command reverts to the default.

Default

no advertise-ipv6-next-hops

Parameters

ipv4

Allows IPv4 unicast routes to be advertised to IPv6-transport peers with an IPv6 address as the BGP next-hop in cases of route origination or **next-hop-self** (configured or automatic). It also allows export policies to change the BGP next-hop of an IPv4 route to an IPv6 address. All of these cases require the remote peer to advertise the necessary extended NH encoding capability. It may be necessary to configure the **forward-ipv4-packets** command under the appropriate **interface>ipv6** contexts in order to enable datapath support for these control plane exchanges.

Platforms

7705 SAR Gen 2

advertise-ipv6-next-hops

Syntax

advertise-ipv6-next-hops [**vpn-ipv6**] [**label-ipv6**] [**evpn**] [**vpn-ipv4**] [**label-ipv4**] [**ipv4**]

no advertise-ipv6-next-hops

Context

[Tree] (config>router>bgp advertise-ipv6-next-hops)

[Tree] (config>router>bgp>group advertise-ipv6-next-hops)

[Tree] (config>router>bgp>group>neighbor advertise-ipv6-next-hops)

Full Context

configure router bgp advertise-ipv6-next-hops

configure router bgp group advertise-ipv6-next-hops

configure router bgp group neighbor advertise-ipv6-next-hops

Description

This command applies to a BGP session established on top of IPv6 transport; BGP routes belonging to the specified families can be advertised with a true IPv6 address when originated or when **next-hop-self** (configured or automatic) is applied.

This command has no effect on routes advertised to IPv4 peers.

When this command is not enabled, the following considerations apply:

- If a VPN IPv6 or label IPv6 route needs to be originated or advertised with **next-hop-self** to an IPv6 transport peer the route is advertised with the IPv4 system address as BGP next-hop (encoded as an IPv4-mapped IPv6 address).
- If a VPN-IPv4 or label IPv4 route needs to be originated or advertised with **next-hop-self** or if an appropriate **extended-nh-encoding** capability was not received from the remote peer, the route is advertised with the IPv4 system address as the BGP next-hop.
- If a VPN IPv4 or label IPv4 route is matched by a BGP export policy entry that tries to change the next-hop to an IPv6 address and an appropriate **extended-nh-encoding** capability was not received from the remote peer, the route is handled as though it was rejected by the policy entry.

The **no** form of this command disables the setting of next hops to a global IPv6 address for the family.

Default

no advertise-ipv6-next-hops

Parameters

vpn-ipv6

Allows VPN IPv6 routes to be advertised to IPv6 transport peers with an IPv6 address as the BGP next-hop in cases of route origination or **next-hop-self** (configured or automatic).

label-ipv6

Allows label IPv6 routes to be advertised to IPv6 transport peers with an IPv6 address as the BGP next-hop in cases of route origination or **next-hop-self** (configured or automatic).

vpn-ipv4

Allows VPN IPv4 routes to be advertised to IPv6 transport peers with an IPv6 address as the BGP next-hop in cases of route origination or **next-hop-self** (configured or automatic). It also allows export policies to change the BGP next-hop of a VPN IPv4 route to an IPv6 address. All of these cases require the remote peer to advertise the necessary extended NH encoding capability.

label-ipv4

Allows label IPv4 routes to be advertised to IPv6 transport peers with an IPv6 address as the BGP next-hop in cases of route origination or **next-hop-self** (configured or automatic). It also allows export policies to change the BGP next-hop of a label IPv4 route to an IPv6 address. All of these cases require the remote peer to advertise the necessary extended NH encoding capability.

ipv4

Instructs BGP to advertise an extended NH encoding capability for NLRI AFI=1, NLRI SAFI=1 and next-hop AFI=2.

evpn

Allows EVPN routes to be advertised to IPv6 transport peers.

Platforms

7705 SAR Gen 2

4.70 advertise-label

advertise-label

Syntax

advertise-label {**per-prefix** | **pop** | **pop-and-forward**}

no advertise-label

Context

[Tree] (config>router>policy-options>policy-statement>entry>action advertise-label)

[Tree] (config>router>policy-options>policy-statement>default-action advertise-label)

Full Context

configure router policy-options policy-statement entry action advertise-label

configure router policy-options policy-statement default-action advertise-label

Description

This command configures the label allocation method for advertised routes. The effect of the **advertise-label** command depends on the context where the associated policy is applied.

Use the **per-prefix** option and configure the command in the default action or entry-specific action of a VRF export policy to advertise every qualifying matched route with a per-prefix label in the resulting VPN-IP routes. In this situation, non-qualifying routes include local interface routes and BGP-VPN routes. The command overrides, for specific routes, the configured label-mode of the exporting VPRN service.

Use the per-prefix option and configure the command in the default action or entry-specific action of a BGP import policy to assign a per-prefix label to qualifying label-IPv4 and label-IPv6 routes when:

- these routes are the best path for their prefix in the respective RIB
- there is a BGP next-hop change

A label-IPv4 or label-IPv6 route advertised with a pre-prefix label supports ECMP forwarding across multiple BGP next-hops.

The **pop** option is applicable in route-table-import policies. The advertised BGP label is programmed for a pop operation when:

- a /32 IPv4 static, OSPF, or IS-IS route is matched and accepted by a label-IPv4 or label-IPv6 RIB route-table-import policy entry or default-action with this command
- the route is a candidate to be advertised as a label-IPv4 or label-IPv6 route (due to a BGP export policy)

When the label-IPv4 RIB imports a /32 static, OSPF, or IS-IS route and then exports the route as a BGP route, the default behavior is to program a swap operation in the datapath, which swaps the BGP label with the tunnel label that takes traffic to the destination of the /32 route.

The **pop-and-forward** option is applicable in route-table-import policies, when these policies match an unlabeled BGP route and apply this policy action.

Use the **pop-and-forward** option to program the label that is advertised in the BGP-LU route to forward the packet according to the resolution of the unlabeled route that triggered the origination of the BGP-LU route. The forwarding is done without an IP FIB lookup, which can be useful in situations where the IP FIB at the exit of the MPLS tunnel is not synchronized with the FIB at the head-end of the MPLS tunnel. The advertisement of a pop-and-forward label overrides the configuration to advertise label-ipv6 routes with an explicit null label and the configuration to advertise BGP-LU with a prefix SID attribute. Those features are not available when using the pop-and-forward label.

Default

no advertise-label

Parameters

per-prefix

Sets the per-prefix label allocation for matched routes. This takes effect only in VRF export policies and BGP import policies, and only for certain types of routes.

pop

Sets the pop label allocation for matched routes. This takes effect only in label-IPv4 route-table-import policies and only applies to /32 IPv4 routes that were learned through static configuration, OSPF, or IS-IS.

pop-and-forward

Sets the pop-and-forward label allocation for matched routes. This takes effect only when an unlabeled BGP IPv4 or IPv6 route is matched by a label-IPv4 or label-IPv6 route-table-import policy.

Platforms

7705 SAR Gen 2

4.71 advertise-ldp-prefix

advertise-ldp-prefix

Syntax

[**no**] **advertise-ldp-prefix**

Context

[Tree] (config>router>bgp>group>neighbor advertise-ldp-prefix)

Full Context

configure router bgp group neighbor advertise-ldp-prefix

Description

This command, when configured for a session that supports the IPv4 labeled-unicast address family, allows (subject to BGP export policies) active /32 LDP FEC prefixes to be advertised to the BGP peer with an RFC 8277 label, even though there may be BGP paths for the same prefix.

Default

no advertise-ldp-prefix

Platforms

7705 SAR Gen 2

4.72 advertise-ne-profile

advertise-ne-profile

Syntax

advertise-ne-profile *name*

no advertise-ne-profile

Context

[Tree] (config>service>vprn>ospf>area advertise-ne-profile)

Full Context

configure service vprn ospf area advertise-ne-profile

Description

This command enables advertising of a specific NE profile using OSPFv2 LSA type 10 opaque.

The **no** version of this command disables advertising of NE profiles.

Default

no advertise-ne-profile

Parameters

name

Specifies the name of the NE profile to be advertised, up to 32 characters.

Platforms

7705 SAR Gen 2

4.73 advertise-passive-only

advertise-passive-only

Syntax

[**no**] **advertise-passive-only**

Context

[Tree] (config>service>vprn>isis advertise-passive-only)

Full Context

configure service vprn isis advertise-passive-only

Description

This command enables IS-IS for the VPRN instance to advertise only prefixes that belong to passive interfaces.

The **no** form of this command disables IS-IS for the VPRN instance from advertising only prefixes that belong to passive interfaces.

Platforms

7705 SAR Gen 2

advertise-passive-only

Syntax

[**no**] **advertise-passive-only**

Context

[Tree] (config>router>isis advertise-passive-only)

Full Context

configure router isis advertise-passive-only

Description

This command enables and disables IS-IS to advertise only prefixes that belong to passive interfaces.

Default

no advertise-passive-only

Platforms

7705 SAR Gen 2

4.74 advertise-router-capability

advertise-router-capability

Syntax

advertise-router-capability {**area** | **as**}

no advertise-router-capability

Context

[Tree] (config>service>vprn>isis advertise-router-capability)

[Tree] (config>service>vprn>isis>level advertise-router-capability)

Full Context

configure service vprn isis advertise-router-capability

configure service vprn isis level advertise-router-capability

Description

This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A new TLV as defined in RFC 4971 advertises the TE Node Capability Descriptor capability.

The parameters (area & as) control the scope of the capabilities advertisements.

The **no** form of this command disables this capability.

Default

no advertise-router-capability

Parameters

area

Capabilities are only advertised within the area of origin.

as

Capabilities are only advertised throughout the entire autonomous system.

Platforms

7705 SAR Gen 2

advertise-router-capability

Syntax

advertise-router-capability

advertise-router-capability {**link** | **area** | **as**}

no advertise-router-capability

Context

[Tree] (config>service>vprn>ospf>area advertise-router-capability)

[Tree] (config>service>vprn>ospf>area>if advertise-router-capability)

[Tree] (config>service>vprn>ospf3 advertise-router-capability)

[Tree] (config>service>vprn>ospf advertise-router-capability)

[Tree] (config>service>vprn>ospf3>area>if advertise-router-capability)

Full Context

configure service vprn ospf area advertise-router-capability

configure service vprn ospf area interface advertise-router-capability

configure service vprn ospf3 advertise-router-capability

configure service vprn ospf advertise-router-capability

configure service vprn ospf3 area interface advertise-router-capability

Description

This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A Router Information (RI) LSA as defined in RFC 4970 advertises the following capabilities:

- OSPF graceful restart capable: no
- OSPF graceful restart helper: yes, when enabled
- OSPF Stub Router support: yes
- OSPF Traffic Engineering support: yes, when enabled
- OSPF point-to-point over LAN: yes
- OSPF Experimental TE: no

The parameters (**link**, **area** and **as**) control the advertisement scope of the router capabilities.

The **no** form of this command disables this capability.

Default

no advertise-router-capability

Parameters

link

Capabilities are only advertised over local link and not flooded beyond.

area

Capabilities are only advertised within the area of origin.

as

Capabilities are only advertised throughout the entire autonomous system.

Platforms

7705 SAR Gen 2

advertise-router-capability

Syntax

advertise-router-capability {**area** | **as**}

no advertise-router-capability

Context

[Tree] (config>router>isis advertise-router-capability)

Full Context

configure router isis advertise-router-capability

Description

This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A TLV as defined in RFC 4971 advertises the TE Node Capability Descriptor capability.

The parameters (area and as) control the scope of the capability advertisements.

The **no** form of this command disables this capability.

Parameters

area

Specifies to only advertise within the area of origin.

as

Specifies to advertise throughout the entire autonomous system.

Platforms

7705 SAR Gen 2

advertise-router-capability

Syntax

[**no**] **advertise-router-capability**

Context

[Tree] (config>router>isis>level advertise-router-capability)

Full Context

configure router isis level advertise-router-capability

Description

This command enables router advertisement capabilities.

The **no** form of this command disables router advertisement capabilities.

Default

advertise-router-capability

Platforms

7705 SAR Gen 2

advertise-router-capability

Syntax

advertise-router-capability {**link** | **area** | **as**}

no advertise-router-capability

Context

[Tree] (config>router>ospf3 advertise-router-capability)

[Tree] (config>router>ospf advertise-router-capability)

Full Context

configure router ospf3 advertise-router-capability

configure router ospf advertise-router-capability

Description

This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A Router Information (RI) LSA as defined in RFC 4970 advertises the following capabilities:

- OSPF graceful restart capable: no

- OSPF graceful restart helper: yes, when enabled
- OSPF stub router support: yes
- OSPF traffic engineering support: yes, when enabled
- OSPF point-to-point over LAN: yes
- OSPF experimental TE: no

The parameters (**link**, **area** and **as**) control the scope of the capability advertisements.

The **no** form of this command disables this capability.

Default

no advertise-router-capability

Parameters

link

capabilities are only advertised over local links and not flooded beyond.

area

capabilities are only advertised within the area of origin.

as

capabilities are advertised throughout the entire autonomous system.

Platforms

7705 SAR Gen 2

advertise-router-capability

Syntax

[**no**] **advertise-router-capability**

Context

[Tree] (config>router>ospf>area advertise-router-capability)

[Tree] (config>router>ospf3>area>interface advertise-router-capability)

[Tree] (config>router>ospf3>area advertise-router-capability)

[Tree] (config>router>ospf>area>interface advertise-router-capability)

Full Context

configure router ospf area advertise-router-capability

configure router ospf3 area interface advertise-router-capability

configure router ospf3 area advertise-router-capability

configure router ospf area interface advertise-router-capability

Description

This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A Router Information (RI) LSA as defined in RFC 4970 advertises the following capabilities:

- OSPF graceful restart capable: no
- OSPF graceful restart helper: yes, when enabled
- OSPF stub router support: yes
- OSPF traffic engineering support: yes, when enabled
- OSPF point-to-point over LAN: yes
- OSPF experimental TE: no

The **no** form of this command disables this capability.

Default

advertise-router-capability

Platforms

7705 SAR Gen 2

4.75 advertise-stale-to-all-neighbors

advertise-stale-to-all-neighbors

Syntax

advertise-stale-to-all-neighbors [**without-no-export**]

no advertise-stale-to-all-neighbors

Context

[Tree] (config>service>vprn>bgp>group>graceful-restart>long-lived advertise-stale-to-all-neighbors)

[Tree] (config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived advertise-stale-to-all-neighbors)

[Tree] (config>service>vprn>bgp>graceful-restart>long-lived advertise-stale-to-all-neighbors)

Full Context

configure service vprn bgp group graceful-restart long-lived advertise-stale-to-all-neighbors

configure service vprn bgp group neighbor graceful-restart long-lived advertise-stale-to-all-neighbors

configure service vprn bgp graceful-restart long-lived advertise-stale-to-all-neighbors

Description

This command allows BGP routes marked as LLGR stale to be advertised to BGP peers that did not advertise the LLGR capability when the session was opened. The **no** version of this command causes advertisement behavior to follow the rule that stale routes cannot be advertised to a peer that does not understand or implement the LLGR capability. Stale routes are withdrawn towards such peers.

When this command is configured with the **without-no-export** option, LLGR stales routes can be advertised to any peer (EBGP or IBGP) that did not signal the LLGR capability. Towards IBGP and confederation-EBGP peers that did not advertise the LLGR capability, the LOCAL_PREFERENCE attribute in the advertised stale routes is automatically set to zero.

When this command is configured without the **without-no-export** option, LLGR stale routes are not advertised to any EBGP peer that did not signal the LLGR capability. Towards IBGP and confederation-EBGP peers that did not advertise the LLGR capability the LOCAL_PREFERENCE attribute in the advertised stale routes is automatically set to zero and a NO_EXPORT standard community is automatically added to the routes.

Default

no advertise-stale-to-all-neighbors

Parameters

without-no-export

Allows LLGR stale routes to be advertised to all peers, such that they can exit the local AS.

Platforms

7705 SAR Gen 2

advertise-stale-to-all-neighbors

Syntax

advertise-stale-to-all-neighbors [**without-no-export** | **no without-no-export**]

no advertise-stale-to-all-neighbors

Context

[Tree] (config>router>bgp>group>neighbor>graceful-restart>long-lived advertise-stale-to-all-neighbors)

[Tree] (config>router>bgp>group>graceful-restart>long-lived advertise-stale-to-all-neighbors)

[Tree] (config>router>bgp>graceful-restart>long-lived advertise-stale-to-all-neighbors)

Full Context

configure router bgp group neighbor graceful-restart long-lived advertise-stale-to-all-neighbors

configure router bgp group graceful-restart long-lived advertise-stale-to-all-neighbors

configure router bgp graceful-restart long-lived advertise-stale-to-all-neighbors

Description

This command allows BGP routes marked as LLGR stale to be advertised to BGP peers that did not advertise the LLGR capability when the session was opened.

When this command is configured with the **without-no-export** option, LLGR stale routes can be advertised to any peer (EBGP or IBGP) that did not signal the LLGR capability. Towards IBGP and confederation-EBGP peers that did not advertise the LLGR capability, the LOCAL_PREFERENCE attribute in the advertised stale routes is automatically set to zero.

When this command is configured without the **without-no-export** option, LLGR stale routes are not advertised to any EBGP peer that did not signal the LLGR capability. Towards IBGP and confederation-EBGP peers that did not advertise the LLGR capability the LOCAL_PREFERENCE attribute in the advertised stale routes is automatically set to zero and a NO_EXPORT standard community is automatically added to the routes.

The **no** version of this command causes advertisement behavior to follow the rule that stale routes cannot be advertised to a peer that does not understand or implement the LLGR capability. Stale routes are withdrawn towards such peers.

Default

no advertise-stale-to-all-neighbors

Parameters

without-no-export

Allows LLGR stale routes to be advertised to all peers, such that they can exit the local AS.

Platforms

7705 SAR Gen 2

4.76 advertise-subnet

advertise-subnet

Syntax

[**no**] **advertise-subnet**

Context

[Tree] (config>service>vprn>ospf>area>if advertise-subnet)

Full Context

configure service vprn ospf area interface advertise-subnet

Description

This command enables advertising point-to-point interfaces as subnet routes (network number and mask). When disabled, point-to-point interfaces are advertised as host routes.

This command is not supported in the OSPF3 context.

The **no** form of this command disables advertising point-to-point interfaces as subnet routes meaning they are advertised as host routes.

Default

advertise-subnet — Advertises point-to-point interfaces as subnet routes.

Platforms

7705 SAR Gen 2

advertise-subnet

Syntax

[**no**] **advertise-subnet**

Context

[Tree] (config>router>ospf>area>interface advertise-subnet)

Full Context

configure router ospf area interface advertise-subnet

Description

This command enables advertising point-to-point interfaces as subnet routes (network number and mask). When disabled, point-to-point interfaces are advertised as host routes.

The **no** form of this command disables advertising point-to-point interfaces as subnet routes meaning they are advertised as host routes.

Default

advertise-subnet

Platforms

7705 SAR Gen 2

4.77 advertise-tunnel-link

advertise-tunnel-link

Syntax

[**no**] **advertise-tunnel-link**

Context

[Tree] (config>router>ospf advertise-tunnel-link)

[Tree] (config>router>isis advertise-tunnel-link)

Full Context

configure router ospf advertise-tunnel-link

configure router isis advertise-tunnel-link

Description

This command enables the forwarding adjacency feature. With this feature, IS-IS or OSPF advertises an RSVP LSP as a link so that other routers in the network can include it in their SPF computations. The RSVP LSP is advertised as an unnumbered point-to-point link and the link LSP or LSA has no Traffic Engineering opaque sub-TLVs, as per RFC 3906. An SR-TE LSP is not supported with forwarding adjacency.

The forwarding adjacency feature can be enabled independently from the IGP shortcut feature in CLI. If both **igp-shortcut** and **advertise-tunnel-link** options are enabled for a given IGP instance, then the **advertise-tunnel-link** takes precedence.

When the forwarding adjacency feature is enabled, each node advertises a p2p unnumbered link for each best metric tunnel to the router ID of any endpoint node. The node does not include the tunnels as IGP shortcuts in SPF computation directly. Instead, when the LSA or LSP that advertises the corresponding P2P unnumbered link is installed in the local routing database, the node performs an SPF using it like any other link LSA or LSP. The bidirectional check of the link requires that a link, regular or tunnel, exists in the reverse direction for the tunnel to be used in SPF.

The **igp-shortcut** option under the LSP name governs the use of the LSP with both the **igp-shortcut** and the **advertise-tunnel-link** options in IGP. In other words, the user can exclude a specific RSVP LSP from being used as a forwarding adjacency by entering the command **config>router>mpls>lsp>no igp-shortcut**.

Support is provided for resolving and forwarding IPv4 and IPv6 prefixes over IPv4 forwarding adjacency RSVP-TE LSP. Specifically, the forwarding adjacency feature supports family IPv4 in OSPFv2, family IPv6 in OSPFv3, families IPv4 and IPv6 in ISIS MT=0, and family IPv6 in ISIS MT=2.

In addition, both IPv4 and IPv6 SR-ISIS tunnels can be resolved and further tunneled over one or more RSVP-TE LSPs used as forwarding adjacencies. This is enabled by configuring both segment routing and forwarding adjacency features within an IS-IS instance in a multi-topology MT=0.

IS-IS forwarding adjacency using the **advertise-tunnel-link** command is not supported in combination with the IS-IS link bundling and the IS-IS metric link quality adjustment features.

The **no** form of this command disables forwarding adjacency and disables the advertisement of RSVP LSP into IGP.

Default

no advertise-tunnel-link

Platforms

7705 SAR Gen 2

4.78 advertised-stale-time

advertised-stale-time

Syntax

advertised-stale-time *seconds*

no advertised-stale-time

Context

[Tree] (config>service>vprn>bgp>graceful-restart>long-lived advertised-stale-time)

[Tree] (config>service>vprn>bgp>group>graceful-restart>long-lived advertised-stale-time)

[Tree] (config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived>family advertised-stale-time)

[Tree] (config>service>vprn>bgp>group>graceful-restart>long-lived>family advertised-stale-time)

[Tree] (config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived advertised-stale-time)

[Tree] (config>service>vprn>bgp>graceful-restart>long-lived>family advertised-stale-time)

Full Context

configure service vprn bgp graceful-restart long-lived advertised-stale-time

configure service vprn bgp group graceful-restart long-lived advertised-stale-time

configure service vprn bgp group neighbor graceful-restart long-lived family advertised-stale-time

configure service vprn bgp group graceful-restart long-lived family advertised-stale-time

configure service vprn bgp group neighbor graceful-restart long-lived advertised-stale-time

configure service vprn bgp graceful-restart long-lived family advertised-stale-time

Description

This command sets the value of the long-lived stale time that is advertised by the router in its LLGR capability. When configured in the long-lived configuration context, **advertised-stale-time** applies to all AFI/SAFI in the advertised LLGR capability except for any AFI/SAFI with a family-specific override. A family-specific override is configured with the **advertised-stale-time** command in a family context.

The **no** version of this command sets the **advertised-stale-time** value to 24 hours (86400 seconds).

Default

no advertised-stale-time

Parameters

seconds

Specifies the advertised long-lived stale time in seconds.

Values 0 to 16777215

Platforms

7705 SAR Gen 2

advertised-stale-time

Syntax

advertised-stale-time *seconds*

no advertised-stale-time

Context

[Tree] (config>router>bgp>group>neighbor>graceful-restart>long-lived advertised-stale-time)

[Tree] (config>router>bgp>group>graceful-restart>long-lived>family advertised-stale-time)

[Tree] (config>router>bgp>graceful-restart>long-lived advertised-stale-time)

[Tree] (config>router>bgp>graceful-restart>long-lived>family advertised-stale-time)

[Tree] (config>router>bgp>group>graceful-restart>long-lived advertised-stale-time)

[Tree] (config>router>bgp>group>neighbor>graceful-restart>long-lived>family advertised-stale-time)

Full Context

configure router bgp group neighbor graceful-restart long-lived advertised-stale-time

configure router bgp group graceful-restart long-lived family advertised-stale-time

configure router bgp graceful-restart long-lived advertised-stale-time

configure router bgp graceful-restart long-lived family advertised-stale-time

configure router bgp group graceful-restart long-lived advertised-stale-time

configure router bgp group neighbor graceful-restart long-lived family advertised-stale-time

Description

This command sets the value of the long-lived stale time that is advertised by the router in its LLGR capability. When configured in the long-lived configuration context, **advertised-stale-time** applies to all AFI/SAFI in the advertised LLGR capability except for any AFI/SAFI with a family-specific override. A family-specific override is configured with the **advertised-stale-time** command in a **family** context.

The **no** version of this command sets the **advertised-stale-time** value to 24 hours (86400 seconds).

Default

no advertised-stale-time

Parameters

seconds

Specifies the advertised long-lived stale time in seconds.

Values 0 to 16777215

Platforms

7705 SAR Gen 2

4.79 agg-rate

agg-rate

Syntax

[**no**] **agg-rate**

Context

[Tree] (config>service>ies>if>sap>egress agg-rate)

Full Context

configure service ies interface sap egress agg-rate

Description

Commands in this context configure aggregation rate parameters. This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, **and queue-frame-based-accounting**.

When specified under a Vport, the **agg-rate**, **port-scheduler-policy** and **scheduler-policy** commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an **agg-rate** or **port-scheduler-policy** involves removing the existing command and applying the new command.

The **no** form of this command disables the aggregation rate.

Platforms

7705 SAR Gen 2

agg-rate

Syntax

[**no**] **agg-rate**

Context

[Tree] (config>service>epipe>sap>egress agg-rate)

Full Context

configure service epipe sap egress agg-rate

Description

This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, and **queue-frame-based-accounting**.

Platforms

7705 SAR Gen 2

agg-rate

Syntax

[**no**] **agg-rate**

Context

[Tree] (config>service>vpls>sap>egress agg-rate)

Full Context

configure service vpls sap egress agg-rate

Description

This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, **and queue-frame-based-accounting**.

Platforms

7705 SAR Gen 2

agg-rate

Syntax

[**no**] **agg-rate**

Context

[Tree] (config>service>vprn>if>sap>egress agg-rate)

Full Context

configure service vprn interface sap egress agg-rate

Description

This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, **and queue-frame-based-accounting**.

Platforms

7705 SAR Gen 2

4.80 aggregate

aggregate

Syntax

aggregate *ip-prefix/ip-prefix-length* [**summary-only**] [**as-set**] [**aggregator** *as-number:ip-address*] [**discard-component-communities**] [**black-hole** [**generate-icmp**]] [**community** *comm-id* [*comm-id*] [**local-preference** *local-pref*]] [**description** *description*] [**tunnel-group** *tunnel-group-id*]

aggregate *ip-prefix/ip-prefix-length* [**summary-only**] [**as-set**] [**aggregator** *as-number:ip-address*] [**discard-component-communities**] [**community** *comm-id* [*comm-id*]] [**indirect** *ip-address*] [**local-preference** *local-pref*]] [**description** *description*] [**tunnel-group** *tunnel-group-id*]

no aggregate *ip-prefix/ip-prefix-length*

Context

[Tree] (config>service>vprn aggregate)

Full Context

configure service vprn aggregate

Description

This command creates an aggregate route. Use this command to automatically install an aggregate route in the routing table when there are one or more component routes. A component route is any route used for forwarding that is a more specific match of the aggregate.

The use of aggregate routes can reduce the number of routes that need to be advertised to neighbor routers, leading to smaller routing table sizes.

Overlapping aggregate routes may be configured; in this case a route becomes a component of only the one aggregate route with the longest prefix match. For example if one aggregate is configured as 10.0.0.0/16 and another as 10.0.0.0/24, then route 10.0.128/17 would be aggregated into 10.0.0.0/16, and route 10.0.0.128/25 would be aggregated into 10.0.0.0/24. If multiple entries are made with the same prefix and the same mask the previous entry is overwritten.

A list of up to 12 BGP communities (any mix of standard, extended, and large communities) may be associated with an aggregate route. These communities can be matched in route policies and are automatically added to BGP routes that are created from the aggregate route.

By default, aggregate routes are not installed in the forwarding table, however there are configuration options that allow an aggregate route to be installed with a black-hole next hop or with an indirect IP address as next hop.

Aggregate routes can be advertised via MP-BGP to other PEs within the network. Aggregate routes advertised using MP-BGP do not include aggregated BGP path attributes from the component routes which were used to activate the aggregate route. The aggregate route will be advertised with the minimal set of path attributes as if the aggregate was originated by the advertising routes. Export route policies should be used to control and modify the advertisement and path attributes of the aggregate routes.

The **no** form of this command removes the aggregate.

Default

no aggregate

Parameters

ip-prefix

The destination address of the aggregate route in dotted decimal notation.

Values

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 32
ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6-prefix-length	0 to 128

the mask associated with the network address expressed as a mask length

Values: 0 to 32

summary-only

This optional parameter suppresses advertisement of more specific component routes for the aggregate.

To remove the **summary-only** option, enter the same aggregate command without the **summary-only** parameter.

as-set

This optional parameter is only applicable to BGP and creates an aggregate where the path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized. Use this feature carefully as it can increase the amount of route churn due to best path changes.

aggregator *as-number*:*ip-address*

This optional parameter specifies the BGP aggregator path attribute to the aggregate route. When configuring the aggregator, a two-octet AS number used to form the aggregate route must be entered, followed by the IP address of the BGP system that created the aggregate route.

discard-component-communities

This optional keyword causes the aggregate to be advertised with only the configured BGP community set, none of the communities from the component routes activating the aggregate are included. (Component attributes are never included in aggregate routes advertised to other PE routers via MP-BGP).

black-hole

This optional parameter installs the aggregate route, when activated, in the FIB with a black-hole next-hop, where packets matching this route are discarded.

generate-icmp

This optional parameter keyword generates an ICMP.

community

This configuration option associates a BGP community with the aggregate route. The community can be matched in route policies and is automatically added to BGP routes exported from the aggregate route.

comm-id

Specifies a BGP community value, up to 72 characters.

Values [*as-num:comm-val* | *well-known-comm* | *ext-comm* | *large-comm*]

where:

- *as-num* — 0 to 65535
- *comm-val* — 0 to 65535
- *well-known-comm* — **null** | **no-export** | **no-export-subconfed** | **no-advertise** | **llgr-stale** | **no-llgr** | **blackhole**
- *ext-comm* — the extended community, defined as one of the following:
  - {*target* | *origin*}:*ip-address*:*comm-val*
  - {*target* | *origin*}:*asnum*:*ext-comm-val*
  - {*target* | *origin*}:*ext-asnum*:*comm-val*
  - **bandwidth**:*asnum*:*val-in-mbps*
  - **ext:4300**:*ovstate*
  - **ext**:*value1*:*value2*
  - color:co-bits:*color-value*

  where:
  - *target* — route target
  - *origin* — route origin
  - *ip-address* — a.b.c.d
  - *ext-comm-val* — 0 to 4294967295
  - *ext-asnum* — 0 to 4294967295
  - *val-in-mbps* — 0 to 16777215
  - *ovstate* — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid)
  - *value1* — 0000 to FFFF
  - *value2* — 0 to FFFFFFFFFFFF
  - *co-bits* — 00, 01, 10 or 11
  - *color-value* — 0 to 4294967295
- *large-comm — asn-or-ex:val-or-ex:val-or-ex*

description

Specifies a text description stored in the configuration file for a configuration context.

local-preference

Specifies a BGP local-preference value with the aggregate route. The local-preference overrides the default local preference value of a BGP route originated by exporting the aggregate route.

Values 0 to 4294967295

indirect *ip-address*

This configuration option specifies that the aggregate route should be installed in the FIB with a next-hop taken from the route used to forward packets to ip-address.

Values		
	ipv4-prefix	a.b.c.d
	ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x: [0 to FFFF]H
		d: [0 to 255]D

tunnel-group-id

Specifies that the MC-IPsec state of the specific tunnel-group is added to the aggregate route.

Values 1 to 16

Platforms

7705 SAR Gen 2

aggregate

Syntax

aggregate *ip-prefix/ip-prefix-length* [**summary-only**] [**as-set**] [**aggregator** *as-number:ip-address*] [**discard-component-communities**] [**black-hole** [**generate-icmp**]] [**community** *comm-id* [*comm-id*]] [**description** *description*] [**local-preference** *local-pref*erence] [**policy** *policy-name*]

aggregate *ip-prefix/ip-prefix-length* [**summary-only**] [**as-set**] [**aggregator** *as-number:ip-address*] [**discard-component-communities**] [**community** *comm-id* [*comm-id*]] [**indirect** *ip-address*] [**description** *description*] [**local-preference** *local-preference*] [**policy** *policy-name*]

no aggregate *ip-prefix/ip-prefix-length*

Context

[Tree] (config>router aggregate)

Full Context

configure router aggregate

Description

This command creates an aggregate route.

Use this command to automatically install an aggregate route in the routing table when there are one or more component routes. A component route is any route used for forwarding that is a more-specific match of the aggregate.

The use of aggregate routes can reduce the number of routes that need to be advertised to neighbor routers, leading to smaller routing table sizes.

Overlapping aggregate routes may be configured; in this case a route becomes a component of only the one aggregate route with the longest prefix match. For example if one aggregate is configured as 10.0.0.0/16 and another as 10.0.0.0/24, then route 10.0.128/17 would be aggregated into 10.0.0.0/16, and route 10.0.0.128/25 would be aggregated into 10.0.0.0/24. If multiple entries are made with the same prefix and the same mask the previous entry is overwritten.

A standard 4-byte BGP community may be associated with an aggregate route in order to facilitate route policy matching.

By default aggregate routes are not installed in the forwarding table, however there are configuration options that allow an aggregate route to be installed with a black-hole next hop or with an indirect IP address as next hop.

The **no** form of this command removes the aggregate.

Default

no aggregate

Parameters

ip-prefix

Specifies the destination address of the aggregate route in dotted decimal notation.

Values	ipv4-prefix	a.b.c.d (host bits must be 0)	
	ipv4-prefix-length	0 to 32	
	ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)	
		x:x:x:x:x:x:d.d.d.d	
		x:	[0 to FFFF]H
		d:	[0 to 255]D
	ipv6-prefix-length	0 to 128	

ip-prefix-length

Specifies the mask associated with the network address expressed as a mask length.

Values 0 to 32

summary-only

Suppresses advertisement of more specific component routes for the aggregate.

To remove the **summary-only** option, enter the same aggregate command without the **summary-only** parameter.

as-set

This optional parameter is only applicable to BGP and creates an aggregate where the path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized. Use this feature carefully as it can increase the amount of route churn due to best path changes.

as-number:ip-address

Specifies the BGP aggregator path attribute to the aggregate route. When configuring the aggregator, a two-octet AS number used to form the aggregate route must be entered, followed by the IP address of the BGP system that created the aggregate route.

discard-component-communities

Causes the aggregate to be advertised with only the configured BGP community set, none of the communities from the component routes activating the aggregate are included.

black-hole

Installs the aggregate route, when activated, in the FIB with a black-hole next-hop, where packets matching this route are discarded.

generate-icmp

Mandatory keyword to generate an ICMP.

community

Associates a BGP community with the aggregate route. The community can be matched in route policies and is automatically added to BGP routes exported from the aggregate route.

comm-id

Specifies a BGP community value, up to 72 characters. A maximum of twelve community IDs can be specified in a single statement.

Values [*as-num:comm-val* | *well-known-comm* | *ext-comm* | *large-comm*]

where:

- *as-num* — 0 to 65535
- *comm-val* — 0 to 65535
- *well-known-comm* — **null** | **no-export** | **no-export-subconfed** | **no-advertise** | **llgr-stale** | **no-llgr** | **blackhole**
- *ext-comm* — the extended community, defined as one of the following:
  - {*target* | *origin*}:*ip-address*:*comm-val*
  - {*target* | *origin*}:*asnum*:*ext-comm-val*
  - {*target* | *origin*}:*ext-asnum*:*comm-val*
  - **bandwidth**:*asnum*:*val-in-mbps*
  - **ext:4300**:*ovstate*
  - **ext**:*value1*:*value2*
  - color:co-bits:*color-value*

where:

- *target* — route target
- *origin* — route origin
- *ip-address* — a.b.c.d
- *ext-comm-val* — 0 to 4294967295
- *ext-asnum* — 0 to 4294967295
- *val-in-mbps* — 0 to 16777215
- *ovstate* — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid)
- *value1* — 0000 to FFFF
- *value2* — 0 to FFFFFFFFFFFF
- *co-bits* — 00, 01, 10 or 11
- *color-value* — 0 to 4294967295

• *large-comm — asn-or-ex:val-or-ex:val-or-ex*

indirect *ip-address*

Specifies that the aggregate route should be installed in the FIB with a next-hop taken from the route used to forward packets to ip-address.

Values		
	ipv4-prefix	a.b.c.d
	ipv6-prefix	x:x:x:x:x:x:x:x
		x:x:x:x:x:x:d.d.d.d
		x: [0 to FFFF]H
		d: [0 to 255]D

description

Specifies a text description stored in the configuration file for a configuration context, up to 80 characters.

local-preference

Specifies a BGP local-preference value with the aggregate route. The local-preference overrides the default local preference value of a BGP route originated by exporting the aggregate route.

Values 0 to 4294967295

policy-name

Specifies the route policy, up to 64 characters.

Platforms

7705 SAR Gen 2

4.81 aggregate-prefix-match

aggregate-prefix-match

Syntax

[**no**] **aggregate-prefix-match**

Context

[Tree] (config>router>ldp aggregate-prefix-match)

Full Context

configure router ldp aggregate-prefix-match

Description

The command enables the use by LDP of the aggregate prefix match procedures.

When this option is enabled, LDP performs the following procedures for all prefixes. When an LSR receives a FEC-label binding from an LDP neighbor for a given specific FEC1 element, it will install the binding in the LDP FIB if:

- It is able to perform a successful longest IP match of the FEC prefix with an entry in the routing table, and
- The advertising LDP neighbor is the next-hop to reach the FEC prefix.

When such a FEC-label binding has been installed in the LDP FIB, then LDP programs an NHLFE entry in the egress data path to forward packets to FEC1. It also advertises a new FEC-label binding for FEC1 to all its LDP neighbors.

When a new prefix appears in the routing table, LDP inspects the LDP FIB to determine if this prefix is a better match (a more specific match) for any of the installed FEC elements. For any FEC for which this is true, LDP may have to update the NHLFE entry for this FEC.

When a prefix is removed from the routing table, LDP inspects the LDP FIB for all FEC elements which matched this prefix to determine if another match exists in the routing table. If so, it updates the NHLFE entry accordingly. If not, it sends a label withdraw message to its LDP neighbors to remove the binding.

When the next hop for a routing prefix changes, LDP updates the LDP FIB entry for the FEC elements which matched this prefix. It also updates the NHLFE entry for these FEC elements accordingly.

The **no** form of this command disables the use by LDP of the aggregate prefix procedures and deletes the configuration. LDP resumes performing exact prefix match for FEC elements.

Default

no aggregate-prefix-match

Platforms

7705 SAR Gen 2

4.82 aggregate-used-paths

aggregate-used-paths

Syntax

aggregate-used-paths *family* [*family*]

no aggregate-used-paths

Context

[Tree] (config>service>vprn>bgp>group>neighbor>link-bandwidth aggregate-used-paths)

[Tree] (config>service>vprn>bgp>group>link-bandwidth aggregate-used-paths)

Full Context

configure service vprn bgp group neighbor link-bandwidth aggregate-used-paths

configure service vprn bgp group link-bandwidth aggregate-used-paths

Description

This command configures BGP to aggregate the bandwidth values from the link-bandwidth extended communities of the used multipaths towards an IP prefix when it is re-advertising a route with next-hop-self towards peers within the scope of the command, as long as the route belongs to one of the listed address families.

Aggregation is not supported unless all of the used multipaths (up to the configured ECMP limit) correspond to received BGP routes with a link-bandwidth extended community. If add-path is also enabled toward the peer, then all of the add-paths advertised to the peer encode the aggregated bandwidth in a link-bandwidth extended community.

Up to three families may be configured.

The **no** form of this command disables aggregation in a next-hop-self scenario and the link-bandwidth extended community in the advertised route is a copy of the link-bandwidth extended community in the received route (which may have been added by import policy or by the effect of the **add-to-received-ebgp** command).

Default

no aggregate-used-paths

Parameters

family

Specifies the address families for which receiving the link-bandwidth extended community from EBGP peers should be supported.

Values ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes.

label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes.

ipv6 — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes.

Platforms

7705 SAR Gen 2

aggregate-used-paths

Syntax

aggregate-used-paths *family* [*family*]

no aggregate-used-paths

Context

[Tree] (config>router>bgp>group>neighbor>link-bandwidth aggregate-used-paths)

[Tree] (config>router>bgp>group>link-bandwidth aggregate-used-paths)

Full Context

configure router bgp group neighbor link-bandwidth aggregate-used-paths

configure router bgp group link-bandwidth aggregate-used-paths

Description

This command configures BGP to aggregate the bandwidth values from the link-bandwidth extended communities of the used multipaths towards an IP prefix when it is re-advertising a route with next-hop-self towards peers within the scope of the command, as long as the route belongs to one of the listed address families.

Aggregation is not supported unless all of the used multipaths (up to the configured ECMP limit) correspond to received BGP routes with a link-bandwidth extended community. If add-path is also enabled toward the peer, then all of the add-paths advertised to the peer encode the aggregated bandwidth in a link-bandwidth extended community.

Up to six families may be configured.

The **no** form of this command disables aggregation in a next-hop-self scenario and the link-bandwidth extended community in the advertised route is a copy of the link-bandwidth extended community in the received route (which may have been added by import policy or by the effect of the **add-to-received-ebgp** command).

Default

no aggregate-used-paths

Parameters

family

Specifies the address families for which receiving the link-bandwidth extended community from EBGP peers should be supported.

Values ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes.

label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes.

vpn-ipv4 — Adds a link-bandwidth extended community to IPv4 VPN (SAFI 128) routes.

ipv6 — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes.

label-ipv6 — Adds a link-bandwidth extended community to labeled-unicast IPv6 routes.

vpn-ipv6 — Adds a link-bandwidth extended community to IPv6 VPN (SAFI 128) routes.

Platforms

7705 SAR Gen 2

4.83 aggregator-id-zero

aggregator-id-zero

Syntax

[**no**] **aggregator-id-zero**

Context

[Tree] (config>service>vprn>bgp>group aggregator-id-zero)

[Tree] (config>service>vprn>bgp>group>neighbor aggregator-id-zero)

[Tree] (config>service>vprn>bgp aggregator-id-zero)

Full Context

configure service vprn bgp group aggregator-id-zero

configure service vprn bgp group neighbor aggregator-id-zero

configure service vprn bgp aggregator-id-zero

Description

This command is used to set the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths.

When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute.

When this command is enabled, BGP adds the router ID to the aggregator path attribute. This command is used at the group level to revert to the value defined under the global level, while this command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of this command used at the global level reverts to default where BGP adds the AS number and router ID to the aggregator path attribute.

The **no** form of this command used at the group level reverts to the value defined at the group level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no aggregator-id-zero — BGP adds the AS number and router ID to the aggregator path attribute.

Platforms

7705 SAR Gen 2

aggregator-id-zero

Syntax

[**no**] **aggregator-id-zero**

Context

[Tree] (config>router>bgp aggregator-id-zero)

[Tree] (config>router>bgp>group aggregator-id-zero)

[Tree] (config>router>bgp>group>neighbor aggregator-id-zero)

Full Context

configure router bgp aggregator-id-zero

configure router bgp group aggregator-id-zero

configure router bgp group neighbor aggregator-id-zero

Description

This command sets the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes for the same prefix with different path attributes.

When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute.

When this command is enabled, BGP adds the router ID to the aggregator path attribute. This command is used at the group level to revert to the value defined under the global level, while this command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of this command used at the global level reverts to default where BGP adds the AS number and router ID to the aggregator path attribute.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no aggregator-id-zero

Platforms

7705 SAR Gen 2

4.84 aging

aging

Syntax

aging *days*

no aging

Context

[Tree] (config>system>security>password aging)

Full Context

configure system security password aging

Description

This command configures the number of days a user password is valid before the user must change their password. This parameter can be used to force the user to change the password at the configured interval. Note the aging starts after the last password configuration or update. This timer is persistence (per user) over a node reboot or activity switch between CPMs. When the user changes the password, the timer is reset to the maximum age. When the password for a user ages out, the user is prompted at login to change the password. Console/SSH/Telnet supports password change prompt.

The **no** form of this command reverts to the default value.

Parameters

days

Specifies the maximum number of days the password is valid.

Values 1 to 500

Note:
This command applies to local users.

Platforms

7705 SAR Gen 2

4.85 ah-ext-hdr

ah-ext-hdr

Syntax

ah-ext-hdr {**true** | **false**}

no ah-ext-hdr

Context

[Tree] (config>filter>ipv6-filter>entry>match ah-ext-hdr)

Full Context

configure filter ipv6-filter entry match ah-ext-hdr

Description

This command enables match on existence of AH Extension Header in the IPv6 filter policy.

The **no** form of this command ignores AH Extension Header presence/absence in a packet when evaluating match criteria of a given filter policy entry.

Default

no ah-ext-hdr

Parameters

true

Matches a packet with an AH Extension Header.

false

Matches a packet without an AH Extension Header.

Platforms

7705 SAR Gen 2

4.86 aigp

aigp

Syntax

[**no**] **aigp**

Context

[Tree] (config>router>bgp>group aigp)

[Tree] (config>router>bgp>group>neighbor aigp)

Full Context

configure router bgp group aigp

configure router bgp group neighbor aigp

Description

This command enables or disables Accumulated IGP (AIGP) path attribute support with one or more BGP peers. BGP path selection among routes with an associated AIGP metric is based on the end-to-end IGP metrics of the different BGP paths, even when these BGP paths span more than one AS and IGP instance.

The effect of disabling AIGP (using the **no** form of this command or implicit) is to remove the AIGP attribute from advertised routes, if present, and to ignore the AIGP attribute in received routes.

Default

no aigp

Platforms

7705 SAR Gen 2

4.87 aigp-metric

aigp-metric

Syntax

aigp-metric *metric*

aigp-metric add

aigp-metric igp

no aigp-metric

Context

[Tree] (config>router>policy-options>policy-statement>default-action aigp-metric)

[Tree] (config>router>policy-options>policy-statement>entry>action aigp-metric)

Full Context

configure router policy-options policy-statement default-action aigp-metric

configure router policy-options policy-statement entry action aigp-metric

Description

This command assigns a BGP AIGP metric to routes matching the entry. The effect of this command on a route matched and accepted by a route policy entry depends on how the policy is applied (BGP import policy vs. BGP export policy), the type of route and the specific form of this command.

In a BGP import policy this command is used to:

- Associate an AIGP metric with an IBGP route received with an empty AS path and no AIGP attribute.
- Associate an AIGP metric with an EBGP route received without an AIGP attribute that has an AS path containing only AS numbers belonging to the local AIGP administrative domain.
- Modify the received AIGP metric value prior to BGP path selection.

In a BGP export policy this command is used to:

- Add the AIGP attribute and set the AIGP metric value in a BGP route originated by exporting a direct, static or IGP route from the routing table.
- Remove the AIGP attribute from a route advertisement to a particular peer.
- Modify the AIGP metric value in a route advertisement to a particular peer.

Default

no aigp-metric

Parameters

metric

Administratively defined metric.

Values 0 to 4294967295

Default name — The AIGP metric parameter variable name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@".

add

Adds the AIGP attribute.

igp

Sets the AIGP metric to the IGP metric.

Platforms

7705 SAR Gen 2

4.88 alarm

alarm

Syntax

alarm *rmon-alarm-id* **variable-oid** *oid-string* **interval** *seconds* [*sample-type*] [**startup-alarm** *alarm-type*] [**rising-event** *rmon-event-id* **rising-threshold** *threshold*] [**falling-event** *rmon-event-id* **falling-threshold** *threshold*] [**owner** *owner-string*]

no alarm *rmon-alarm-id*

Context

[Tree] (config>system>thresholds>rmon alarm)

Full Context

configure system thresholds rmon alarm

Description

The alarm command configures an entry in the RMON-MIB alarmTable. The alarm command controls the monitoring and triggering of threshold crossing events. In order for notification or logging of a threshold crossing event to occur there must be at least one associated rmon>event configured.

The agent periodically takes statistical sample values from the MIB variable specified for monitoring and compares them to thresholds that have been configured with the alarm command. The alarm command configures the MIB variable to be monitored, the polling period (interval), sampling type (absolute or delta value), and rising and falling threshold parameters. If a sample has crossed a threshold value, the associated event is generated.

Use the **no** form of this command to remove an rmon-alarm-id from the configuration.

Parameters

rmon-alarm-id

Specifies a numerical identifier for the alarm being configured. The number of alarms that can be created is limited to 1200. Alarm ID values above 65400 are used for dynamic system threshold commands and should be avoided.

Values 1 to 65535

oid-string

Specifies the SNMP object identifier of the particular variable to be sampled. Only SNMP variables that resolve to an ASN.1 primitive type of integer (integer, Integer32, Counter32, Counter64, Gauge, or TimeTicks) may be sampled. The oid-string, up to 255 characters, may be expressed using either the dotted string notation or as object

name plus dotted instance identifier. For example, "1.3.6.1.2.1.2.2.1.10.184582144" or "ifInOctets.184582144".

seconds

Specifies the polling period over which the data is sampled and compared with the rising and falling thresholds. When setting this interval value, care should be taken in the case of 'delta' type sampling - the interval should be set short enough that the sampled variable is very unlikely to increase or decrease by more than 2147483647 - 1 during a single sampling interval. Care should also be taken not to set the interval value too low to avoid creating unnecessary processing overhead.

Values 1 to 2147483647

sample-type

Specifies the method of sampling the selected variable and calculating the value to be compared against the thresholds.

Values absolute — Specifies that the value of the selected variable will be compared directly with the thresholds at the end of the sampling interval.

delta — Specifies that the value of the selected variable at the last sample will be subtracted from the current value, and the difference compared with the thresholds.

Default absolute

alarm-type

Specifies the alarm that may be sent when this alarm is first created.

If the first sample is greater than or equal to the rising threshold value and **startup-alarm** is equal to **rising** or **either**, then a single rising threshold crossing event is generated.

If the first sample is less than or equal to the falling threshold value and **startup-alarm** is equal to **falling** or **either**, a single falling threshold crossing event is generated.

Values rising, falling, either

Default either

rising-event *rmon-event-id*

Specifies the identifier of the **rmon>event** that specifies the action to be taken when a rising threshold crossing event occurs.

If there is no corresponding event configured for the specified rmon-event-id, then no association exists and no action is taken.

If the **rising-event** *rmon-event-id* has a value of zero (0), no associated event exists.

If a **rising-event** *rmon-event-id* is configured, the CLI requires a **rising-threshold** to also be configured.

Values 0 to 65535

Default 0

rising-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is greater than or equal to this threshold and the associated startup-alarm is equal to **rising** or **either**.

After a rising threshold crossing event is generated, another such event will not be generated until the sampled value falls below this threshold and reaches less than or equal the **falling-threshold** value.

Values -2147483648 to 2147483647

Default 0

falling-event *rmon-event-id*

Specifies the identifier of the **rmon>event** that specifies the action to be taken when a falling threshold crossing event occurs. If there is no corresponding event configured for the specified rmon-event-id, then no association exists and no action is taken. If the **falling-event** has a value of zero (0), no associated event exists.

If a **falling-event** is configured, the CLI requires a **falling-threshold** to also be configured.

Values 0 to 65535

Default 0

falling-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is less than or equal to this threshold and the associated **startup-alarm** is equal to **falling** or **either**.

After a falling threshold crossing event is generated, another such event will not be generated until the sampled value rises above this threshold and reaches greater than or equal the rising-threshold value.

Values -2147483648 to 2147483647

Default 0

owner-string

Specifies the owner string; the owner identifies the creator of this alarm. It defaults to "TiMOS CLI". This parameter is defined primarily to allow entries that have been created in the RMON-MIB alarmTable by remote SNMP managers to be saved and reloaded in a CLI configuration file. The owner will not normally be configured by CLI users and can be a maximum of 80 characters long.

Default TiMOS CLI

Configuration example

```
alarm 3 variable-oid ifInOctets.184582144 interval 20 sample-type delta
start-alarm either rising-event 5 rising-threshold 10000 falling-event 5
falling-threshold 9000 owner "TiMOS CLI"
```

Platforms

7705 SAR Gen 2

4.89 alarms

alarms

Syntax

alarms

Context

[Tree] (config>system alarms)

Full Context

configure system alarms

Description

Commands in this context configure facility alarm parameters. Alarm support is intended to cover a focused subset of router states that are likely to indicate service impacts (or imminent service impacts) related to the overall state of hardware assemblies (cards, fans, links, and so on).

Platforms

7705 SAR Gen 2

4.90 alias

alias

Syntax

alias *alias-name alias-command-name*

no alias *alias-name*

Context

[Tree] (environment alias)

Full Context

environment alias

Description

This command enables the substitution of a command line (or part of a command line) by an alias. Use this command to create alternative or easier to remember or understand names for an entity or command string. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. The special characters forward slash (/) and backslash (\) cannot be used as the first character inside an alias string. An alias can contain a double quote character by preceding the quote with a backslash (\) character (for example, **alias my-alias "| match \"***string***\""**). Only a single command can be present in the command string (the command can be long with many parameters but there is no support for aliases that include multiple CLI commands or lines). This command can be entered in any context but must be created in the **root environment** context.

For example, to create an alias named **soi** to display OSPF interfaces, enter the following command:

alias soi "show router ospf interface”

Complex aliases can be created to have shortcuts for customized show routine output.

environment alias my-summary "| match expression \"----|Description|Interface|Admin State|Oper State|Transceiver Type|Optical Compliance|Link Length\" | match invert-match expression \"Ethernet Interface|OTU Interface\" | match invert-match expression \"----\" post-lines 1"

and then used like this:

show port detail my-summary

Parameters

alias-name

Specifies the alias name, up to 80 characters. Do not use a valid command string for the name of the alias. If the alias specified is an actual command, this causes the command to be replaced by the alias.

alias-command-name

Specifies the command name to be associated, up to 320 characters.

Platforms

7705 SAR Gen 2

4.91 align

align

Syntax

[**no**] **align**

Context

[Tree] (config>log>acct-policy align)

Full Context

configure log accounting-policy align

Description

This command enables alignment of statistics collection to the nearest interval within an hour. Enabling the alignment allows statistics collection into an accounting file that is being synchronized across multiple network nodes in the network.

The **no** form of this command disables alignment of statistics collection.

Default

no align

Platforms

7705 SAR Gen 2

4.92 all

all

Syntax

all [**detail**]

no all

Context

[Tree] (debug>router>rsvp>event all)

[Tree] (debug>router>mpls>event all)

Full Context

debug router rsvp event all

debug router mpls event all

Description

This command debugs all events.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about all events.

Platforms

7705 SAR Gen 2

all

Syntax

all [**detail**]

no all

Context

[Tree] (debug>router>rsvp>packet all)

Full Context

debug router rsvp packet all

Description

This command debugs all packets.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about all RSVP packets.

Platforms

7705 SAR Gen 2

all

Syntax

all [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]

no all

Context

[Tree] (debug>router>pim all)

Full Context

debug router pim all

Description

This command enables debugging for all the PIM modules.

The **no** form of this command disables debugging PIM modules.

Parameters

grp-ip-address

Debugs information associated with all PIM modules.

Values IPv4 or IPv6 address

ip-address

Debugs information associated with all PIM modules.

Values IPv4 or IPv6 address

detail

Debugs detailed information on all PIM modules.

Platforms

7705 SAR Gen 2

all

Syntax

[**no**] **all**

Context

[Tree] (debug>router>rpki-session>packet all)

Full Context

debug router rpki-session packet all

Description

This command enables debugging for all RPKI packets.

The **no** form of this command disables debugging for all RPKI packets.

Platforms

7705 SAR Gen 2

all

Syntax

all [**detail**]

no all

Context

[Tree] (debug>router>pcep>pcc all)

[Tree] (debug>router>pcep>pcc>conn all)

Full Context

debug router pcep pcc all

debug router pcep pcc connection all

Description

This command enables debugging for all PCEP PCC or connection events.

The **no** form of this command disables debugging.

Parameters

detail

Keyword used to specify detailed information about all events.

Platforms

7705 SAR Gen 2

5 a Commands – Part II

5.1 all-events

all-events

Syntax

all-events

Context

[Tree] (debug>service>id>stp all-events)

Full Context

debug service id stp all-events

Description

This command enables STP debugging for all events.

The **no** form of the command disables debugging.

Platforms

7705 SAR Gen 2

5.2 all-l1isis

all-l1isis

Syntax

all-l1isis *ieee-address*

no all-l1isis

Context

[Tree] (config>service>vprn>isis all-l1isis)

Full Context

configure service vprn isis all-l1isis

Description

This command specifies the MAC address to use for the VPRN instance of the Layer 1 IS-IS routers. The MAC address should be a multicast address.

The **no** form of this command reverts to the default value.

Default

all-l1isis 01:80:c2:00:00:14

Parameters

ieee-address

Specifies the destination MAC address for all Layer 1 I-IS neighbors on the link for this ISIS instance.

Platforms

7705 SAR Gen 2

all-l1isis

Syntax

all-l1isis *ieee-address*

no all-l1isis

Context

[Tree] (config>router>isis all-l1isis)

Full Context

configure router isis all-l1isis

Description

This command enables you to specify the MAC address to use for all Layer 1 IS-IS routers. The MAC address should be a multicast address.

The **no** form of this command reverts to the default value.

Default

01:80:c2:00:00:14

Parameters

ieee-address

Specifies the destination MAC address for all Layer 1 I-IS neighbors on the link for this IS-IS instance.

Platforms

7705 SAR Gen 2

5.3 all-l2isis

all-l2isis

Syntax

all-l2isis *ieee-address*

no all-l2isis

Context

[Tree] (config>service>vprn>isis all-l2isis)

Full Context

configure service vprn isis all-l2isis

Description

This command specifies the MAC address to use for Layer 2 IS-IS routers for the VPRN instance. The MAC address should be a multicast address.

The **no** form of this command reverts to the default value.

Default

all-l2isis 01:80:c2:00:00:15

Parameters

ieee-address

Specifies the destination MAC address for all Layer 2 ISIS neighbors on the link for this ISIS instance.

Platforms

7705 SAR Gen 2

all-l2isis

Syntax

all-l2isis *ieee-address*

no all-l2isis

Context

[Tree] (config>router>isis all-l2isis)

Full Context

configure router isis all-l2isis

Description

This command enables you to specify the MAC address to use for all Layer 2 IS-IS routers. The MAC address should be a multicast address.

The **no** form of this command reverts to the default value.

Default

01:80:c2:00:00:15

Parameters

ieee-address

Specifies the destination MAC address for all Layer 2 IS-IS neighbors on the link for this IS-IS instance.

Platforms

7705 SAR Gen 2

5.4 all-octets-offered-count

all-octets-offered-count

Syntax

[**no**] **all-octets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>queue>i-counters all-octets-offered-count)

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters all-octets-offered-count)

Full Context

configure log accounting-policy custom-record queue i-counters all-octets-offered-count

configure log accounting-policy custom-record ref-queue i-counters all-octets-offered-count

Description

This command includes all octets offered in the count.

The **no** form of this command excludes the octets offered in the count.

Default

no all-octets-offered-count

Platforms
7705 SAR Gen 2

5.5 all-packets-offered-count

all-packets-offered-count

Syntax
[**no**] **all-packets-offered-count**

Context
[Tree] (config>log>acct-policy>cr>queue>i-counters all-packets-offered-count)
[Tree] (config>log>acct-policy>cr>ref-queue>i-counters all-packets-offered-count)

Full Context
configure log accounting-policy custom-record queue i-counters all-packets-offered-count
configure log accounting-policy custom-record ref-queue i-counters all-packets-offered-count

Description
This command includes all packets offered in the count.
The **no** form of this command excludes the packets offered in the count.

Default
no all-packets-offered-count

Platforms
7705 SAR Gen 2

5.6 allocate-dual-sids

allocate-dual-sids

Syntax
[**no**] **allocate-dual-sids**

Context
[Tree] (config>router>ospf>segm-rtng>adj-sid allocate-dual-sids)
[Tree] (config>router>isis>segm-rtng>adj-sid allocate-dual-sids)

Full Context

configure router ospf segment-routing adjacency-sid allocate-dual-sids

configure router isis segment-routing adjacency-sid allocate-dual-sids

Description

This command enables the support of two SR-MPLS adjacency SIDs per interface. A protected and unprotected adjacency SID is instantiated and advertised. If an SR-MPLS adjacency SID already exists, an additional complementary (protected or unprotected) adjacency SID is created on the interface.

The **no** form of this command disables the support of two SR-MPLS adjacency SIDs per interface.

Default

no allocate-dual-sids

Platforms

7705 SAR Gen 2

5.7 allow-bgp-to-igp-export

allow-bgp-to-igp-export

Syntax

[no] allow-bgp-to-igp-export

Context

[Tree] (config>router allow-bgp-to-igp-export)

Full Context

configure router allow-bgp-to-igp-export

Description

This command enables the export of base BGP RTM routes into the IGP routing instance within the base router. This command applies to already exported BGP prefixes and to newly received BGP prefixes.

Default

allow-bgp-to-igp-export

Platforms

7705 SAR Gen 2

5.8 allow-boot-license-violations

allow-boot-license-violations

Syntax

[**no**] **allow-boot-license-violations**

Context

[Tree] (config>system allow-boot-license-violations)

Full Context

configure system allow-boot-license-violations

Description

This command configures whether the system should allow successful execution of the bootup configuration file when it contains license violations. When enabled, the system will not error on any configuration that causes a license violation and as a result permits the system to come into service. However, if violations are detected, the system reboots after a period of time if the violations are not fixed.

Platforms

7705 SAR Gen 2

5.9 allow-directed-broadcasts

allow-directed-broadcasts

Syntax

[**no**] **allow-directed-broadcasts**

Context

[Tree] (config>router>if allow-directed-broadcasts)

[Tree] (config>service>vprn>nw-if allow-directed-broadcasts)

[Tree] (config>service>vprn>if allow-directed-broadcasts)

[Tree] (config>service>ies>if allow-directed-broadcasts)

Full Context

configure router interface allow-directed-broadcasts

configure service vprn network-interface allow-directed-broadcasts

configure service vprn interface allow-directed-broadcasts

configure service ies interface allow-directed-broadcasts

Description

This command enables the forwarding of directed broadcasts out of the IP interface.

A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The **allow-directed-broadcasts** command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface.

When enabled, a frame destined to the local subnet on this IP interface is sent as a subnet broadcast out this interface. Care should be exercised when allowing directed broadcasts as it is a well-known mechanism used for denial-of-service attacks.

When disabled, directed broadcast packets discarded at this egress IP interface are counted in the normal discard counters for the egress SAP.

Note:
Allowing directed broadcasts is a well-known mechanism used for denial-of-service attacks.

By default, directed broadcasts are not allowed and are discarded at this egress IP interface.

The **no** form of this command disables the forwarding of directed broadcasts out of the IP interface. All broadcasts are dropped.

Default

no allow-directed-broadcasts — Directed broadcasts are dropped.

Platforms

7705 SAR Gen 2

5.10 allow-egress-remark-dscp

allow-egress-remark-dscp

Syntax

[**no**] **allow-egress-remark-dscp**

Context

[Tree] (config>oam-pm>session>ip allow-egress-remark-dscp)

Full Context

configure oam-pm session ip allow-egress-remark-dscp

Description

This command instructs the egress QoS process to modify the DSCP based on the egress QoS configuration. This command exposes the DSCP to egress DSCP processing rules.

The **no** form of this command instructs the egress QoS process to ignore the DSCP and allow it to bypass egress QoS. If the **config**>**qos**>**network**>**egress**>**remark force** command is configured for the network egress QoS profile, the egress QoS process is applied and the DSCP can be overwritten regardless of the **allow-egress-remark-dscp** configuration.

Platforms

7705 SAR Gen 2

5.11 allow-export-bgp-vpn

allow-export-bgp-vpn

Syntax

[**no**] **allow-export-bgp-vpn**

Context

[Tree] (config>service>vprn allow-export-bgp-vpn)

Full Context

configure service vprn allow-export-bgp-vpn

Description

This command allows routes leaked from another local VPRN service to be re-exported by this VPRN in the form of new VPN-IP routes. The service label, route targets, and BGP next-hop of the re-advertised routes are based on the configuration and default values of the re-exporting VPRN.

When re-exporting leaked routes, the following restrictions apply.

- The **allow-export-bgp-vpn** command is not configurable in combination with any of the following commands: **carrier-carrier-vpn** (CSC), **label-mode next-hop** (LPN), **type** {**hub** | **spoke** | **subscriber-split-horizon**}, **redundant-interface**, and **export-inactive-bgp**.
- Re-exported routes always have the per-VRF label of the exporting VPRN; label-per-prefix advertisement is not supported.
- The best-external (inactive BGP) routes leaked by another VPRN cannot be re-exported by a VPRN configured with **allow-export-bgp-vpn**.

Caution:
When a VPRN configured with **allow-export-bgp-vpn** advertises a leaked route, the **split-horizon** context is lost. A re-exported route can be easily advertised back to the sending peer unless this is blocked by BGP export policies. This can cause route flaps or other similar instability. In addition, **allow-export-bgp-vpn** may never be used in a VPRN service with a

route distinguisher that is used in other PEs attached to the same service; if the same route distinguisher is used in this case, there is constant route flap.

If the **no** form of this command is configured, leaked routes cannot be re-advertised as VPN-IP routes; they can only be re-advertised to PE-CE BGP peers of the VPRN.

Default

no allow-export-bgp-vpn

Platforms

7705 SAR Gen 2

5.12 allow-flex-algo-fallback

allow-flex-algo-fallback

Syntax

[**no**] **allow-flex-algo-fallback**

Context

[Tree] (config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family allow-flex-algo-fallback)

[Tree] (config>router>bgp>next-hop-resolution>shortcut-tunnel>family allow-flex-algo-fallback)

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel allow-flex-algo-fallback)

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel allow-flex-algo-fallback)

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel allow-flex-algo-fallback)

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel allow-flex-algo-fallback)

Full Context

configure router bgp next-hop-resolution labeled-routes transport-tunnel family allow-flex-algo-fallback

configure router bgp next-hop-resolution shortcut-tunnel family allow-flex-algo-fallback

configure service vpls bgp-evpn mpls auto-bind-tunnel allow-flex-algo-fallback

configure service vprn bgp-ipvpn mpls auto-bind-tunnel allow-flex-algo-fallback

configure service epipe bgp-evpn mpls auto-bind-tunnel allow-flex-algo-fallback

configure service vprn bgp-evpn mpls auto-bind-tunnel allow-flex-algo-fallback

Description

This command configures a router to relax the strictly enforced Flex-Algorithm aware autobind, which is enabled through an import policy configured with the **action flex-algo** command.

If the **allow-flex-algo-fallback** command is enabled, the BGP router can autobind to a fallback algorithm 0 tunnel if no target Flex-Algorithm tunnel is available. If the **allow-flex-algo-fallback** command is disabled, the BGP autobind is strictly enforced to an intended Flex-Algorithm tunnel, which may cause traffic loss if no corresponding Flex-Algorithm tunnel exists.

The **no** form of this command removes the **allow-flex-algo-fallback** command from the configuration.

Default

no allow-flex-algo-fallback

Platforms

7705 SAR Gen 2

allow-flex-algo-fallback

Syntax

allow-flex-algo-fallback

Context

[Tree] (config>service>vprn>auto-bind-tunnel allow-flex-algo-fallback)

Full Context

configure service vprn auto-bind-tunnel allow-flex-algo-fallback

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

5.13 allow-fragmentation

allow-fragmentation

Syntax

[**no**] **allow-fragmentation**

Context

[Tree] (config>service>pw-template allow-fragmentation)

[Tree] (config>service>sdp allow-fragmentation)

Full Context

configure service pw-template allow-fragmentation

configure service sdp allow-fragmentation

Description

This command disables the setting of the **do-not-fragment** bit in the IP header of GRE encapsulated service traffic. This feature is only applicable to GRE SDPs and will be applied to all service traffic using the associated GRE SDP.

The **no** form of this command removes the command from the active configuration and returns the associated SDP to its default which is to set the **do-not-fragment** bit in all GRE encapsulated service traffic.

Default

no allow-fragmentation

Platforms

7705 SAR Gen 2

5.14 allow-ftp

allow-ftp

Syntax

[**no**] **allow-ftp**

Context

[Tree] (config>service>vprn>management allow-ftp)

Full Context

configure service vprn management allow-ftp

Description

This commands allows access to the FTP server from VPRN.

The **no** form of this command removes FTP access for this VPRN.

Platforms

7705 SAR Gen 2

allow-ftp

Syntax

[**no**] **allow-ftp**

Context

[Tree] (config>system>security>management allow-ftp)

Full Context

configure system security management allow-ftp

Description

This command allows access to the FTP server from Base and Management routers if it is operationally up.

The **no** form of this command disallows access to the FTP server.

Default

allow-ftp

Platforms

7705 SAR Gen 2

5.15 allow-grpc

allow-grpc

Syntax

[**no**] **allow-grpc**

Context

[Tree] (config>system>security>management allow-grpc)

Full Context

configure system security management allow-grpc

Description

This command allows access to the gRPC server from Base and Management routers if it is operationally up.

The **no** form of this command disallows access to the gRPC server.

Platforms
7705 SAR Gen 2

allow-grpc

Syntax
[**no**] **allow-grpc**

Context
[Tree] (config>service>vprn>management allow-grpc)

Full Context
configure service vprn management allow-grpc

Description
This commands allows access to the GRPC server from VPRN.

The **no** form of this command removes GRPC access for this VPRN.

Platforms
7705 SAR Gen 2

5.16 allow-icmp-redirect

allow-icmp-redirect

Syntax
[**no**] **allow-icmp-redirect**

Context
[Tree] (config>router allow-icmp-redirect)

Full Context
configure router allow-icmp-redirect

Description
This command allows ICMP redirects received on the management interface.

The **no** form of this command drops the ICMP redirects received on the management interface.

Platforms
7705 SAR Gen 2

5.17 allow-icmp6-redirect

allow-icmp6-redirect

Syntax

[**no**] **allow-icmp-redirect**

Context

[Tree] (config>router allow-icmp6-redirect)

Full Context

configure router allow-icmp6-redirect

Description

This command allows IPv6 ICMP redirects received on the management interface.

The **no** form of this command drops the IPv6 ICMP redirects received on the management interface.

Platforms

7705 SAR Gen 2

5.18 allow-immediate

allow-immediate

Syntax

[**no**] **allow-immediate**

Context

[Tree] (config>system>management-interface>cli>classic-cli allow-immediate)

Full Context

configure system management-interface cli classic-cli allow-immediate

Description

This command enables write access in the classic CLI configuration branch without having to use the classic CLI **candidate edit** functionality.

The **no** form of this command blocks write access and configuration changes in the classic CLI configuration branch, and the classic CLI configuration branch is read-only. This enforces using the classic

CLI **candidate edit** functionality, including **candidate commit**, to modify the router configuration, instead of allowing immediate line-by-line configuration changes.

Default

allow-immediate

Platforms

7705 SAR Gen 2

5.19 allow-ip-int-bind

allow-ip-int-bind

Syntax

[**no**] **allow-ip-int-bind**

Context

[Tree] (config>service>vpls allow-ip-int-bind)

Full Context

configure service vpls allow-ip-int-bind

Description

The allow-ip-int-bind command that sets a flag on the VPLS or I-VPLS service that enables the ability to attach an IES or VPRN IP interface to the VPLS service in order to make the VPLS service routable. When the allow-ip-int-bind command is not enabled, the VPLS service cannot be attached to an IP interface.

VPLS Configuration Constraints for Enabling allow-ip-int-bind

When attempting to set the allow-ip-int-bind VPLS flag, the system first checks to see if the correct configuration constraints exist for the VPLS service and the network ports. The following VPLS features must be disabled or not configured for the allow-ip-int-bind flag to set:

- SAP ingress QoS policies applied to the VPLS SAPs cannot have MAC match criteria defined
- The VPLS service type cannot be B-VPLS or M-VPLS
- MVR from Routed VPLS and to another SAP is not supported
- Enhanced and Basic Subscriber Management (ESM and BSM) features
- Network domain on SDP bindings

Once the VPLS allow-ip-int-bind flag is set on a VPLS service, the above features cannot be enabled on the VPLS service.

Network Port Hardware Constraints

The system also checks to ensure that all ports configured in network mode are associated with FlexPath2 forwarding planes. If a port is currently in network mode and the port is associated with a FlexPath1

forwarding plane, the allow-ip-int-bind command will fail. Once the allow-ip-int-bind flag is set on any VPLS service, attempting to enable network mode on a port associated with a FlexPath1 forwarding plane will fail.

VPLS SAP Hardware Constraints

Besides VPLS configuration and network port hardware association, the system also checks to that all SAPs within the VPLS are created on Ethernet ports and the ports are associated with FlexPath2 forwarding planes. Certain Ethernet ports and virtual Ethernet ports are not supported which include CCAG virtual ports (VSM based). If a SAP in the VPLS exists on an unsupported port type or is associated with a FlexPath1 forwarding plane, the allow-ip-int-bind command will fail. Once the allow-ip-int-bind flag is set on the VPLS service, attempting to create a VPLS SAP on the wrong port type or associated with a FlexPath1 forwarding plane will fail.

VPLS Service Name Bound to IP Interface without allow-ip-int-bind flag Set

If a service name is applied to a VPLS service and that service name is also bound to an IP interface but the allow-ip-int-bind flag has not been set on the VPLS service context, the system attempt to resolve the service name between the VPLS service and the IP interface will fail. After the allow-ip-int-bind flag is successfully set on the VPLS service, either the service name on the VPLS service must be removed and reapplied or the IP interface must be re-initialized using the **shutdown** / **no shutdown** commands. This will cause the system to reattempt the name resolution process between the IP interface and the VPLS service.

The **no** form of this command resets the allow-ip-int-bind flag on the VPLS service. If the VPLS service currently has an IP interface from an IES or VPRN service attached, the no allow-ip-int-bind command will fail. Once the allow-ip-int-bind flag is reset on the VPLS service, the configuration and hardware restrictions associated with setting the flag are removed. The port network mode hardware restrictions are also removed.

Platforms

7705 SAR Gen 2

5.20 allow-ipv6-udp-checksum-zero

allow-ipv6-udp-checksum-zero

Syntax

[**no**] **allow-ipv6-udp-checksum-zero**

Context

[Tree] (config>router>twamp-light>reflector allow-ipv6-udp-checksum-zero)

[Tree] (config>service>vprn>twamp-light>reflector allow-ipv6-udp-checksum-zero)

Full Context

configure router twamp-light reflector allow-ipv6-udp-checksum-zero

configure service vprn twamp-light reflector allow-ipv6-udp-checksum-zero

Description

This command configures the acceptance of IPv6 packets with UDP checksums of 0.This optional configuration allows the router to process arriving IPv6 TWAMP Test packets that contain IPv6 UDP checksum of 0x0000. The UDP port specific to this TWAMP Light test bypasses the default discard IPv6 UDP checksum 0x0000. If this optional command is not configured, IPv6 UDP checksum 0x000 arriving packets are discarded.

The **no** form of this command reverts to the default value, discarding packets that arrive with an IPv6 UDP checksum of 0x0000.

Default

no allow-ipv6-udp-checksum-zero

Platforms

7705 SAR Gen 2

5.21 allow-lease-query

allow-lease-query

Syntax

[**no**] **allow-lease-query**

Context

[Tree] (config>service>vprn>dhcp6>server allow-lease-query)

[Tree] (config>router>dhcp6>server allow-lease-query)

Full Context

configure service vprn dhcp6 local-dhcp-server allow-lease-query

configure router dhcp6 local-dhcp-server allow-lease-query

Description

If enabled, the local DHCPv6 server will handle and reply to lease query messages.

The **no** form of this command disables lease query support.

Platforms

7705 SAR Gen 2

5.22 allow-local-management

allow-local-management

Syntax

[**no**] **allow-local-management**

Context

[Tree] (config>service>vprn>grt>enable-grt allow-local-management)

Full Context

configure service vprn grt-lookup enable-grt allow-local-management

Description

This command enables the support of specific management protocols over VPRN interfaces that terminate on Base routing context IPv4 and IPv6 interface addresses, including Base loopback and system addresses. Global Routing Table (GRT) leaking is used to enable the visibility and access of the Base interface addresses in the VPRN. The supported protocols are Telnet, FTP, SNMP, TACACS+, RADIUS (IPv4 only, not IPv6), SSH (including applications that ride over the standard SSH TCP port 22 such as SCP and SFTP) and NETCONF (configured on port 22 or 830).

Ping and traceroute responses from the Base router interfaces are supported but are not configurable.

The **allow-local-management** command does not control the support for management protocols terminating on VPRN interfaces directly. See "Node Management using VPRN" in the *7705 SAR Gen 2 Layer 3 Services Guide: IES and VPRN* for more information. Also, see the **access** command in the **config>service>vprn>snmp** context, and the commands in the **config>service>vprn>management** context.

Platforms

7705 SAR Gen 2

5.23 allow-netconf

allow-netconf

Syntax

[**no**] **allow-netconf**

Context

[Tree] (config>system>security>management allow-netconf)

Full Context

configure system security management allow-netconf

Description

This command allows access to the NETCONF server from Base and Management routers if it is operationally up.

The **no** form of this command disallows access to the NETCONF server.

Platforms

7705 SAR Gen 2

allow-netconf

Syntax

[**no**] **allow-netconf**

Context

[Tree] (config>service>vprn>management allow-netconf)

Full Context

configure service vprn management allow-netconf

Description

This commands allows access to the NETCONF server from VPRN.

The **no** form of this command removes NETCONF access for this VPRN.

Platforms

7705 SAR Gen 2

5.24 allow-reverse-route-override

allow-reverse-route-override

Syntax

allow-reverse-route-override [*type*]

no allow-reverse-route-override

Context

[Tree] (config>service>vprn>ipsec allow-reverse-route-override)

Full Context

configure service vprn ipsec allow-reverse-route-override

Description

This command allows a new dynamic LAN-to-LAN tunnel that terminates in the private VPRN service to be created with an overlapping reverse route.

The **no** form of this command reverts to the default value.

Default

no allow-reverse-route-override

Parameters

type

Specifies the action to take when the system accepts a new reverse route.

Values same-idi — Specifies that the system accepts a new reverse route and removes the existing route only if the IDi of the new tunnel is the same as existing route.

any-idi — Specifies that the system accepts a new reverse route and removes the existing route regardless of the IDi.

Platforms

7705 SAR Gen 2

5.25 allow-sr-over-srte

allow-sr-over-srte

Syntax

[**no**] **allow-sr-over-srte**

Context

[Tree] (config>router>ospf>igp-sc allow-sr-over-srte)

[Tree] (config>router>isis>igp-sc allow-sr-over-srte)

Full Context

configure router ospf igp-shortcut allow-sr-over-srte

configure router isis igp-shortcut allow-sr-over-srte

Description

This command enables the SR-TE LSPs as eligible SRv4 or SRv6 IGP shortcuts.

For SR-MPLS SRv4 and SRv6, IGP shortcuts can only use SR-TE LSPs with **allow-sr-over-srte** explicitly enabled that have an adjacency SID as top SID in the SR-TE LSP. IPv4 and IPv6 addresses can use all available SR-TE LSPs as shortcuts regardless of the explicit **allow-sr-over-srte** configuration.

Under ECMP, when IGP **allow-sr-over-srte** is configured, preference is given to the SR-TE LSPs with **allow-sr-over-srte** explicitly configured over the LSPs that do not have **allow-sr-over-srte** configured.

The **no** form of this command disables the eligibility.

Default

no allow-sr-over-srte

Platforms

7705 SAR Gen 2

5.26 allow-ssh

allow-ssh

Syntax

[**no**] **allow-ssh**

Context

[Tree] (config>service>vprn>management allow-ssh)

Full Context

configure service vprn management allow-ssh

Description

This command allows configuration of the SSH parameters.

The **no** form of this command disallows configuration of the SSH parameters.

Platforms

7705 SAR Gen 2

allow-ssh

Syntax

[**no**] **allow-ssh**

Context

[Tree] (config>system>security>management allow-ssh)

Full Context

configure system security management allow-ssh

Description

This command allows the SSH parameters to be configured from Base and Management routers.

The **no** form of this command disallows SSH parameters from being configured.

Default

allow-ssh

Platforms

7705 SAR Gen 2

5.27 allow-static

allow-static

Syntax

allow-static

no allow-static

Context

[Tree] (config>router>bgp>next-hop-res>labeled-routes allow-static)

Full Context

configure router bgp next-hop-resolution labeled-routes allow-static

Description

This command allows the BGP next-hop of label-IPv4, label-IPv6, VPN-IPv4, and VPN-IPv6 routes received from any EBGP or IBGP peer to be resolved using static routes, except for static default routes (0/0 and ::/0).

A static route is less preferred than a local or interface route for resolving the BGP next-hop of labeled route, but more preferred than other IGP routes or tunnels.

Note:
A label-IPv4 or label-IPv6 route can be resolved by a static blackhole route, even when the **allow-static** command is not configured, but only if the static blackhole route is the longest prefix match (LPM) static route for the BGP next-hop address.

Default

no allow-static

Platforms

7705 SAR Gen 2

5.28 allow-telnet

allow-telnet

Syntax

[**no**] **allow-telnet**

Context

[Tree] (config>service>vprn>management allow-telnet)

Full Context

configure service vprn management allow-telnet

Description

This command allows access to the Telnet server from a VPRN.

The **no** form of this command removes the Telnet access.

Platforms

7705 SAR Gen 2

allow-telnet

Syntax

[**no**] **allow-telnet**

Context

[Tree] (config>system>security>management allow-telnet)

Full Context

configure system security management allow-telnet

Description

This command allows access to the Telnet server from Base and Management routers if it is operationally up.

The **no** form of this command disallows access to the Telnet server.

Default
allow-telnet

Platforms
7705 SAR Gen 2

5.29 allow-telnet6

allow-telnet6

Syntax
[**no**] **allow-telnet6**

Context
[Tree] (config>service>vprn>management allow-telnet6)

Full Context
configure service vprn management allow-telnet6

Description
This command allows access to the Telnet IPv6 server from a VPRN.

The **no** form of this command removes the Telnet IPv6 access.

Platforms
7705 SAR Gen 2

allow-telnet6

Syntax
[**no**] **allow-telnet6**

Context
[Tree] (config>system>security>management allow-telnet6)

Full Context
configure system security management allow-telnet6

Description
This command allows access to the Telnet IPv6 server from Base and Management routers if it is operationally up.

The **no** form of this command disallows access to the Telnet IPv6 server.

Default

allow-telnet6

Platforms

7705 SAR Gen 2

5.30 allow-unresolved-leaking

allow-unresolved-leaking

Syntax

[**no**] **allow-unresolved-leaking**

Context

[Tree] (config>router>bgp>next-hop-res allow-unresolved-leaking)

Full Context

configure router bgp next-hop-resolution allow-unresolved-leaking

Description

This command instructs BGP, in the base router instance, to allow its routes to be leaked to other (VPRN) BGP instances, even if the routes to be leaked do not have a BGP next hop that can be resolved by the base instance.

By default, BGP routes cannot be leaked to another BGP instance unless they are resolvable by the instance that receives them.

The **no** form of this command provides the default behavior.

Default

no allow-unresolved-leaking

Platforms

7705 SAR Gen 2

5.31 allow-unsecure-connection

allow-unsecure-connection

Syntax

[**no**] **allow-unsecure-connection**

Context

[Tree] (config>system>grpc allow-unsecure-connection)

Full Context

configure system grpc allow-unsecure-connection

Description

This command enables unsecure operation of gRPC connections. This means that TCP connections are not encrypted, including username and password information.

This command can be enabled only if there is no TLS profile assigned to the gRPC server.

The **no** form of this command enables TLS encryption on gRPC connections.

Default

no allow-unsecure-connection

Platforms

7705 SAR Gen 2

allow-unsecure-connection

Syntax

[**no**] **allow-unsecure-connection**

Context

[Tree] (config>system>management-interface>remote-management allow-unsecure-connection)

Full Context

configure system management-interface remote-management allow-unsecure-connection

Description

This command enables unsecure operation of all remote manager connections. In an unsecured operation, connections are not encrypted, including the username and password information.

This command and **client-tls-profile** are mutually exclusive. This means it can be used only if there are no TLS profiles assigned to the server.

If this command is also configured in the **config**>**system**>**management-interface**>**remote-management**> **manager** context, that configuration takes precedence.

The **no** form of this command disables unsecured connections.

Default

no allow-unsecure-connection

Platforms

7705 SAR Gen 2

allow-unsecure-connection

Syntax

[**no**] **allow-unsecure-connection**

Context

[Tree] (config>system>management-interface>remote-management>manager allow-unsecure-connection)

Full Context

configure system management-interface remote-management manager allow-unsecure-connection

Description

This command allows an unsecured connection to the remote managers; the TCP connection is not encrypted. This includes username and password information.

This command and **client-tls-profile** are mutually exclusive.

This command takes precedence over the same command configured in the **config**> **system**>**management-interface**>**remote-management** context, if applicable.

The **no** form of this command disables unsecured connections for the specified manager.

Default

no allow-unsecure-connection

Platforms

7705 SAR Gen 2

allow-unsecure-connection

Syntax

[**no**] **allow-unsecure-connection**

Context

[Tree] (config>system>telemetry>destination-group allow-unsecure-connection)

Full Context

configure system telemetry destination-group allow-unsecure-connection

Description

This command enables an unsecured connection for a specified destination group.

This command is mutually exclusive with the **tls-client-profile** command.

The **no** form of this command disables unsecured connections for the specified destination group.

Default

no allow-unsecure-connection

Platforms

7705 SAR Gen 2

allow-unsecure-connection

Syntax

[**no**] **allow-unsecure-connection**

Context

[Tree] (config>system>grpc-tunnel>destination-group allow-unsecure-connection)

Full Context

configure system grpc-tunnel destination-group allow-unsecure-connection

Description

This command enables an unsecured connection for a specified destination group, which allows a gRPC tunnel to run without a secured transport protocol. Data is transferred in unencrypted form.

This command is mutually exclusive with the **tls-client-profile** command.

The **no** form of this command disables unsecured connections for the specified destination group.

Default

no allow-unsecure-connection

Platforms

7705 SAR Gen 2

5.32 allow-unsecured-msgs

allow-unsecured-msgs

Syntax

[**no**] **allow-unsecured-msgs**

Context

[Tree] (config>service>ies>if>ipv6>secure-nd allow-unsecured-msgs)

Full Context

configure service ies interface ipv6 secure-nd allow-unsecured-msgs

Description

This command specifies whether unsecured messages are accepted. When Secure Neighbor Discovery (SeND) is enabled, only secure messages are accepted by default.

The **no** form of this command disables accepting unsecured messages.

Platforms

7705 SAR Gen 2

allow-unsecured-msgs

Syntax

[**no**] **allow-unsecured-msgs**

Context

[Tree] (config>service>vprn>if>send allow-unsecured-msgs)

Full Context

configure service vprn interface ipv6 secure-nd allow-unsecured-msgs

Description

This command specifies whether unsecured messages are accepted. When Secure Neighbor Discovery (SeND) is enabled, only secure messages are accepted by default.

The **no** form of this command disables accepting unsecured messages.

Platforms

7705 SAR Gen 2

allow-unsecured-msgs

Syntax

[**no**] **allow-unsecured-msgs**

Context

[Tree] (config>router>if>ipv6>secure-nd allow-unsecured-msgs)

Full Context

configure router interface ipv6 secure-nd allow-unsecured-msgs

Description

This command specifies whether unsecured messages are accepted. When Secure Neighbor Discovery (SeND) is enabled, only secure messages are accepted by default.

The **no** form of this command disables accepting unsecured messages.

Platforms

7705 SAR Gen 2

5.33 allow-user-name

allow-user-name

Syntax

[**no**] **allow-user-name**

Context

[Tree] (config>system>security>password>complexity-rules allow-user-name)

Full Context

configure system security password complexity-rules allow-user-name

Description

The user name is allowed to be used as part of the password.

The **no** form of this command does not allow user name to be used as password.

Default

no allow-user-name

Platforms

7705 SAR Gen 2

5.34 allowed-peer-as

allowed-peer-as

Syntax

[**no**] **allowed-peer-as** *min-as-number* [**max** *max-as-number*]

Context

[Tree] (config>service>vprn>bgp>group>dynamic-neighbor>match>prefix allowed-peer-as)

Full Context

configure service vprn bgp group dynamic-neighbor match prefix allowed-peer-as

Description

This command configures a single peer AS value or a contiguous range of peer AS values to associate with a prefix from which dynamic BGP sessions can be accepted.

If an incoming dynamic BGP session is associated with the prefix then the peer's AS, as reported in the OPEN message, is checked against the list of allowed-peer-as values. If the peer AS is not contained in one of the **allowed-peer-as** commands, then the connection is rejected with a Bad_Peer_AS error. If there is no **allowed-peer-as** configuration in the matched prefix, then the ASN in the peer's OPEN message, is checked against the group level peer-as.

The **no** form of this command removes an allowed-peer-as entry.

Default

no allowed-peer-as

Parameters

min-as-number

Specifies an allowed peer AS value as well as the start of an allowed range if the *max-as-number* value is also configured.

Values 1 to 4294967295

max-as-number

Specifies the end of an allowed range.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

allowed-peer-as

Syntax

[**no**] **allowed-peer-as** *min-as-number* [**max** *max-as-number*]

Context

[Tree] (config>router>bgp>group>dynamic-neighbor>match>prefix allowed-peer-as)

Full Context

configure router bgp group dynamic-neighbor match prefix allowed-peer-as

Description

This command configures a single peer AS value or a contiguous range of peer AS values to associate with a prefix from which dynamic BGP sessions can be accepted.

If an incoming dynamic BGP session is associated with the prefix, then the peer's AS, as reported in the OPEN message, is checked against the list of allowed-peer-as values. If the peer AS is not contained in one of the **allowed-peer-as** commands, then the connection is rejected with a Bad_Peer_AS error. If there is no **allowed-peer-as** configuration in the matched prefix, then the ASN in the peer's OPEN message, is checked against the group level peer-as.

The **no** form of this command removes an allowed-peer-as entry.

Default

no allowed-peer-as

Parameters

min-as-number

Specifies an allowed peer AS value as well as the start of an allowed range if the *max-as-number* value is also configured.

Values 1 to 4294967295

max-as-number

Specifies the end of an allowed range.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

allowed-peer-as

Syntax

[**no**] **allowed-peer-as** *min-as-number* [**max** *max-as-number*]

Context

[Tree] (config>router>bgp>group>dynamic-neighbor>interface allowed-peer-as)

[Tree] (config>service>vprn>bgp>group>dynamic-neighbor>interface allowed-peer-as)

Full Context

configure router bgp group dynamic-neighbor interface allowed-peer-as

configure service vprn bgp group dynamic-neighbor interface allowed-peer-as

Description

This command configures a singular allowed peer AS value or a range of acceptable values.

The **no** form of this command removes an allowed peer AS value or range of acceptable values.

Parameters

min-as-number

Specifies an allowed peer AS value as well as the start of an allowed range if the *max-as-number* value is also configured.

Values 1 to 4294967295

max-as-number

Specifies the end of an allowed range.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

5.35 allowed-source-macs

allowed-source-macs

Syntax

allowed-source-macs

Context

[Tree] (config>port>ethernet>dot1x>per-host-authentication allowed-source-macs)

Full Context

configure port ethernet dot1x per-host-authentication allowed-source-macs

Description

Commands in this context add the source MAC addresses of the hosts to the allowed MAC list.

Platforms

7705 SAR Gen 2

5.36 always-compare-med

always-compare-med

Syntax

always-compare-med {**zero** | **infinity**}

no always-compare-med strict-as {**zero** | **infinity**}

no always-compare-med

Context

[Tree] (config>router>bgp>best-path-selection always-compare-med)

[Tree] (config>service>vprn>bgp>path-selection always-compare-med)

Full Context

configure router bgp best-path-selection always-compare-med

configure service vprn bgp best-path-selection always-compare-med

Description

This command configures the comparison of BGP routes based on the MED attribute. The default behavior of SR OS (equivalent to the **no** form of this command) is to only compare two routes on the basis of MED if they have the same neighbor AS (the first non-confed AS in the received AS_PATH attribute). Also by default, a route without a MED attribute is handled the same as though it had a MED attribute with the value 0. The **always-compare-med** command without the **strict-as** keyword allows MED to be compared even if the paths have a different neighbor AS; in this case, if neither **zero** nor **infinity** is specified, the **zero** option is inferred, meaning a route without a MED is handled the same as though it had a MED attribute with the value 0. When the **strict-as** keyword is present, MED is only compared between paths from the same neighbor AS, and in this case, **zero** or **infinity** is mandatory and tells BGP how to interpret paths without a MED attribute.

Default

no always-compare-med

Parameters

zero

Specifies that for routes learned without a MED attribute that a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred.

infinity

Specifies for routes learned without a MED attribute that a value of infinity (2^32-1) is used in the MED comparison. This in effect makes these routes the least desirable.

strict-as

Specifies that the BGP MED values are only compared if the route comes from the same neighbor AS.

Platforms

7705 SAR Gen 2

5.37 always-display

always-display

Syntax

always-display

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>info-output always-display)

Full Context

configure system management-interface cli md-cli environment info-output always-display

Description

Commands in this context configure the elements that are always displayed in the **info** output of an MD-CLI session, regardless of whether the **detail** option is used.

Platforms

7705 SAR Gen 2

5.38 always-set-sender-for-ir

always-set-sender-for-ir

Syntax

[**no**] **always-set-sender-for-ir**

Context

[Tree] (config>system>security>pki>ca-profile>cmpv2 always-set-sender-for-ir)

Full Context

configure system security pki ca-profile cmpv2 always-set-sender-for-ir

Description

This command specifies to always set the sender field in CMPv2 header of all Initial Registration (IR) messages with the subject name. By default, the sender field is only set if an optional certificate is specified in the CMPv2 request.

Default

no always-set-sender-for-ir

Platforms

7705 SAR Gen 2

5.39 ancp

ancp

Syntax

ancp

Context

[Tree] (config>system>persistence ancp)

Full Context

configure system persistence ancp

Description

This command configures ANCP persistence parameters.

Platforms

7705 SAR Gen 2

5.40 anycast

anycast

Syntax

[**no**] **anycast** *rp-ip-address*

Context

[Tree] (config>service>vprn>pim>rp anycast)

Full Context

configure service vprn pim rp anycast

Description

This command configures a PIM anycast protocol instance for the RP being configured. Anycast enables fast convergence when a PIM RP router fails by allowing receivers and sources to rendezvous at the closest RP.

The **no** form of this command removes the anycast instance from the configuration.

Parameters

rp-ip-address

Configure the loopback IP address shared by all routes that form the RP set for this anycast instance. Only a single address can be configured. If another anycast command is entered with an address then the old address will be replaced with the new address. If no ip-address is entered then the command is simply used to enter the anycast CLI level.

Values Any valid loopback address configured on the node.

Platforms

7705 SAR Gen 2

anycast

Syntax

anycast *ipv6-address*

no anycast *ipv6-address*

Context

[Tree] (config>service>vprn>pim>rp>ipv6 anycast)

Full Context

configure service vprn pim rp ipv6 anycast

Description

This command configures an IPv6 PIM anycast protocol instance for the RP being configured. Anycast enables fast convergence when a PIM RP router fails by allowing receivers and sources to rendezvous at the closest RP.

The **no** form of this command removes the anycast instance from the configuration.

Parameters

ipv6-address

Configures the loopback IP address shared by all routes that form the RP set for this anycast instance. Only a single address can be configured. If another anycast command is

entered with an address then the old address will be replaced with the new address. If no address is entered then the command is simply used to enter the anycast CLI context.

Values

ipv6-address	: x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x [0 to FFFF]H
	d [0 to 255]D

Platforms

7705 SAR Gen 2

anycast

Syntax

[**no**] **anycast** *rp-ip-address*

Context

[Tree] (config>router>pim>rp anycast)

Full Context

configure router pim rp anycast

Description

This command configures a PIM anycast protocol instance for the RP being configured. Anycast enables fast convergence when a PIM RP router fails by allowing receivers and sources to rendezvous at the closest RP.

The **no** form of this command removes the anycast instance from the configuration.

Parameters

rp-ip-address

Specifies the loopback IP address shared by all routes that form the RP set for this anycast instance. Only a single address can be configured. If another anycast command is entered with an address then the old address will be replaced with the new address. If no ip-address is entered then the command is simply used to enter the anycast CLI level.

Values Any valid loopback address configured on the node.

Platforms

7705 SAR Gen 2

anycast

Syntax

[**no**] **anycast** *ipv6-address*

Context

[Tree] (config>router>pim>rp>ipv6 anycast)

Full Context

configure router pim rp ipv6 anycast

Description

This command configures a PIM anycast protocol instance for the RP being configured. Anycast enables fast convergence when a PIM RP router fails by allowing receivers and sources to rendezvous at the closest RP.

The **no** form of this command removes the anycast instance from the configuration.

Parameters

ipv6-address

Specifies the loopback IPv6 address shared by all routes that form the RP set for this anycast instance. Only a single address can be configured. If another anycast command is entered with an address then the old address is replaced with the new address. If no *ipv6-address* is entered then the command is simply used to enter the anycast CLI level.

Values Any valid loopback address configured on the node.

Platforms

7705 SAR Gen 2

5.41 app-route-notifications

app-route-notifications

Syntax

app-route-notifications

Context

[Tree] (config>log app-route-notifications)

Full Context

configure log app-route-notifications

Description

Specific system applications in SR OS can take action based on a route to certain IP destinations being available. This CLI branch contains configuration related to these route availability notifications. A delay can be configured between the time that a route is determined as available in the CPM, and the time that the application is notified of the available route. For example, this delay may be used to increase the chances that other system modules (such as IOMs/XCMs/MDAs/XMAs) are fully programmed with the new route before the application takes action. Currently, the only application that acts upon these *route available* or *route changed* notifications with their configurable delays is the SNMP replay feature, which receives notifications of route availability to the SNMP trap receiver destination IP address.

Platforms

7705 SAR Gen 2

5.42 application

application

Syntax

application *dscp-app-name* **dscp** {*dscp-value* | *dscp-name*}

application *dot1p-app-name* **dot1p** *dot1p-priority*

no application {*dscp-app-name* | *dot1p-app-name*}

Context

[Tree] (config>service>vprn>sgt-qos application)

[Tree] (config>router>sgt-qos application)

Full Context

configure service vprn sgt-qos application

configure router sgt-qos application

Description

This command configures DSCP/dot1p remarking for self-generated application traffic. When an application is configured using this command, the specified DSCP name is used for all packets generated by this application within the router instance it is configured. The instances can be base router, vprn, or management.

Using the value configured in this command:

- sets the DSCP bits in the IP packet
- maps to the FC. This value will be signaled from the CPM to the egress forwarding complex.
- based on this signaled FC, the egress forwarding complex QoS policy sets the Ethernet 802.1p and MPLS EXP bits. This includes ARP, PPPoE, and IS-IS packets that do not carry DSCP bits.
- configure the DSCP value in the egress IP header. The egress QoS policy does not overwrite this value.

Only one DSCP name can be configured per application, if multiple entries are configured, the subsequent entry overrides the previous configured entry.

The **no** form of this command reverts back to the default value.

Parameters

dscp-app-name

Specifies the DSCP application name.

Values Some of the following values may only apply to specific products. Refer to the *SR OS R25.x.Rx Software Release Notes* for details about application support for different SR OS products:

bfd, bgp, bmp, call-trace, cflowd, dhcp, diameter, dns, ftp, grpc, gtp, http, icmp, igmp, igmp-reporter, l2tp, ldp, mld, mpls-udp-return, msdp, mtrace2, ndis, ntp, ospf, pcep, pim, ptp, radius, rip, rsvp, sflow, snmp, snmp-notification, srrp, ssh, syslog, tacplus, telnet, tftp, traceroute, vrrp

dscp-value

Specifies a value when this packet egresses; the respective egress policy should provide the mapping for the DSCP value to either LSP-EXP bits or IEEE 802.1p (dot1p) bits as appropriate. Otherwise, the default mapping applies.

Values 0 to 63

dscp-name

Specifies the DSCP name.

Values none, be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

dot1p-priority

Specifies the dot1p priority.

Values none, 0 to 7

dot1p-app-name

Specifies the dot1p application name.

Values Some of the following values may only apply to specific products. Refer to the *SR OS R25.x.Rx Software Release Notes* for details about application support for different SR OS products:

arp, isis, pppoe

Platforms

7705 SAR Gen 2

application

Syntax

application *app* [*ip-int-name* | *ip-address*]

no application *app*

Context

[Tree] (config>service>vprn>source-address application)

Full Context

configure service vprn source-address application

Description

This command specifies the source address and application name.

The **no** form of this command removes the interface name or IP address from the command.

Parameters

app

Specifies the application name.

Values cflowd, ntp, ping, ptp, snmptrap, ssh, telnet, traceroute, icmp-error

ip-int-name

Specifies the name of the IP interface, up to 32 characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

ip-address

Specifies the source IP address.

Values

ipv4-address: a.b.c.d

Platforms

7705 SAR Gen 2

application

Syntax

application {**eq** | **neq**} *application-id*

no application

Context

[Tree] (config>service>vprn>log>filter>entry>match application)

Full Context

configure service vprn log filter entry match application

Description

This command adds an OS application as an event filter match criterion.

An OS application is the software entity that reports the event. Applications include IP, MPLS, OSPF, CLI, SERVICES and so on Only one application can be specified. The latest **application** command overwrites the previous command.

The **no** form of this command removes the application as a match criterion.

Default

no application — no application match criterion is specified

Parameters

eq | neq

The operator specifying the type of match.

Values

eq	equal to
neq	not equal to

application-id

The application name string.

Values port, ppp, rip, route, policy, rsvp, security, snmp, stp, svcmgr, system, user, vrrp, vrtr

Platforms

7705 SAR Gen 2

application

Syntax

application {**eq** | **neq**} *application-id*

no application

Context

[Tree] (config>log>filter>entry>match application)

Full Context

configure log filter entry match application

Description

This command adds an OS application as an event filter match criterion.

An OS application is the software entity that reports the event. Applications include IP, MPLS, OSPF, CLI, SERVICES and so on. Only one application can be specified. The latest **application** command overwrites the previous command.

The **no** form of this command removes the application as a match criterion.

Parameters

eq | neq

Specifies the operator match type. Valid operators are listed in Table 16: Valid Operators.

Table 16: Valid Operators

Operator	Notes
eq	equal to
neq	not equal to

application-id

The application name string.

Values application_assurance, aps, bgp, cflowd, chassis, debug, dhcp, dhcps, diameter, dynsvc, efm_oam, elmi, ering, eth_cfm, etun, fiter, gsmp, igh, igmp, igmp_snooping, ip, ipsec, isis, l2tp, lag, ldp, li, lldp, logger, mcpath, mc_redundancy, mirror, mld, mld_snooping, mpls, mpls_tp, msdp, nat, ntp, oam, open_flow, ospf, pim, pim_snooping, port, ppp, pppoe, ptp, radius, rip, rip_ng, route_policy, rsvp, security, snmp, stp, svcmgr, system, user, video, vrrp, vrtr, wlan_gw, wpp

Platforms

7705 SAR Gen 2

application

Syntax

application *app* [*ip-int-name* | *ip-address*]

no application *app*

Context

[Tree] (config>system>security>source-address application)

Full Context

configure system security source-address application

Description

This command configures the source IP address specified by the **source-address** command.

The **no** form of this command removes the interface name or IP address from the command.

Parameters

app

Specifies the application name.

Values cflowd, dns, ftp, ntp, ldap, ping, ptp, radius, sflow, snmptrap, sntp, ssh, syslog, tacplus, telnet, traceroute, mcreporter, icmp-error

ip-int-name

Specifies the name of the IP interface, up to 32 characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

ip-address

Specifies the source IP address.

Values *ipv4-address*: a.b.c.d

Platforms

7705 SAR Gen 2

application

Syntax

application *application* **[keychain** *keychain-name***]**

no application *application*

Context

[Tree] (config>redundancy>multi-chassis>peer>sync>transport-encryption application)

Full Context

configure redundancy multi-chassis peer sync transport-encryption application

Description

This command configures transport encryption.

The **no** form of this command removes the specified application.

Parameters

application

Specifies a Multi-Chassis Synchronization (MCS) client application

keychain-name

Specifies a keychain name, up to 32 characters

Platforms

7705 SAR Gen 2

5.43 application-link-attributes

application-link-attributes

Syntax

[**no**] **application-link-attributes**

Context

[Tree] (config>router>isis>traffic-engineering-options application-link-attributes)

Full Context

configure router isis traffic-engineering-options application-link-attributes

Description

Commands in this context configure the advertisement of the TE attributes of each link on a per-application basis. Two applications are supported in SR OS: RSVP-TE and SR-TE.

The legacy mode of advertising TE attributes that is used in RSVP-TE is still supported but it can be disabled by using the **no legacy** command, which also enables per-application TE attribute advertisement for RSVP-TE.

The **no** form of this command deletes the context.

Default

no application-link-attributes

Platforms

7705 SAR Gen 2

5.44 application6

application6

Syntax

application6 *app ipv6-address*

no application6 *app*

Context

[Tree] (config>service>vprn>source-address application6)

Full Context

configure service vprn source-address application6

Description

This command specifies the IPv6 source address and application.

The **no** form of this command removes the application and IPv6 address from the configuration.

Parameters

app

Specifies the application name.

Values cflowd, ntp, ping, ptp, snmptrap, ssh, telnet, traceroute, icmp6-error

ipv6-address

Specifies the IPv6 address.

Values

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

Platforms

7705 SAR Gen 2

application6

Syntax

application6 *app ipv6-address*

no application6

Context

[Tree] (config>system>security>source-address application6)

Full Context

configure system security source-address application6

Description

This command configures the application to use the source IPv6 address specified by the **source-address** command.

The **no** form of this command removes the application and IPv6 address from the configuration.

Parameters

app

Specifies the application name.

Values cflowd, dns, ftp, ldap, ntp, ping, ptp, radius, sflow, snmptrap, ssh, syslog, tacplus, telnet, traceroute, icmp6-error

ipv6-address

Specifies the IPv6 address.

Values

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

Platforms

7705 SAR Gen 2

5.45 apply-bgp-nh-override

apply-bgp-nh-override

Syntax

[**no**] **apply-bgp-nh-override**

Context

[Tree] (config>service>vprn>pim apply-bgp-nh-override)

Full Context

configure service vprn pim apply-bgp-nh-override

Description

This command forces the RPF check to be performed via IPv4 VPN AF next-hop and not via IPv4 VPN AF VRF import extended community.

Default

no apply-bgp-nh-override

Platforms

7705 SAR Gen 2

5.46 apply-path

apply-path

Syntax

[**no**] **apply-path**

Context

[Tree] (config>filter>match-list>ip-prefix-list apply-path)

[Tree] (config>filter>match-list>ipv6-prefix-list apply-path)

Full Context

configure filter match-list ip-prefix-list apply-path

configure filter match-list ipv6-prefix-list apply-path

Description

Commands in this context configure the auto-generation of address prefixes for IPv4 or IPv6 address prefix match lists. The context in which the command is executed governs whether IPv4 or IPv6 prefixes will be auto-generated.

The **no** form of this command removes all auto-generation configuration under the **apply-path** context.

Default

no apply path

Platforms

7705 SAR Gen 2

5.47 apply-to

apply-to

Syntax

apply-to {**all** | **none**}

Context

[Tree] (config>service>vprn>pim apply-to)

Full Context

configure service vprn pim apply-to

Description

This command creates a PIM interface with default parameters.

If a manually created interface or modified interface is deleted, the interface will be recreated when the **apply-to** command is executed. If PIM is not required on a specific interface, then execute a **shutdown** command.

The **apply-to** command is saved first in the PIM configuration structure, all subsequent commands either create new structures or modify the defaults as created by the **apply-to** command.

Default

apply-to none

Parameters

all

Specifies that all VPRN and non-VPRN interfaces are automatically applied in PIM.

none

No interfaces are automatically applied in PIM. PIM interfaces must be manually configured.

Platforms

7705 SAR Gen 2

apply-to

Syntax

apply-to {**ies** | **non-ies** | **all** | **none**}

Context

[Tree] (config>router>pim apply-to)

Full Context

configure router pim apply-to

Description

This command creates a PIM interface with default parameters.

If a manually created or a modified interface is deleted, the interface is recreated when (re)processing the **apply-to** command and if PIM is not required on a specific interface a shutdown should be executed.

The **apply-to** command is first saved in the PIM configuration structure. Then, all subsequent commands either create new structures or modify the defaults as created by the apply-to command.

Default

apply-to none

Parameters

ies

Specifies to apply all IES interfaces in PIM.

non-ies

Specifies to apply non-IES interfaces created in PIM.

all

Specifies to apply all IES and non-IES interfaces created in PIM.

none

Removes all interfaces that are not manually created or modified. It also removes explicit no interface commands if present.

Platforms

7705 SAR Gen 2

5.48 arbiter

arbiter

Syntax

arbiter *arbiter-name* [**create**]

no arbiter *arbiter-name*

Context

[Tree] (config>qos>plcr-ctrl-plcy>tier arbiter)

Full Context

configure qos policer-control-policy tier arbiter

Description

This command is used to create an arbiter within the context of **tier 1** or **tier 2**. An arbiter is a child policer bandwidth control object that manages the throughput of a set of child policers. An arbiter allows child policers or other arbiters to parent to one of eight strict levels. Each arbiter is itself parented to either another tiered arbiter or to the **root** arbiter.

The root arbiter starts with its defined maximum rate and distributes the bandwidth to its directly attached child policers and arbiters beginning with priority 8. As the children at each priority level are distributed bandwidth according to their needs and limits, the root proceeds to the next lower priority until either all children's needs are met or it runs out of bandwidth. The bandwidth given to a tiered arbiter is then divided between that arbiter's children (child policers or a tier 2 arbiter) in the same fashion. A tiered arbiter may also have a rate limit defined that limits the amount of bandwidth it may receive from its parent.

An arbiter that is currently parented by another arbiter cannot be deleted.

Each time the **policer-control-policy** is applied to either a SAP, or a subscriber (through association with a sub-profile that has the policy applied), or a multiservice site, an instance of the parent policer and the arbiters is created.

Any child policer that uses the arbiter's name in its parenting command will be associated with the arbiter instance. The child policer will also become associated with any arbiter to which its parent arbiter is parented (grandparent). Having child policers parented to an arbiter does not prevent that arbiter from being removed from the **policer-control-policy**. When removed, the child policers become orphaned.

You can create up to 31 tiered arbiters within the **policer-control-policy** on either tier 1 or tier 2 (in addition to the arbiter).

The **no** form of this command is used to remove an arbiter from tier 1 or tier 2. If the specified arbiter does not exist, the command returns without an error. If the specified arbiter is currently specified as the parent for another arbiter, the command will fail. When an arbiter is removed from a **policer-control-policy**, all instances of the arbiter will also be removed. Any child policers currently parented to the arbiter instance will become orphans and will not be bandwidth managed by the policer control policy instances parent policer.

Parameters

arbiter-name

Any unique name within the policy. Up to 31 arbiters may be created.

Platforms

7705 SAR Gen 2

5.49 area

area

Syntax

[**no**] **area** *area-id*

Context

[Tree] (config>service>vprn>ospf3 area)

[Tree] (config>service>vprn>ospf area)

Full Context

configure service vprn ospf3 area

configure service vprn ospf area

Description

This command creates the context to configure an OSPF area. An area is a collection of network segments within an AS that have been administratively grouped together. The area ID can be specified in dotted decimal notation or as a 32-bit decimal integer.

The **no** form of this command deletes the specified area from the configuration. Deleting the area also removes the OSPF configuration of all the interfaces, virtual-links, sham-links, address-ranges and so on, that are currently assigned to this area.

Default

no area — No OSPF areas are defined.

Parameters

area-id

The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

Values 0.0.0.0 to 255.255.255.255 (dotted decimal)
0 to 4294967295 (decimal integer)

Platforms

7705 SAR Gen 2

area

Syntax

[**no**] **area** *area-id*

Context

[Tree] (config>router>ospf area)

[Tree] (config>router>ospf3 area)

Full Context

configure router ospf area

configure router ospf3 area

Description

This command creates the context to configure an OSPF or OSPF3 area. An area is a collection of network segments within an AS that have been administratively grouped together. The area ID can be specified in dotted decimal notation or as a 32-bit decimal integer.

The **no** form of this command deletes the specified area from the configuration. Deleting the area also removes the OSPF configuration of all the interfaces, virtual-links, and address-ranges and so on, that are currently assigned to this area.

Default

no area

Parameters

area-id

The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

Values 0.0.0.0 to 255.255.255.255 (dotted decimal), 0 to 4294967295 (decimal integer)

Platforms

7705 SAR Gen 2

area

Syntax

area [*area-id*]

no area

Context

[Tree] (debug>router>ospf3 area)

[Tree] (debug>router>ospf area)

Full Context

debug router ospf3 area

debug router ospf area

Description

This command enables debugging for an OSPF area.

Parameters

area-id

Specifies the OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

Values ip-address — a.b.c.d

area — 0 to 4294967295

Platforms

7705 SAR Gen 2

area

Syntax

area *area-id*

no area

Context

[Tree] (config>router>policy-options>policy-statement>entry>from area)

Full Context

configure router policy-options policy-statement entry from area

Description

This command configures an OSPF area as a route policy match criterion.

This match criterion is only used in export policies.

All OSPF routes (internal and external) are matched using this criterion if the best path for the route is by the specified area.

The **no** form of this command removes the OSPF area match criterion.

Default

no area

Parameters

area-id

Specifies the OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

Values 0.0.0.0 to 255.255.255.255 (dotted decimal), 0 to 4294967295 (decimal)

Platforms

7705 SAR Gen 2

5.50 area-id

area-id

Syntax

[**no**] **area-id** *area-address*

Context

[Tree] (config>service>vprn>isis area-id)

Full Context

configure service vprn isis area-id

Description

This command configures the area ID portion of NSAP addresses for the VPRN instance. This identifies a point of connection to the network, such as a router interface, and is called a Network Service Access Point (NSAP). Addresses in the IS-IS protocol are based on the ISO NSAP addresses and Network Entity Titles (NETs), not IP addresses.

A maximum of 3 **area addresses** can be configured for the VPRN instance.

NSAP addresses are divided into three parts.

- Area ID — A variable length field between 1 and 13 bytes long. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.
- System ID — A six-byte system identification. When not configured, the system ID is derived from the configurations for **configure router isis router-id**, **configure router router-id**, or **system address** *ipv4 address*. If the previous commands are not configured, the system ID defaults to the last four octets of the chassis MAC address.
- Selector ID — A one-byte selector identification that must contain zeros when configuring a NET. This value is not configurable. The selector ID is always 00.

The NET is constructed like an NSAP but the selector byte contains a 00 value. NET addresses are exchanged in hello and LSP PDUs. All net addresses configured on the node are advertised to its neighbors.

For Level 1 interfaces, neighbors can have different area IDs, but, they must have at least one area ID (AFI + area) in common. Sharing a common area ID, they become neighbors and area merging between the potentially different areas can occur.

For Level 2 (only) interfaces, neighbors can have different area IDs. However, if they have no area IDs in common, they become only Level 2 neighbors and Level 2 LSPs are exchanged.

For Level 1 and Level 2 interfaces, neighbors can have different area IDs. If they have at least one area ID (AFI + area) in common, they become neighbors. In addition to exchanging Level 2 LSPs, area merging between potentially different areas can occur.

If multiple **area-id** commands are entered, the system ID of all subsequent entries must match the first **area** address.

The **no** form of this command removes the area address.

Platforms

7705 SAR Gen 2

area-id

Syntax

[**no**] **area-id** *area-address*

Context

[Tree] (config>router>isis area-id)

Full Context

configure router isis area-id

Description

This command was previously named the **net** *network-entity-title* command. The **area-id** command allows you to configure the area ID portion of NSAP addresses which identifies a point of connection to the network, such as a router interface, and is called a Network Service Access Point (NSAP). Addresses in the IS-IS protocol are based on the ISO NSAP addresses and Network Entity Titles (NETs), not IP addresses.

A maximum of three **area addresses** can be configured.

NSAP addresses are divided into three parts.

- Area ID — A variable length field between 1 and 13 bytes long. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.
- System ID — A six-byte system identification. When not configured, the system ID is derived from the configurations for **configure router isis router-id**, **configure router router-id**, or **system address** *ipv4 address*. If the previous commands are not configured, the system ID defaults to the last four octets of the chassis MAC address.
- Selector ID — A one-byte selector identification that must contain zeros when configuring a NET. This value is not configurable. The selector ID is always 00.

The NET is constructed like an NSAP but the selector byte contains a 00 value. NET addresses are exchanged in hello and LSP PDUs. All net addresses configured on the node are advertised to its neighbors.

For Level 1 interfaces, neighbors can have different area IDs, but, they must have at least one area ID (AFI + area) in common. Sharing a common area ID, they become neighbors and area merging between the potentially different areas can occur.

For Level 2 (only) interfaces, neighbors can have different area IDs. However, if they have no area IDs in common, they become only Level 2 neighbors and Level 2 LSPs are exchanged.

For Level 1 and Level 2 interfaces, neighbors can have different area IDs. If they have at least one area ID (AFI + area) in common, they become neighbors. In addition to exchanging Level 2 LSPs, area merging between potentially different areas can occur.

If multiple **area-id** commands are entered, the system ID of all subsequent entries must match the first **area** address.

The **no** form of this command removes the area address.

Parameters

area-address

Specifies a 1 — 13-byte address. Of the total 20 bytes comprising the NET, only the first 13 bytes can be manually configured. As few as one byte can be entered or, at most, 13 bytes. If less than 13 bytes are entered, the rest is padded with zeros.

Platforms

7705 SAR Gen 2

5.51 area-range

area-range

Syntax

area-range *ip-prefix/prefix-length* [**advertise** | **not-advertise**]

no area-range *ip-prefixImask*

area-range *ipv6-prefix/prefix-length* [**advertise** | **not-advertise**]

no area-range *ipv6-prefix/prefix-length*

Context

[Tree] (config>service>vprn>ospf>area>nssa area-range)

[Tree] (config>service>vprn>ospf>area area-range)

[Tree] (config>service>vprn>ospf3>area area-range)

[Tree] (config>service>vprn>ospf3>area>nssa area-range)

Full Context

configure service vprn ospf area nssa area-range

configure service vprn ospf area area-range

configure service vprn ospf3 area area-range

configure service vprn ospf3 area nssa area-range

Description

This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, it is configured to be advertised or not advertised into other areas. Multiple range commands are used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.

ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.

The **no** form of this command deletes the range (non) advertisement.

Default

no area-range

Parameters

ipv6-prefix/prefix-length

The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.

Values	ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x: [0 to FFFF]H
		d: [0 to 255]D
	ipv6-prefix-length	0 to 128

mask

The subnet mask for the range expressed as a decimal integer mask length or in dotted decimal notation.

Values 0 to 32 (mask length), 0.0.0.0 to 255.255.255.255 (dotted decimal)

advertise | not-advertise

Specifies whether or not to advertise the summarized range of addresses into other areas. The **advertise** keyword indicates the range will be advertised, and the keyword **not-advertise** indicates the range will not be advertised.

The default is **advertise**.

Platforms

7705 SAR Gen 2

area-range

Syntax

area-range *ip-prefixlmask* [**advertise** | **not-advertise**]

no area-range *ip-prefixlmask*

Context

[Tree] (config>router>ospf>area area-range)

[Tree] (config>router>ospf>area>nssa area-range)

Full Context

configure router ospf area area-range

configure router ospf area nssa area-range

Description

This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, the range is configured to be advertised or not advertised into other areas. Multiple range commands may be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.

ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.

The **no** form of this command deletes the range (non) advertisement.

Default

no area-range

Parameters

ip-prefix

Specifies the IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.

Values ip-prefix/mask: ip-prefix a.b.c.d (host bits must be 0)

mask

Specifies the subnet mask for the range expressed as a decimal integer mask length or in dotted decimal notation.

Values 0 to 32 (mask length), 0.0.0.0 to 255.255.255.255 (dotted decimal)

advertise | not-advertise

Specifies whether to advertise the summarized range of addresses into other areas. The **advertise** keyword indicates the range will be advertised, and the keyword **not-advertise** indicates the range will not be advertised.

Default advertise

Platforms

7705 SAR Gen 2

area-range

Syntax

area-range *ipv4-prefix/mask* | *ipv6-prefix/prefix-length* [**advertise** | **not-advertise**]

no area-range *ipv4-prefix/mask* | *ipv6-prefixIprefix-length*

Context

[Tree] (config>router>ospf3>area>nssa area-range)

[Tree] (config>router>ospf3>area area-range)

Full Context

configure router ospf3 area nssa area-range

configure router ospf3 area area-range

Description

This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, the range is configured to be advertised or not advertised into other areas. Multiple range commands may be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.

ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.

The **no** form of this command deletes the range (non) advertisement.

Default

no area-range

Parameters

ip-prefix/prefix-length

Specifies the IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.

Values ip-prefix/mask:

- ip-prefix a.b.c.d (host bits must be 0)

ipv6-prefix:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

prefix-length: 0 to 128

advertise | not-advertise

Specifies whether or not to advertise the summarized range of addresses into other areas. The **advertise** keyword indicates the range will be advertised, and the keyword **not-advertise** indicates the range will not be advertised.

Default advertise

Platforms

7705 SAR Gen 2

area-range

Syntax

area-range [*ip-address*]

no area-range

Context

[Tree] (debug>router>ospf3 area-range)

[Tree] (debug>router>ospf area-range)

Full Context

debug router ospf3 area-range

debug router ospf area-range

Description

This command enables debugging for an OSPF area range.

Parameters

ip-address

Specifies the IPv4 or IPv6 address for the range used by the ABR to advertise the area into another area.

Values ipv4-address:

- a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

5.52 arp

arp

Syntax

arp

Context

[Tree] (config>service>vprn>if>vpls>evpn arp)

[Tree] (config>service>ies>if>vpls>evpn arp)

Full Context

configure service vprn interface vpls evpn arp

configure service ies interface vpls evpn arp

Description

Commands in this context configure ARP host route parameters.

Platforms

7705 SAR Gen 2

arp

Syntax

arp

Context

[Tree] (debug>router>ip arp)

Full Context

debug router ip arp

Description

This command configures route table debugging.

Platforms

7705 SAR Gen 2

5.53 arp-host-route

arp-host-route

Syntax

arp-host-route

Context

[Tree] (config>service>vprn>if arp-host-route)

[Tree] (config>service>ies>if arp-host-route)

Full Context

configure service vprn interface arp-host-route

configure service ies interface arp-host-route

Description

Commands in this context configure ARP host routes to populate.

Platforms

7705 SAR Gen 2

5.54 arp-learn-unsolicited

arp-learn-unsolicited

Syntax

[**no**] **arp-learn-unsolicited**

Context

[Tree] (config>service>vprn>if arp-learn-unsolicited)

[Tree] (config>router>if arp-learn-unsolicited)

[Tree] (config>service>ies>if arp-learn-unsolicited)

Full Context

configure service vprn interface arp-learn-unsolicited

configure router interface arp-learn-unsolicited

configure service ies interface arp-learn-unsolicited

Description

This command allows the ARP application to learn new entries based on any received ARP message (GARP, ARP-Request, or ARP-Reply, such as any frame with ethertype 0x0806).

The **no** form of this command disables the above behavior and causes ARP entries to only be learned when needed, that is, when the router receives an ARP-reply after an ARP-request triggered by received traffic.

Platforms

7705 SAR Gen 2

5.55 arp-limit

arp-limit

Syntax

arp-limit *limit* [**log-only**] [**threshold** *percent*]

no arp-limit

Context

[Tree] (config>service>ies>interface arp-limit)

Full Context

configure service ies interface arp-limit

Description

This command configures the maximum amount of dynamic IPv4 ARP entries that can be learned on an IP interface.

When the number of dynamic ARP entries reaches the configured percentage of this limit, a log event is raised. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations will be dropped. Entries that have already been learned will be refreshed.

The **no** form of this command removes the **arp-limit**.

Default

no arp-limit

Parameters

log-only

Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned.

percent

The threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

Default 90

limit

The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic ARP learning is disabled and no dynamic ARP entries are learned.

Values 0 to 524288

Platforms

7705 SAR Gen 2

arp-limit

Syntax

arp-limit *limit* [**log-only**] [**threshold** *percent*]

no arp-limit

Context

[Tree] (config>service>vprn>if arp-limit)

Full Context

configure service vprn interface arp-limit

Description

This command configures the maximum amount of dynamic IPv4 ARP entries that can be learned on an IP interface.

When the number of dynamic ARP entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations will be dropped. Entries that have already been learned will be refreshed.

The **no** form of this command removes the **arp-limit**.

Default

90 percent

Parameters

log-only

Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned.

percent

The threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

limit

The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic ARP learning is disabled and no dynamic ARP entries are learned.

Values 0 to 524288

Platforms

7705 SAR Gen 2

arp-limit

Syntax

arp-limit *limit* [**log-only**] [**threshold** *percent*]

no arp-limit

Context

[Tree] (config>router>if arp-limit)

Full Context

configure router interface arp-limit

Description

This command configures the maximum amount of dynamic IPv4 ARP entries that can be learned on an IP interface.

When the number of dynamic ARP entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations will be dropped. Entries that have already been learned will be refreshed.

The **no** form of this command removes the **arp-limit.**

Default

no arp-limit

Parameters

limit

The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic ARP learning is disabled and no dynamic ARP entries are learned.

Values 0 to 524288

log-only

Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned.

percent

The threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

Platforms

7705 SAR Gen 2

5.56 arp-nd-extended-community-advertisement

arp-nd-extended-community-advertisement

Syntax

[**no**] **arp-nd-extended-community-advertisement**

Context

[Tree] (config>service>vpls>bgp-evpn arp-nd-extended-community-advertisement)

Full Context

configure service vpls bgp-evpn arp-nd-extended-community-advertisement

Description

This command enables the advertisement of the RFC 9047 ARP/ND extended community along with the MAC/IP routes that are advertised for local static and dynamic proxy ARP or ND entries. This command also controls the processing of the ARP/ND extended community and the selection of ARP or ND entries based on the inmutable flag.

The **no** form of this command disables the advertisement of the RFC 9047 ARP/ND extended community.

Default

no arp-nd-extended-community-advertisement

Platforms

7705 SAR Gen 2

5.57 arp-nd-only-with-fdb-advertisement

arp-nd-only-with-fdb-advertisement

Syntax

[**no**] **arp-nd-only-with-fdb-advertisement**

Context

[Tree] (config>service>vpls>bgp-evpn arp-nd-only-with-fdb-advertisement)

Full Context

configure service vpls bgp-evpn arp-nd-only-with-fdb-advertisement

Description

This command enables the router to advertise local ARP/ND entries of VPRN interfaces using this VPLS BGP-EVPN service when the corresponding local MAC is programmed in the FDB.

The **no** form of this command disables the advertisement of the ARP/ND entries.

Default

no arp-nd-only-with-fdb-advertisement

Platforms

7705 SAR Gen 2

5.58 arp-populate

arp-populate

Syntax

[**no**] **arp-populate**

Context

[Tree] (config>service>vprn>if arp-populate)

[Tree] (config>service>ies>if arp-populate)

Full Context

configure service vprn interface arp-populate

configure service ies interface arp-populate

Description

This command, when enabled, disables dynamic learning of ARP entries. Instead, the ARP table is populated with static and dynamic entries from the DHCP Lease State Table (enabled with **lease-populate**), and optionally with static entries entered with the **static-host** command.

The host's IP address and MAC address are placed in the system ARP cache as a managed entry. Static hosts must be defined on the interface using the **static-host** command. Dynamic hosts are enabled on the system through enabling lease-populate in the IP interface DHCP context.

In the event that both a static host and a dynamic host share the same IP and MAC address, the system's ARP cache retains the host information until both the static and dynamic information are removed.

Both static and dynamic hosts override static ARP entries. Static ARP entries are marked as inactive when they conflict with static or dynamic hosts and will be repopulated once all static and dynamic host information for the IP address are removed. Since static ARP entries are not possible when static subscriber hosts are defined or when DHCP lease state table population is enabled, conflict between static ARP entries and the arp-populate function is not an issue.

Enabling the **arp-populate** command removes any dynamic ARP entries learned on this interface from the ARP cache.

The **arp-populate** command fails if an existing static ARP entry exists for this interface.

When **arp-populate** is enabled, the system does not send out ARP requests for hosts that are not in the ARP cache. Only statically configured and DHCP learned hosts are reachable through an IP interface with **arp-populate** enabled. The **arp-populate** command can only be enabled on IES and VPRN interfaces supporting Ethernet encapsulation.

The **no** form of this command disables ARP cache population functions for static and dynamic hosts on the interface. All static and dynamic host information for this interface is removed from the system's ARP cache. Any existing static ARP entries previously inactive due to static or dynamic hosts will be populated in the system ARP cache.

Default

no arp-populate

Platforms

7705 SAR Gen 2

5.59 arp-proactive-refresh

arp-proactive-refresh

Syntax

[**no**] **arp-proactive-refresh**

Context

[Tree] (config>service>ies>if arp-proactive-refresh)

Full Context

configure service ies interface arp-proactive-refresh

Description

This command enables the router to always send out a single refresh message with no entries 30 seconds prior to the timeout of the entry.

The **no** form of this command sets the default behavior, in which an entry is marked as stale 30 seconds prior to age-out, and the router only sends an ARP request to refresh the entry if the IOM receives traffic that uses it. If so, the IOM asks the ARP application to send a refresh message. With **arp-proactive-refresh** enabled, the ARP module sends a refresh message regardless of whether the IOM receives traffic.

Platforms

7705 SAR Gen 2

arp-proactive-refresh

Syntax

[**no**] **arp-proactive-refresh**

Context

[Tree] (config>service>vprn>if arp-proactive-refresh)

Full Context

configure service vprn interface arp-proactive-refresh

Description

This command enables the router to always send out a refresh message 30 seconds prior to the timeout of the entry (a single refresh message with no retries).

The **no** form of this command sets the default behavior, in which an entry is marked as stale 30 seconds prior to age-out, and the router only sends an ARP request to refresh the entry if the IOM receives traffic that uses it. If so, the IOM asks the ARP application to send a refresh message. With **arp-proactive-refresh** enabled, the ARP module sends a refresh message regardless of the IOM receiving traffic.

Platforms

7705 SAR Gen 2

arp-proactive-refresh

Syntax

[**no**] **arp-proactive-refresh**

Context

[Tree] (config>router>if arp-proactive-refresh)

Full Context

configure router interface arp-proactive-refresh

Description

This command enables the router to always send out a refresh message 30 seconds prior to the timeout of the entry (a single refresh message with no retries).

The **no** form of this command sets the default behavior, in which an entry is marked as stale 30 seconds prior to age-out, and the router only sends an ARP request to refresh the entry if the IOM receives traffic that uses it. If so, the IOM asks the ARP application to send a refresh message. With **arp-proactive-refresh** enabled, the ARP module sends a refresh message regardless of the IOM receiving traffic.

Platforms

7705 SAR Gen 2

5.60 arp-retry-timer

arp-retry-timer

Syntax

arp-retry-timer *timer-multiple*

no arp-retry-timer

Context

[Tree] (config>service>ies>if arp-retry-timer)

Full Context

configure service ies interface arp-retry-timer

Description

This command allows the arp retry timer to be configured to a specific value.

The timer value is entered as a multiple of 100 ms. So a timer value of 1, means the ARP timer will be set to 100 ms.

The **no** form of this command removes the command from the active configuration and returns the ARP retry timer to its default value of 5 seconds.

Default

arp-retry-timer 50

Parameters

timer-multiple

Specifies the multiple of 100 ms that the ARP retry timer will be configured as.

Values 1 to 300 (equally a timer range of 100 ms to 30,000 ms)

Platforms

7705 SAR Gen 2

arp-retry-timer

Syntax

arp-retry-timer *timer-multiple*

no arp-retry-timer

Context

[Tree] (config>service>vprn>if arp-retry-timer)

[Tree] (config>service>vprn>network-interface arp-retry-timer)

Full Context

configure service vprn interface arp-retry-timer

configure service vprn network-interface arp-retry-timer

Description

This command allows the arp retry timer to be configured to a specific value.

The timer value is entered as a multiple of 100 ms. So a timer value of 1, means the ARP timer will be set to 100 ms.

The **no** form of this command removes the command from the active configuration and returns the ARP retry timer to its default value of 5 s.

Default

arp-retry-timer 50

Parameters

timer-multiple

Specifies the multiple of 100 ms that the ARP retry timer will be configured as.

Values 1 to 300 (equally a timer range of 100 ms to 30 000 ms)

Platforms

7705 SAR Gen 2

arp-retry-timer

Syntax

arp-retry-timer *timer-multiple*

no arp-retry-timer

Context

[Tree] (config>router>if arp-retry-timer)

Full Context

configure router interface arp-retry-timer

Description

This command allows the arp retry timer to be configured to a specific value.

The timer value is entered as a multiple of 100 ms. So a timer value of 1, means the ARP timer will be set to 100 ms.

The **no** form of this command removes the command from the active configuration and returns the ARP retry timer to its default value of 5 seconds.

Default

arp-retry-timer 50

Parameters

timer-multiple

Specifies the multiple of 100 ms that the ARP retry timer will be configured as.

Values 1 to 300 (equally a timer range of 100 ms to 30,000 ms)

Platforms

7705 SAR Gen 2

5.61 arp-timeout

arp-timeout

Syntax

arp-timeout *seconds*

no arp-timeout

Context

[Tree] (config>service>vprn>if arp-timeout)

[Tree] (config>service>ies>if arp-timeout)

Full Context

configure service vprn interface arp-timeout

configure service ies interface arp-timeout

Description

This command configures the minimum time in seconds an ARP entry learned on the IP interface is stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

When the **arp-populate** and **lease-populate** commands are enabled on an interface, the ARP table entries will no longer be dynamically learned, but instead by snooping DHCP ACK message from a DHCP server. In this case the configured **arp-timeout** value has no effect.

The default value for **arp-timeout** is 14400 seconds (4 hours).

The **no** form of this command reverts to the default value.

Default

arp-timeout 14400

Parameters

seconds

Specifies the minimum number of seconds a learned ARP entry is stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 to 65535

Platforms

7705 SAR Gen 2

arp-timeout

Syntax

arp-timeout *seconds*

no arp-timeout

Context

[Tree] (config>service>vpls>interface arp-timeout)

Full Context

configure service vpls interface arp-timeout

Description

This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

The default value for **arp-timeout** is 14400 seconds (4 hours).

The **no** form of this command restores **arp-timeout** to the default value.

Default

arp-timeout 14400

Parameters

seconds

The minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 to 65535

Platforms

7705 SAR Gen 2

arp-timeout

Syntax

arp-timeout *seconds*

no arp-timeout

Context

[Tree] (config>router>if arp-timeout)

Full Context

configure router interface arp-timeout

Description

This command configures the minimum time, in seconds, an ARP entry learned on the IP interface is stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host. Otherwise, the ARP entry is aged from the ARP table. If the **arp-timeout** value is set to 0 seconds, ARP aging is disabled.

The **no** form of this command reverts to the default value.

Default

no arp-timeout

Parameters

seconds

The minimum number of seconds a learned ARP entry is stored in the ARP table, expressed as a decimal integer. A value of 0 specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 to 65535

Platforms

7705 SAR Gen 2

5.62 as-override

as-override

Syntax

[**no**] **as-override**

Context

[Tree] (config>service>vprn>bgp>group as-override)

[Tree] (config>service>vprn>bgp>group>neighbor as-override)

Full Context

configure service vprn bgp group as-override

configure service vprn bgp group neighbor as-override

Description

This command replaces all instances of the peer's AS number with the local AS number in a BGP route's AS_PATH.

This command breaks BGP's loop detection mechanism. It should be used carefully.

Default

no as-override

Platforms

7705 SAR Gen 2

as-override

Syntax

[**no**] **as-override**

Context

[Tree] (config>router>bgp>group>neighbor as-override)

[Tree] (config>router>bgp>group as-override)

Full Context

configure router bgp group neighbor as-override

configure router bgp group as-override

Description

This command enables BGP to monitor the outbound routes toward the peer and whenever there is a route with the peer's autonomous system number (ASN) in the AS_PATH, all occurrences are removed and replaced with the advertising router's local ASN (or its confederation ID if the peer is outside the confederation).

In the group context, the **no** form of this command disables the functionality. In the neighbor context, the **no** form of this command causes the setting to be inherited from the group level.

Default

no as-override

Platforms

7705 SAR Gen 2

5.63 as-path

as-path

Syntax

[**no**] **as-path** *name*

Context

[Tree] (config>router>policy-options as-path)

Full Context

configure router policy-options as-path

Description

This command creates a route policy AS path to use in route policy entries.

The **no** form of this command deletes the AS path.

Default

no as-path

Parameters

name

The AS path regular expression name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

as-path

Syntax

as-path *name*

no as-path

Context

[Tree] (config>router>policy-options>policy-statement>entry>from as-path)

Full Context

configure router policy-options policy-statement entry from as-path

Description

This command configures an AS path regular expression statement as a match criterion for the route policy entry.

If no AS path criterion is specified, any AS path is considered to match.

AS path regular expression statements are configured at the global route policy level (**config>router>policy-options>as-path** *name*).

The **no** form of this command removes the AS path regular expression statement as a match criterion.

Default

no as-path

Parameters

name

Specifies the AS path regular expression name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@," "start@variable@end", " @variable@end", or "start@variable@".

Platforms

7705 SAR Gen 2

as-path

Syntax

as-path {**add** | **replace**} *name*

no as-path

Context

[Tree] (config>router>policy-options>policy-statement>default-action as-path)

[Tree] (config>router>policy-options>policy-statement>entry>action as-path)

Full Context

configure router policy-options policy-statement default-action as-path

configure router policy-options policy-statement entry action as-path

Description

This command assigns a BGP AS path list to routes matching the route policy statement entry.

If no AS path list is specified, the AS path attribute is not changed.

The **no** form of this command disables the AS path list editing action from the route policy entry.

Default

no as-path

Parameters

add

Specifies that the AS path list is to be prepended to an existing AS list.

replace

Specifies AS path list replaces any existing as path attribute.

name

Specifies the AS path list name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters

must be enclosed by at-signs (@) and may be midstring; for example, "@variable@," "start@variable@end"," @variable@end", or "start@variable@".

The *name* specified must already be defined.

Platforms

7705 SAR Gen 2

5.64 as-path-group

as-path-group

Syntax

[**no**] **as-path-group** *name*

Context

[Tree] (config>router>policy-options as-path-group)

Full Context

configure router policy-options as-path-group

Description

This command creates a route policy AS path regular expression statement to use in route policy entries.

The **no** form of this command deletes the AS path regular expression statement.

Default

no as-path-group

Parameters

name

Specifies the AS path regular expression name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@”.

Platforms

7705 SAR Gen 2

as-path-group

Syntax

as-path-group *name*

no as-path-group *name*

Context

[Tree] (config>router>policy-options>policy-statement>entry>from as-path-group)

Full Context

configure router policy-options policy-statement entry from as-path-group

Description

This command creates a route policy AS path regular expression statement to use in route policy entries.

The **no** form of this command deletes the AS path regular expression statement.

Default

no as-path-group

Parameters

name

Specifies the AS path regular expression name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@," "start@variable@end", " @variable@end", or "start@variable@".

Platforms

7705 SAR Gen 2

5.65 as-path-ignore

as-path-ignore

Syntax

as-path-ignore [**ipv4**] [**ipv6**] [**label-ipv4**] [**label-ipv6**]

no as-path-ignore

Context

[Tree] (config>service>vprn>bgp>path-selection as-path-ignore)

Full Context

configure service vprn bgp best-path-selection as-path-ignore

Description

This command configures whether AS path length is considered in the selection of the best BGP route for a prefix.

If an address family is listed in this command, the length of AS paths is not a factor in the route selection process for routes of that address family.

The **no** form of this command removes the parameter from the configuration.

Default

no as-path-ignore

Parameters

ipv4

Specifies that the AS path length is ignored for all unlabeled unicast IPv4 routes.

ipv6

Specifies that the AS path length is ignored for all unlabeled unicast IPv6 routes.

label-ipv4

Specifies that the AS path length is ignored for all labeled unicast IPv4 routes.

label-ipv6

Specifies that the AS path length is ignored for all labeled unicast IPv6 routes.

Platforms

7705 SAR Gen 2

as-path-ignore

Syntax

as-path-ignore [**ipv4**] [**label-ipv4**] [**vpn-ipv4**] [**ipv6**] [**label-ipv6**] [**vpn-ipv6**] [**mcast-ipv4**] [**mcast-ipv6**] [**mvpn-ipv4**] [**mvpn-ipv6**] [**l2-vpn**]

no as-path-ignore

Context

[Tree] (config>router>bgp>best-path-selection as-path-ignore)

Full Context

configure router bgp best-path-selection as-path-ignore

Description

This command configures whether AS path length is considered in the selection of the best BGP route for a prefix.

If an address family is listed in this command, then the length of AS paths is not a factor in the route selection process for routes of that address family.

The **no** form of this command removes the parameter from the configuration.

Default

no as-path-ignore

Parameters

ipv4

Specifies that the AS-path length will be ignored for all unlabeled unicast IPv4 routes.

label-ipv4

Specifies that the AS-path length will be ignored for all labeled-unicast IPv4 routes.

vpn-ipv4

Specifies that the length AS-path will be ignored for all VPN IPv4 (SAFI 128) routes.

ipv6

Specifies that the AS-path length will be ignored for all unlabeled unicast IPv6 routes.

label-ipv6

Specifies that the AS-path length will be ignored for all labeled-unicast IPv6 routes.

vpn-ipv6

Specifies that the AS-path length will be ignored for all VPN IPv6 (SAFI 128) routes.

mcast-ipv4

Specifies that the AS-path length will be ignored for all IPv4 multicast routes.

mcast-ipv6

Specifies that the AS-path length will be ignored for all IPv6 multicast routes.

mvpn-ipv4

Specifies that the AS-path length will be ignored for all IPv4 MVPN routes.

mvpn-ipv6

Specifies that the AS-path length will be ignored for all IPv6 MVPN routes.

l2-vpn

Specifies that the AS-path length will be ignored for all L2-VPN NLRIs.

Platforms

7705 SAR Gen 2

5.66 as-path-length

as-path-length

Syntax

as-path-length *length* [**equal** | **or-higher** | **or-lower**] [**unique**]

no as-path-length

Context

[Tree] (config>router>policy-options>policy-statement>entry>from as-path-length)

Full Context

configure router policy-options policy-statement entry from as-path-length

Description

This command matches BGP routes based on their AS path length (the number of AS numbers in the AS_PATH).

If no comparison qualifiers are present (**equal**, **or-higher**, **or-lower**), then **equal** is the implied default.

Confederation member AS numbers in the AS_PATH do not count towards the total. An AS_SET element is considered to have a length of 1.

The **unique** option counts.

A non-BGP route does not match a policy entry if it contains the **as-path-length** command.

Default

no as-path-length

Parameters

length

Specifies the length of the AS path.

Values 0 to 255, or a parameter name delimited by starting and ending at-sign (@) characters

equal

Specifies that matched routes should have the same number of AS path elements as the value specified.

or-higher

Specifies that matched routes should have the same or a greater number of AS path elements as the value specified.

or-lower

Specifies that matched routes should have the same or a lower number of AS path elements as the value specified.

unique

Specifies that only the unique AS numbers should be counted (that is, multiple occurrences of the same AS number in the sequence count as one).

Platforms

7705 SAR Gen 2

5.67 as-path-prepend

as-path-prepend

Syntax

as-path-prepend *as-path* [*repeat*]

as-path-prepend most-recent [*repeat*]

no as-path-prepend

Context

[Tree] (config>router>policy-options>policy-statement>default-action as-path-prepend)

[Tree] (config>router>policy-options>policy-statement>entry>action as-path-prepend)

Full Context

configure router policy-options policy-statement default-action as-path-prepend

configure router policy-options policy-statement entry action as-path-prepend

Description

The command prepends a BGP AS number once or numerous times to the AS path attribute of routes matching the route policy statement entry.

If an AS number is not configured, the AS path is not changed.

If the optional *number* is specified, then the AS number is prepended as many times as indicated by the number.

The **no** form of this command disables the AS path prepend action from the route policy entry.

Default

no as-path-prepend

Parameters

as-path

Specifies the AS number to prepend expressed as a decimal integer.

Values 1 to 4294967295

param-name — Specifies the AS path parameter variable name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@”.

repeat

Specifies the number of times to prepend the specified AS number expressed as a decimal integer.

Values 1 to 50

param-name — Specifies the AS path parameter variable name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@".

most-recent

Specifies that the most recent AS number must be prepended to the AS-Path attribute of the route.

Platforms

7705 SAR Gen 2

5.68 asbr

asbr

Syntax

[**no**] **asbr** [**trace-path** *domain-id*]

no asbr

[**no**] **asbr**

Context

[Tree] (config>router>ospf asbr)

[Tree] (config>router>ospf3 asbr)

Full Context

configure router ospf asbr

configure router ospf3 asbr

Description

This command configures the router as an Autonomous System Boundary Router (ASBR) if the router is to be used to export routes from the Routing Table Manager (RTM) into this instance of OSPF. After a router is configured as an ASBR, the export policies into this OSPF domain take effect. If no policies are configured, no external routes are redistributed into the OSPF domain.

The **no** form of this command removes the ASBR status and withdraws the routes redistributed from the Routing Table Manager into this instance of OSPF from the link state database.

When configuring multiple instances of OSPF, there is a risk of loops because networks are advertised by multiple domains configured with multiple interconnections to one another. To prevent this from happening, all routers in a domain should be configured with the same domain ID. Each domain (OSPF-instance) should be assigned a specific bit value in the 32-bit tag mask.

When an external route is originated by an ASBR using an internal OSPF route in a given domain, the corresponding bit is set in the AS-external LSA. As the route gets redistributed from one domain to another, more bits are set in the tag mask, each corresponding to the OSPF domain the route visited. Route redistribution looping is prevented by checking the corresponding bit as part of the export policy; if the bit corresponding to the announcing OSPF process is already set, the route is not exported there.

Domain IDs are incompatible with any other use of normal tags. The domain ID should be configured with a value between 1 and 31 by each router in a given OSPF domain (OSPF Instance).

When an external route is originated by an ASBR using an internal OSPF route in a given domain, the corresponding (1-31) bit is set in the AS-external LSA.

As the route gets redistributed from one domain to another, more bits are set in the tag mask, each corresponding to the OSPF domain the route visited. Route redistribution looping is prevented by checking the corresponding bit as part of the export policy; if the bit corresponding to the announcing OSPF process is already set, the route is not exported there.

Default

no asbr

Parameters

domain-id

Specifies the domain ID.

Values 1 to 31

Default 0

Platforms

7705 SAR Gen 2

5.69 assert

assert

Syntax

assert [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]

no assert

Context

[Tree] (debug>router>pim assert)

Full Context

debug router pim assert

Description

This command enables debugging for PIM assert mechanism.

The **no** form of this command disables PIM assert debugging.

Parameters

grp-ip-address

Debugs information associated with the PIM assert mechanism.

Values multicast group address (ipv4, ipv6)

ip-address

Debugs information associated with the PIM assert mechanism.

Values source address (ipv4, ipv6)

detail

Debugs detailed information on the PIM assert mechanism.

Platforms

7705 SAR Gen 2

5.70 assert-period

assert-period

Syntax

assert-period *assert-period*

no assert-period

Context

[Tree] (config>service>vprn>pim>if assert-period)

Full Context

configure service vprn pim interface assert-period

Description

This command configures the period in seconds for periodic refreshes of PIM Assert messages on an interface.

The **no** form of this command reverts to the default.

Default

assert-period 60

Parameters

assert-period

Specifies the period, in seconds, for periodic refreshes of PIM Assert messages on an interface.

Values 1 to 300

Platforms

7705 SAR Gen 2

assert-period

Syntax

assert-period *assert-period*

no assert-period

Context

[Tree] (config>router>pim>interface assert-period)

Full Context

configure router pim interface assert-period

Description

This command configures the period for periodic refreshes of PIM Assert messages on an interface.

The **no** form of this command removes the assert-period from the configuration.

Default

no assert-period

Parameters

assert-period

Specifies the period, in seconds, for periodic refreshes of PIM Assert messages on an interface.

Values 1 to 300

Platforms

7705 SAR Gen 2

5.71 assignment

assignment

Syntax

assignment {**port** *port-id* | **card** *slot-number*}

no assignment

Context

[Tree] (config>service>cust>multi-service-site assignment)

Full Context

configure service customer multi-service-site assignment

Description

This command assigns a multi-service customer site to a specific chassis slot, port, or channel. This allows the system to allocate the resources necessary to create the virtual schedulers defined in the ingress and egress scheduler policies as they are specified. This also verifies that each SAP assigned to the site exists within the context of the proper customer ID and that the SAP was configured on the proper slot, port, or channel. The assignment must be given prior to any SAP associations with the site.

The **no** form of this command removes the port, channel, or slot assignment. If the customer site has not yet been assigned, the command has no effect and returns without any warnings or messages.

Default

no assignment

Parameters

port-id

Assigns the multi-service customer site to the port-id or port-id.channel-id given. When the multi-service customer site is assigned to a specific port or channel, all SAPs associated with this customer site must be on a service owned by the customer and created on the defined port or channel. The defined port or channel must already have been pre-provisioned on the system but need not be installed when the customer site assignment is made.

Syntax: *port-id*[:encap-val]

Values

port-id	slot/mda/port[.channel]	
aps-id	aps-group-id[.channel]	
	aps keyword	
	group-id	1 to 128

		eth-tunnel-id	eth-tunnel-<id>	
			eth-tunnel	keyword
			id	1 to 1024
		lag-id	lag-id	
			lag	keyword
			id	1 to 800
			id	1 to 1024
		eth-sat-id	esat-<id>/<slot>/[u]<port>	
			esat	keyword
			id	1 to 20
			u	keyword for up-link port
		tdm-sat-id	tsat-<id>/<slot>/[<u>]<port>.<channel>	
			tsat	keyword
			id	1 to 20
			u	keyword for up-link port
		pxc-id	psc-id.sub-port	
			pxc psc-id.sub-port	
			pxc	keyword
			id: 1 to 64	
			sub-port: a, b	
		pw-id	pw-<id>	
			pw	keyword
			id	1 to 32767
	slot-number	1 to 10		
	fpe-id	1 to 64		

slot-number

Assigns the multi-service customer site to the slot-number given. When the multi-service customer site is assigned to a specific slot in the chassis, all SAPs associated with this customer site must be on a service owned by the customer and created on the defined chassis slot. The defined slot must already be pre-provisioned on the system but need not be installed when the customer site assignment is made.

Values Any pre-provisioned slot number for the chassis type that allows SAP creation.

1 to 20

fpe-id

Specifies the multi-service-site (MSS) assignment to an FPE object for the purpose of controlling aggregated bandwidth across a set of PW SAPs.

Values 1 to 64

Platforms

7705 SAR Gen 2

5.72 association-id

association-id

Syntax

association-id *association-id*

no association-id

Context

[Tree] (config>router>pcep>pcc>pce-assoc>div association-id)

Full Context

configure router pcep pcc pce-associations diversity association-id

Description

This command configures the diversity association ID. The user must specify an association ID.

The **no** form of the command removes the association ID from the diversity association.

Default

no association-id

Parameters

association-id

Specifies the diversity association ID.

Values 1 to 65535

Platforms

7705 SAR Gen 2

association-id

Syntax

association-id *association-id*

no association-id

Context

[Tree] (config>router>pcep>pcc>pce-assoc>plcy association-id)

Full Context

configure router pcep pcc pce-associations policy association-id

Description

This command configures the policy association ID. The user must specify an association ID.

The **no** form of the command removes the association ID from the policy association.

Default

no association-id

Parameters

association-id

Specifies the policy association ID.

Values 1 to 65535

Platforms

7705 SAR Gen 2

5.73 association-source

association-source

Syntax

association-source *ip-address*

no association-source

Context

[Tree] (config>router>pcep>pcc>pce-assoc>div association-source)

Full Context

configure router pcep pcc pce-associations diversity association-source

Description

This command configures the source IP address of the diversity association.

The **no** form of the command removes the IP address from the diversity association.

Default

no association-source

Parameters

ip-address

Specifies the source IP address.

Values

ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

Platforms

7705 SAR Gen 2

association-source

Syntax

association-source *ip-address*

no association-source

Context

[Tree] (config>router>pcep>pcc>pce-assoc>plcy association-source)

Full Context

configure router pcep pcc pce-associations policy association-source

Description

This command configures the source IP address of the policy association.

The **no** form of the command removes IP address from the policy association.

Default

no association-source

Parameters

ip-address

Specifies the source IP address.

Values

ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

Platforms

7705 SAR Gen 2

5.74 asynchronous-execution

asynchronous-execution

Syntax

asynchronous-execution *seconds*

asynchronous-execution never

Context

[Tree] (config>system>management-interface>ops>global-timeout asynchronous-execution)

Full Context

configure system management-interface operations global-timeouts asynchronous-execution

Description

This command configures the period of time that operations launched as "asynchronous" are allowed to execute before being automatically stopped by the SR OS.

An asynchronous operation is not deleted from the system when it is stopped. See the **asynchronous-retention** command.

If a specific execution timeout is not included in the request for a particular asynchronous operation, this system-level timeout applies.

Note:
This execution timeout is part of the general global operations infrastructure and is separate and independent from any operation-specific timeouts (for example, the **ping** operation also has its own **timeout** parameter).

Default

asynchronous-execution 3600

Parameters

seconds

Specifies the period of time, in seconds, that asynchronous operations are allowed to execute.

Values 1 to 604800

never

Keyword to specify that an execution timeout is not applied to asynchronous operations.

Platforms

7705 SAR Gen 2

5.75 asynchronous-retention

asynchronous-retention

Syntax

asynchronous-retention *seconds*

asynchronous-retention never

Context

[Tree] (config>system>management-interface>ops>global-timeout asynchronous-retention)

Full Context

configure system management-interface operations global-timeouts asynchronous-retention

Description

This command configures the period of time that data related to operations launched as "asynchronous" is retained in the system. After the retention timeout expires, all information related to the operation is deleted, including any status information and result data.

If a specific retention timeout is not included in the request for a particular asynchronous operation, this system-level timeout applies.

Default

asynchronous-retention 86400

Parameters

seconds

Specifies the period of time, in seconds, that data related to asynchronous operations is retained in the system.

Values 1 to 604800

never

Keyword to specify that data related to asynchronous operations will persist in memory until explicitly deleted.

Platforms

7705 SAR Gen 2

5.76 attempts

attempts

Syntax

attempts *count* [**time** *minutes1* [**lockout** *minutes2*]

no attempts

Context

[Tree] (config>system>security>password attempts)

Full Context

configure system security password attempts

Description

This command configures a threshold value of unsuccessful login attempts allowed in a specified time frame.

If the threshold is exceeded, the user is locked out for a specified time period.

If multiple **attempts** commands are entered, each command overwrites the previously entered command.

The **no attempts** command resets all values to default.

Note:
This command applies to a local user, in addition to users on RADIUS, TACACS, and LDAP.

Default

attempts 3 time 5 lockout 10

Parameters

count

Specifies the number of unsuccessful login attempts allowed for the specified **time**. This is a mandatory value that must be explicitly entered.

Values 1 to 64

minutes

Specifies the period of time, in minutes, that a specified number of unsuccessful attempts can be made before the user is locked out.

Values 0 to 60

minutes

Specifies the lockout period, in minutes, during which the user is not allowed to login.

Values 0 to 1440, or infinite

If the user exceeds the attempted **count** times in the specified **time**, then that user is locked out from any further login attempts for the configured lockout time period.

Values 0 to 1440

Values infinite; user is locked out and must wait until manually unlocked before any further attempts.

Platforms

7705 SAR Gen 2

attempts

Syntax

attempts [*count*] [**time** *minutes1*] [**lockout** *minutes2*]

no attempts

Context

[Tree] (config>system>security>snmp attempts)

Full Context

configure system security snmp attempts

Description

This command configures a threshold value of unsuccessful SNMPv2 or SNMPv3 connection attempts allowed in a specified time frame. The command parameters are used to counter denial of service (DoS) attacks through SNMP.

If the threshold is exceeded, the host is locked out for the lockout time period.

The **no** form of the command restores the default values.

Default

attempts 20 time 5 lockout 10

Parameters

count

Specifies the number unsuccessful SNMP attempts allowed for the specified **time**.

Values 1 to 64

minutes1

Specifies period of time, in minutes, that a specified number of unsuccessful attempts can be made before the host is locked out.

Values 0 to 60

minutes2

Specifies the lockout period in minutes where the host is not allowed to login. When the host exceeds the attempted count times in the specified time, then that host is locked out from any further login attempts for the configured time period.

Values 0 to 144**0**

Platforms

7705 SAR Gen 2

5.77 attrib

attrib

Syntax

attrib [**+r** | **-r**] *file-url*

attrib

Context

[Tree] (file attrib)

Full Context

file attrib

Description

This command sets or clears/resets the read-only attribute for a file in the local file system. To list all files and their current attributes enter **attrib** or **attrib x** where **x** is either the filename or a wildcard (*).

When an **attrib** command is entered to list a specific file or all files in a directory, the file's attributes are displayed with or without an "R" preceding the filename. The "R" implies that the **+r** is set and that the file is read-only. Files without the "R" designation implies that the **-r** is set and that the file is read-write-all. For example:

```
ALA-1>file cf3:\ # attrib
cf3:\bootlog.txt
cf3:\bof.cfg
cf3:\boot.ldr
cf3:\sr1.cfg
cf3:\test
cf3:\bootlog_prev.txt
cf3:\BOF.SAV
```

Parameters

file-url

Specifies the URL for the local file.

Values

local-url	[*cflash-id*/][*file-path*] up to 200 characters, including cflash-id directory length 99 chars max each
remote-url	[{ftp:// | tftp://}*login*:*pswd*@*remote-locn*/][*file-path*]
	up to 247 characters
	directory length up to 199 characters
remote-locn	[*hostname* | *ipv4-address* | [*ipv6-address*]]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D
	interface - up to 32 characters, for link local addresses 255
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

+r

Sets the read-only attribute on the specified file.

-r

Clears/resets the read-only attribute on the specified file.

Platforms

7705 SAR Gen 2

5.78 attribute-propagation

attribute-propagation

Syntax

[**no**] **attribute-propagation**

Context

[Tree] (config>service>system>bgp-evpn>ad-per-evi-routes attribute-propagation)

Full Context

configure service system bgp-evpn ad-per-evi-routes attribute-propagation

Description

This command enables attribute propagation in multi-instance Epipe services.

The **no** form of this command disables the propagation of attributes, including D-PATH, even if the **domain-id** is configured in the service.

Default

no attribute-propagation

Platforms

7705 SAR Gen 2

5.79 attribute-set

attribute-set

Syntax

attribute-set

Context

[Tree] (config>service>vprn>bgp attribute-set)

Full Context

configure service vprn bgp attribute-set

Description

Commands in this context configure the handling of attribute set (ATTR_SET) attributes in BGP routes received from PE-CE peers of the VPRN.

ATTR_SET is an optional transitive BGP path attribute standardized by RFC 6368 that is added to BGP Layer 3 VPN routes to provide logical separation between the BGP domain of a customer and the BGP domain of a service provider.

Platforms

7705 SAR Gen 2

attribute-set

Syntax

attribute-set

Context

[Tree] (config>service>vprn>bgp-ipvpn attribute-set)

Full Context

configure service vprn bgp-ipvpn attribute-set

Description

Commands in this context configure the handling of attribute set (ATTR_SET) attributes attached to VPN-IP routes imported into or exported from the VPRN.

ATTR_SET is an optional transitive BGP path attribute standardized by RFC 6368 that is added to BGP Layer 3 VPN routes to provide logical separation between the BGP domain of a customer and the BGP domain of a service provider.

Platforms

7705 SAR Gen 2

5.80 attribute-uniform-propagation

attribute-uniform-propagation

Syntax

[**no**] **attribute-uniform-propagation**

Context

[Tree] (config>service>system>bgp-evpn>ip-prefix-routes>iff attribute-uniform-propagation)

Full Context

configure service system bgp-evpn ip-prefix-routes interface-ful attribute-uniform-propagation

Description

This command enables the uniform propagation of BGP attributes for EVPN Interface-ful (EVPN-IFF) routes. EVPN-IFF is used in R-VPLS services with **bgp-evpn**>**ip-route-advertisement**. When enabled, the received EVPN-IFF routes for the R-VPLS can be propagated with the original BGP path attributes into EVPN-IFL, IPVPN, EVPN-IFF (in other R-VPLS services), or BGP IP routes advertised for the attached VPRN. This command also enables the attribute propagation in the opposite direction; for example, from EVPN-IFL, IPVPN, IP, or EVPN-IFF routes into EVPN-IFF routes.

The propagation is in accordance with the uniform mode defined in *draft-ietf-bess-evpn-ipvpn-interworking*.

The **no** form of this command re-originates the BGP path attributes when propagating EVPN-IFF routes into other inter-subnet forwarding families.

Default

no attribute-uniform-propagation

Platforms

7705 SAR Gen 2

5.81 augment-route-table

augment-route-table

Syntax

[**no**] **augment-route-table**

Context

[Tree] (config>router>isis>loopfree-alternates augment-route-table)

Full Context

configure router isis loopfree-alternates augment-route-table

Description

This command enables IS-IS to attach Remote LFA specific information to RTM entries for use by other protocols. This command requires **configure router isis lfa remote-lfa** to be enabled. Currently only LDP makes use of this additional information.

The **no** form of this command disables IS-IS to attach Remote LFA specific information to RTM entries for use by other protocols.

Platforms

7705 SAR Gen 2

augment-route-table

Syntax

[**no**] **augment-route-table**

Context

[Tree] (config>router>ospf>loopfree-alternates augment-route-table)

Full Context

configure router ospf loopfree-alternates augment-route-table

Description

This command enables OSPF to attach Remote LFA (rLFA) information to RTM entries for use by other protocols. Before this command is configured, the **configure router ospf lfa remote-lfa** command, must be enabled on the system. Currently, only LDP makes use of this additional information.

The **no** form of this command disables the attachment of rLFA-specific information to RTM entries for use by other protocols.

Default

no augment-route-table

Platforms

7705 SAR Gen 2

5.82 auth

auth

Syntax

[**no**] **auth**

Context

[Tree] (debug>router>rsvp>event auth)

Full Context

debug router rsvp event auth

Description

This command debugs auth events.

The **no** form of the command disables the debugging.

Platforms

7705 SAR Gen 2

auth

Syntax

[**no**] **auth** [**neighbor** *ip-int-name* | *ip-address*]

Context

[Tree] (debug>router>rip auth)

Full Context

debug router rip auth

Description

This command enables debugging for RIP authentication.

Parameters

ip-int-name* | *ip-address

Debugs the RIP authentication for the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

5.83 auth-keychain

auth-keychain

Syntax

auth-keychain *name*

Context

[Tree] (config>service>vprn>bgp>group auth-keychain)

[Tree] (config>service>vprn>bgp>group>neighbor auth-keychain)

[Tree] (config>service>vprn>bgp auth-keychain)

Full Context

configure service vprn bgp group auth-keychain

configure service vprn bgp group neighbor auth-keychain

configure service vprn bgp auth-keychain

Description

This command configures the BGP authentication key for all peers.

The keychain allows the rollover of authentication keys during the lifetime of a session.

Default

no auth-keychain

Parameters

name

Specifies the name of an existing keychain, up to 32 characters, to use for the specified TCP session or sessions.

Platforms

7705 SAR Gen 2

auth-keychain

Syntax

auth-keychain *name*

Context

[Tree] (config>service>vprn>isis auth-keychain)

[Tree] (config>service>vprn>isis>level auth-keychain)

Full Context

configure service vprn isis auth-keychain

configure service vprn isis level auth-keychain

Description

This command configures an authentication keychain to use for the protocol interface for the VPRN instance. The keychain allows the rollover of authentication keys during the lifetime of a session.

Default

no auth-keychain

Parameters

name

Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

Platforms

7705 SAR Gen 2

auth-keychain

Syntax

auth-keychain *name*

Context

[Tree] (config>router>isis auth-keychain)

[Tree] (config>router>isis>level auth-keychain)

Full Context

configure router isis auth-keychain

configure router isis level auth-keychain

Description

This command configures an authentication keychain to use for the protocol interface. The keychain allows the rollover of authentication keys during the lifetime of a session.

Parameters

name

Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

Platforms

7705 SAR Gen 2

auth-keychain

Syntax

auth-keychain *name*

Context

[Tree] (config>service>vprn>ospf>area>sham-link auth-keychain)

[Tree] (config>service>vprn>ospf>area>if auth-keychain)

[Tree] (config>service>vprn>ospf>area>virtual-link auth-keychain)

Full Context

configure service vprn ospf area sham-link auth-keychain

configure service vprn ospf area interface auth-keychain

configure service vprn ospf area virtual-link auth-keychain

Description

This command enables the authentication keychain.

Parameters

name

Specifies the name of the authentication keychain, up to 32 characters.

Platforms

7705 SAR Gen 2

auth-keychain

Syntax

auth-keychain *name*

Context

[Tree] (config>router>ldp>tcp-session-params auth-keychain)

[Tree] (config>router>ldp>tcp-session-params>peer-transport auth-keychain)

Full Context

configure router ldp tcp-session-parameters auth-keychain

configure router ldp tcp-session-parameters peer-transport auth-keychain

Description

This command configures the TCP authentication keychain to use for the TCP session. The per-peer authentication configuration takes precedence over the global authentication configuration.

Parameters

name

Specifies the name of the keychain, up to 32 characters. This keychain is used for the specified TCP session or sessions, and allows the rollover of authentication keys during the lifetime of a session. The peer address used must be the TCP session transport address.

Platforms

7705 SAR Gen 2

auth-keychain

Syntax

auth-keychain *name*

Context

[Tree] (config>router>rsvp>interface auth-keychain)

Full Context

configure router rsvp interface auth-keychain

Description

This command configures an authentication keychain to use for authentication of protocol messages sent and received over the associated interface. The keychain must include a valid entry to properly authenticate protocol messages, including a key, specification of a supported authentication algorithm, and beginning time. Each entry may also include additional options to control the overall lifetime of each entry to allow for the seamless rollover of without affecting the protocol adjacencies.

The **no** form of the auth-keychain command removes the association between the routing protocol and any keychain currently used.

Default

no auth-keychain

Parameters

name

Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

Platforms

7705 SAR Gen 2

auth-keychain

Syntax

auth-keychain *name*

Context

[Tree] (config>router>bgp>group auth-keychain)

[Tree] (config>router>bgp auth-keychain)

[Tree] (config>router>bgp>group>neighbor auth-keychain)

Full Context

configure router bgp group auth-keychain

configure router bgp auth-keychain

configure router bgp group neighbor auth-keychain

Description

This command configures a TCP authentication keychain to use for the session. The keychain allows the rollover of authentication keys during the lifetime of a session.

Default

no auth-keychain

Parameters

name

Specifies the name of the keychain, up to 32 characters, to use for the specified TCP session or sessions.

Platforms

7705 SAR Gen 2

auth-keychain

Syntax

auth-keychain

Context

[Tree] (config>router>ospf>area>virtual-link auth-keychain)

[Tree] (config>router>ospf>area>interface auth-keychain)

Full Context

configure router ospf area virtual-link auth-keychain

configure router ospf area interface auth-keychain

Description

This command configures an authentication keychain to use for the protocol interface. The keychain allows the rollover of authentication keys during the lifetime of a session.

The **no** form of this command removes the association to a previously specified keychain.

Default

no auth-keychain

Parameters

name

Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

Platforms

7705 SAR Gen 2

auth-keychain

Syntax

auth-keychain *name*

Context

[Tree] (config>router>pcep>pcc>peer auth-keychain)

Full Context

configure router pcep pcc peer auth-keychain

Description

This command specifies a keychain to be used for TCP-AO authentication between the PCC and the PCE. The keychain must first be configured in the **configure system security keychain** context.

Default

no auth-keychain

Parameters

name

Specifies the name of the keychain, up to 32 characters.

Platforms

7705 SAR Gen 2

5.84 auth-method

auth-method

Syntax

auth-method {**psk** | **plain-psk-xauth** | **cert-auth** | **psk-radius** | **cert-radius** | **eap** | **auto-eap-radius** | **auto-eap**}

no auth-method

Context

[Tree] (config>ipsec>ike-policy auth-method)

Full Context

configure ipsec ike-policy auth-method

Description

This command specifies the authentication method used with this IKE policy.

The **no** form of this command removes the parameter from the configuration.

Default

no auth-method

Parameters

psk

Both client and gateway authenticate each other by a hash derived from a pre-shared secret. Both client and gateway must have the PSK. This work with both IKEv1 and IKEv2

plain-psk-xauth

Both client and gateway authenticate each other by pre-shared key and RADIUS. This work with IKEv1 only.

psk-radius

Use the pre-shared-key and RADIUS to authenticate. IKEv2 remote-access tunnel only.

cert-radius

Use the certificate, public/private key and RADIUS to authenticate. IKEv2 remote-access tunnel only.

eap

Use the EAP to authenticate peer. IKEv2 remote-access tunnel only

auto-eap-radius

Use EAP or potentially other method to authenticate the peer. IKEv2 remote-access tunnel only. Also see **config>ipsec>ike-policy auto-eap-method** and **config>ipsec>ike-policy auto-eap-own-method**.

auto-eap

Use the EAP or potentially other RADIUS-related method to authenticate the peer. IKEv2 remote-access tunnel only. Also see **config>ipsec>ike-policy auto-eap-method** and **config>ipsec>ike-policy auto-eap-own-method**.

Platforms

7705 SAR Gen 2

5.85 auth-port

auth-port

Syntax

auth-port *port*

no auth-port

Context

[Tree] (config>service>vprn>radius-server>server auth-port)

[Tree] (config>router>radius-server>server auth-port)

Full Context

configure service vprn radius-server server auth-port

configure router radius-server server auth-port

Description

This command specifies the UDP listening port for RADIUS authentication requests.

The **no** form of this commands resets the UDP port to its default value (1812)

Default

auth-port 1812

Parameters

port

Specifies the UDP listening port for accounting requests of the external RADIUS server.

Values 1 to 65535

Platforms

7705 SAR Gen 2

5.86 authenticate

authenticate

Syntax

[**no**] **authenticate**

Context

[Tree] (config>service>vprn>ntp authenticate)

Full Context

configure service vprn ntp authenticate

Description

This command enables authentication for the NTP server.

Platforms

7705 SAR Gen 2

5.87 authenticate-client

authenticate-client

Syntax

authenticate-client

Context

[Tree] (config>system>security>tls>server-tls-profile authenticate-client)

Full Context

configure system security tls server-tls-profile authenticate-client

Description

Commands in this context configure client authentication parameters.

Platforms

7705 SAR Gen 2

5.88 authentication

authentication

Syntax

authentication bidirectional *sa-name*

authentication inbound *sa-name* **outbound** *sa-name*

no authentication

Context

[Tree] (config>service>vprn>ospf3>area>if authentication)

[Tree] (config>service>vprn>ospf3>area>virtual-link authentication)

Full Context

configure service vprn ospf3 area interface authentication

configure service vprn ospf3 area virtual-link authentication

Description

This command configures OPSFv3 confidentiality authentication.

The **no** form of this command removes the SA name from the configuration.

Parameters

bidirectional *sa-name*

Specifies the IPsec security association name in case the OSPFv3 traffic on the interface has to be authenticated.

inbound *sa-name*

Specifies the IPsec security association name in case the OSPFv3 traffic on the interface has to be authenticated.

outbound *sa-name*

Specifies the IPsec security association name in case the OSPFv3 traffic on the interface has to be authenticated.

Platforms

7705 SAR Gen 2

authentication

Syntax

authentication *ascii-algorithm* **ascii-key** *ascii-string* [**hash** | **hash2** | **custom**]

authentication *auth-algorithm* **hex-key** *hex-string* [**hash** | **hash2** | **custom**]

no authentication

Context

[Tree] (config>ipsec>static-sa authentication)

Full Context

configure ipsec static-sa authentication

Description

This command configures the authentication algorithm to use for an IPsec manual SA.

Default

no authentication

Parameters

auth-algorithm

Specifies the authentication algorithm to be used.

Values mda5, sha1

ascii-string

Specifies an ASCII key; 16 characters for **md5** and 20 characters for **sha1**.

hex-string

Specifies a HEX key; 32 hex nibbles for **md5** and 40 hex nibbles for **sha1**.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication

Syntax

authentication none

authentication *authentication-protocol authentication-key* [**privacy-none**] [**hash** | **hash2** | **custom**]

authentication *authentication-protocol authentication-key* **privacy** *privacy-protocol privacy-key* [**hash** | **hash2** | **custom**]

no authentication

Context

[Tree] (config>system>security>user>snmp authentication)

Full Context

configure system security user snmp authentication

Description

This command configures the SNMPv3 authentication and privacy protocols for the user to communicate with the router. The keys are stored in an encrypted format in the configuration.

The keys configured with these commands must be localized keys, which are a hash of the SNMP engine ID and a password. The password is not entered directly in this command. Use the **tools perform system management-interface snmp generate-key** command to generate localized authentication and privacy keys.

Default

authentication none

Parameters

none

Keyword to specify that no authentication protocol is used. If **none** is specified, privacy cannot be configured.

authentication-protocol

Specifies the SNMPv3 authentication protocol.

Values **hmac-md5-96** — Specifies use of the HMAC-MD5-96 authentication protocol.

hmac-sha1-96 — Specifies use of the HMAC-SHA-96 authentication protocol.

hmac-sha2-224 — Specifies use of the HMAC-SHA-224 authentication protocol.

hmac-sha2-256 — Specifies use of the HMAC-SHA-256 authentication protocol.

hmac-sha2-384 — Specifies use of the HMAC-SHA-384 authentication protocol.

hmac-sha-512 — Specifies use of the HMAC-SHA-512 authentication protocol.

authentication-key

Specifies the localized authentication key, which is entered as a hexadecimal string; the character length depends on the specified authentication protocol. The following table lists the authentication protocol key lengths.

Table 17: Authentication protocol key lengths

Authentication protocol	Character lengths
HMAC-MD5-96	32
HMAC-SHA-96	40
HMAC-SHA-224	56
HMAC-SHA-256	64
HMAC-SHA-384	96
HMAC-SHA-512	128

privacy-none

Keyword to specify that a privacy protocol is not used in the communication.

Default privacy none

privacy-protocol

Specifies the SNMPv3 privacy protocol.

Values **cbc-des** — Specifies the use of the CBC-DES privacy protocol.

cfb128-aes-128 — Specifies the use of the CFB128-AES-128 privacy protocol.

cfb128-aes-192 — Specifies the use of the CFB128-AES-192 privacy protocol.

cfb128-aes-256 — Specifies the use of the CFB128-AES-256 privacy protocol.

privacy-key

Specifies the localized privacy key, which is entered as a hexadecimal string; the character length depends on the specified privacy protocol. The following table lists the privacy protocol key lengths.

Table 18: Privacy protocol key lengths

Privacy protocol	Character length
CBC-DES	32
CFB128-AES-128	32
CFB128-AES-192	48
CFB128-AES-256	64

hash

Keyword that specifies the key is entered in an encrypted form. If the **hash** or **hash2** keyword is not specified, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Keyword that specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone; that is, the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** keyword is not specified, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Keyword that specifies the custom encryption to the management interface.

Platforms

7705 SAR Gen 2

authentication

Syntax

authentication bidirectional *sa-name*

authentication [**inbound** *sa-name* **outbound** *sa-name*]

no authentication

Context

[Tree] (config>router>ospf3>area>virtual-link authentication)

[Tree] (config>router>ospf3>area>interface authentication)

Full Context

configure router ospf3 area virtual-link authentication

configure router ospf3 area interface authentication

Description

This command configures the password used by the OSPF3 interface or virtual-link to send and receive OSPF3 protocol packets on the interface when simple password authentication is configured.

All neighboring routers must use the same type of authentication and password for proper protocol communication.

By default, no authentication key is configured.

The **no** form of this command removes the authentication.

Default

no authentication

Parameters

bidirectional sa-name

Specifies bidirectional OSPF3 authentication.

inbound sa-name

Specifies the inbound security association (SA) name for OSPF3 authentication.

outbound sa-name

Specifies the outbound SA name for OSPF3 authentication.

Platforms

7705 SAR Gen 2

5.89 authentication-check

authentication-check

Syntax

[**no**] **authentication-check**

Context

[Tree] (config>service>vprn>isis authentication-check)

Full Context

configure service vprn isis authentication-check

Description

This command sets an authentication check to reject PDUs that do not match the type or key requirements for the VPRN instance.

The default behavior when authentication is configured is to reject all IS-IS protocol PDUs that have a mismatch in either the authentication type or authentication key.

When **no authentication-check** is configured, authentication PDUs are generated and IS-IS PDUs are authenticated on receipt. However, mismatches cause an event to be generated and will not be rejected.

The **no** form of this command allows authentication mismatches to be accepted and generates a log event.

Default

authentication-check — Rejects authentication mismatches.

Platforms

7705 SAR Gen 2

authentication-check

Syntax

[**no**] **authentication-check**

Context

[Tree] (config>service>vprn>ntp authentication-check)

Full Context

configure service vprn ntp authentication-check

Description

This command provides the option to skip the rejection of NTP PDUs that do not match the authentication key-id, type or key requirements. The default behavior when authentication is configured is to reject all NTP protocol PDUs that have a mismatch in either the authentication key-id, type or key.

When **authentication-check** is enabled, NTP PDUs are authenticated on receipt. However, mismatches cause a counter to be increased, one counter for type and one for key-id, one for type, value mismatches. These counters are visible in a show command.

The **no** form of this command allows authentication mismatches to be accepted; the counters however are maintained.

Default

authentication-check — Rejects authentication mismatches.

Platforms

7705 SAR Gen 2

authentication-check

Syntax

[**no**] **authentication-check**

Context

[Tree] (config>system>time>ntp authentication-check)

Full Context

configure system time ntp authentication-check

Description

This command provides the option to skip the rejection of NTP PDUs that do not match the authentication key-id, type or key requirements. The default behavior when authentication is configured is to reject all NTP protocol PDUs that have a mismatch in either the authentication key-id, type or key.

When **authentication-check** is enabled, NTP PDUs are authenticated on receipt. However, mismatches cause a counter to be increased, one counter for type and one for key-id, one for type, value mismatches. These counters are visible in a show command.

The **no** form of this command allows authentication mismatches to be accepted; the counters however are maintained.

Default

authentication-check

Platforms

7705 SAR Gen 2

authentication-check

Syntax

[**no**] **authentication-check**

Context

[Tree] (config>router>isis authentication-check)

Full Context

configure router isis authentication-check

Description

This command sets an authentication check to reject PDUs that do not match the type or key requirements.

The default behavior when authentication is configured is to reject all IS-IS protocol PDUs that have a mismatch in either the authentication type or authentication key.

When **no authentication-check** is configured, authentication PDUs are generated and IS-IS PDUs are authenticated on receipt. However, mismatches cause an event to be generated and will not be rejected.

The **no** form of this command allows authentication mismatches to be accepted and generates a log event.

Default

authentication-check

Platforms

7705 SAR Gen 2

5.90 authentication-key

authentication-key

Syntax

authentication-key {*authentication-key* | *hash-key*} [**hash** | **hash2** | **custom**]

no authentication-key

Context

[Tree] (config>redundancy>multi-chassis>peer authentication-key)

Full Context

configure redundancy multi-chassis peer authentication-key

Description

This command configures the authentication key used between this node and the multi-chassis peer. The authentication key can be any combination of letters or numbers. The **no** form of the command removes the authentication key.

Default

no authentication-key

Parameters

authentication-key

Specifies the authentication key. Allowed values are any string up to 20 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 33 (hash1-key) or 55 (hash2-key) characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2** | **custom**]

no authentication-key

Context

[Tree] (config>subscr-mgmt>rip-policy authentication-key)

Full Context

configure subscriber-mgmt rip-policy authentication-key

Description

This command configures the BGP authentication key.

Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD5 message-based digest. The authentication key can be any combination of letters or numbers from 1 to 16.

The **no** form of this command removes the authentication password from the configuration and effectively disables authentication.

Default

Authentication is disabled and the authentication password is empty.

Parameters

authentication-key

Specifies the authentication key. The key can be any combination of ASCII characters up to 255 characters (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 342 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys

are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key {*authentication-key* | *hash-key*} [**hash** | **hash2**| **custom**]

no authentication-key

Context

[Tree] (config>service>ies>if>vrrp authentication-key)

Full Context

configure service ies interface vrrp authentication-key

Description

The **authentication-key** command, within the **vrrp** *virtual-router-id* context, is used to assign a simple text password authentication key to generate master VRRP advertisement messages and validating received VRRP advertisement messages.

The authentication-key command is one of the few commands not affected by the presence of the owner keyword. If simple text password authentication is not required, the authentication-key command is not required. If the command is re-executed with a different password key defined, the new key will be used immediately. If a no authentication-key command is executed, the password authentication key is restored to the default value. The authentication-key command may be executed at any time.

To change the current in-use password key on multiple virtual router instances:

- Identify the current master
- Shutdown the virtual router instance on all backups
- Execute the authentication-key command on the master to change the password key
- Execute the authentication-key command and no shutdown command on each backup key

The **no** form of the command removes the authentication key.

Default

No default. The authentication data field contains the value 0 in all 16 octets.

Parameters

authentication-key

The *key* parameter identifies the simple text password used when VRRP Authentication Type 1 is enabled on the virtual router instance. Type 1 uses a string eight octets long that is inserted into all transmitted VRRP advertisement messages and compared against all received VRRP advertisement messages. The authentication data fields are used to transmit the key.

The *key* parameter is expressed as a string consisting up to eight alpha-numeric characters. Spaces must be contained in quotation marks (" "). The quotation marks are not considered part of the string.

The string is case sensitive and is left-justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet (starting with IETF RFC bit position 0) containing the first character. The second field holds the fifth through eighth characters. Any unspecified portion of the authentication data field is padded with the value 0 in the corresponding octet.

Values Any 7-bit printable ASCII character.

Exceptions:	Double quote (")	ASCII 34
	Carriage Return	ASCII 13
	Line Feed	ASCII 10
	Tab	ASCII 9
	Backspace	ASCII 8

hash-key

The hash key. The key can be any combination of ASCII characters up to 22 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key {*authentication-key* | *hash-key*} [**hash** | **hash2** | **custom**]

no authentication-key

Context

[Tree] (config>service>vprn>bgp authentication-key)

[Tree] (config>service>vprn>bgp>group authentication-key)

[Tree] (config>service>vprn>bgp>group>neighbor authentication-key)

Full Context

configure service vprn bgp authentication-key

configure service vprn bgp group authentication-key

configure service vprn bgp group neighbor authentication-key

Description

This command configures the BGP authentication key.

Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD5 message-based digest. The authentication key can be any combination of letters or numbers from 1 to 16.

The **no** form of this command removes the authentication password from the configuration and effectively disables authentication.

Default

no authentication-key

Parameters

authentication-key

Specifies an authentication key. The key can be up to 255 characters (unencrypted).

hash-key

The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2** | **custom**]

no authentication-key

Context

[Tree] (config>service>vprn>if>vrrp authentication-key)

Full Context

configure service vprn interface vrrp authentication-key

Description

The **authentication-key** command, within the **vrrp** *virtual-router-id* context, is used to assign a simple text password authentication key to generate master VRRP advertisement messages and validate received VRRP advertisement messages.

The **authentication-key** command is one of the few commands not affected by the presence of the **owner** keyword. If simple text password authentication is not required, this command is not required. If the command is re-executed with a different password key defined, the new key will be used immediately. If a no **authentication-key** command is executed, the password authentication key is restored to the default value. The **authentication-key** command may be executed at any time.

To change the current in-use password key on multiple virtual router instances:

- Identify the current master
- Shut down the virtual router instance on all backups
- Execute the **authentication-key** command on the master to change the password key
- Execute the **authentication-key** command and the **no shutdown** command on each backup key

The **no** form of this command restores the default null string to the value of key.

Parameters

authentication-key

The *key* parameter identifies the simple text password used when VRRP Authentication Type 1 is enabled on the virtual router instance. Type 1 uses a string eight octets long that is inserted into all transmitted VRRP advertisement messages and compared against all received VRRP advertisement messages. The authentication data fields are used to transmit the key.

The *key* parameter is expressed as a string consisting of up to eight alpha-numeric characters. Spaces must be contained in quotation marks (" ”). The quotation marks are not considered part of the string.

The string is case sensitive and is left-justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet (starting with IETF RFC bit position 0) containing the first character. The second field holds the fifth through eighth characters. Any unspecified portion of the authentication data field is padded with the value 0 in the corresponding octet.

Values Any 7-bit printable ASCII character.

Exceptions:	Double quote (")	ASCII 34
	Carriage Return	ASCII 13
	Line Feed	ASCII 10
	Tab	ASCII 9
	Backspace	ASCII 8

hash-key

The hash key. The key can be any combination of ASCII characters up to 22 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ”)

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key {*authentication-key* | *hash-key*} [**hash** | **hash2** | **custom**]

no authentication-key

Context

[Tree] (config>service>vprn>isis authentication-key)

[Tree] (config>service>vprn>isis>level authentication-key)

Full Context

configure service vprn isis authentication-key

configure service vprn isis level authentication-key

Description

This command sets the authentication key used to verify PDUs sent by neighboring routers on the interface for the VPRN instance.

Neighboring routers use passwords to authenticate PDUs sent from an interface. For authentication to work, both the authentication *key* and the authentication *type* on a segment must match. The OSPF Commands statement must also be included.

To configure authentication on the global level, configure this command in the **config>router>isis** context. When this parameter is configured on the global level, all PDUs are authenticated including the Hello PDU.

To override the global setting for a specific level, configure the **authentication-key** command in the **config>router>isis>level** context. When configured within the specific level, hello PDUs are not authenticated.

The **no** form of this command removes the authentication key.

Default

no authentication-key — No authentication key is configured.

Parameters

authentication-key

The authentication key. The key can be any combination of ASCII characters up to 255 characters in length (un-encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key *key-id* **key** *key* [**hash** | **hash2** | **custom**] **type** {**des** | **message-digest**}

no authentication-key *key-id*

Context

[Tree] (config>service>vprn>ntp authentication-key)

Full Context

configure service vprn ntp authentication-key

Description

This command sets the authentication key-id, type and key used to authenticate NTP PDUs sent by the broadcast server function toward external clients or to authenticate NTP PDUs received from external unicast clients within the VPRN routing instance. For authentication to work, the authentication key-id, type, and key value must match.

The **no** form of this command removes the authentication key.

Parameters

key-id

Configure the authentication key-id that will be used by the node when transmitting or receiving Network Time Protocol packets.

Entering the authentication-key command with a key-id value that matches an existing configuration key will result in overriding the existing entry.

Recipients of the NTP packets must have the same authentication key-id, type, and key value in order to use the data transmitted by this node. This is an optional parameter.

Values 1 to 255

key

The authentication key associated with the configured key-id, the value configured in this parameter is the actual value used by other network elements to authenticate the NTP packet.

The key can be any combination of ASCII characters up to 8 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (".").

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

type

This parameter determines if DES or message-digest authentication is used.

This is a required parameter; either DES or message-digest must be configured.

Values des — Specifies that DES authentication is used for this key.

message-digest — Specifies that MD5 authentication in accordance with RFC 2104 is used for this key.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key {*authentication-key* | *hash-key*} [**hash** | **hash2** | **custom**]

no authentication-key

Context

[Tree] (config>service>vprn>ospf>area>sham-link authentication-key)

[Tree] (config>service>vprn>ospf>area>virtual-link authentication-key)

[Tree] (config>service>vprn>ospf>area>if authentication-key)

Full Context

configure service vprn ospf area sham-link authentication-key

configure service vprn ospf area virtual-link authentication-key

configure service vprn ospf area interface authentication-key

Description

This command configures the password used by the OSPF interface or virtual-link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.

This command is not valid in the OSPF3 context.

All neighboring routers must use the same type of authentication and password for proper protocol communication. If the **authentication-type** is configured as password, then this key must be configured.

By default, no authentication key is configured.

This command is not supported in the OSPF context.

The **no** form of this command removes the authentication key.

Default

no authentication-key — No authentication key is defined.

Parameters

authentication-key

The authentication key. The key can be any combination of ASCII characters up to 8 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

The hash key. The key can be any combination of ASCII characters up to 22 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key {*authentication-key* | *hash-key*} [**hash** | **hash2** | **custom**]

no authentication-key

Context

[Tree] (config>service>vprn>rip>group authentication-key)

[Tree] (config>service>vprn>rip>group>neighbor authentication-key)

[Tree] (config>service>vprn>rip authentication-key)

Full Context

configure service vprn rip group authentication-key

configure service vprn rip group neighbor authentication-key

configure service vprn rip authentication-key

Description

This command sets the authentication password to be passed between RIP neighbors.

The authentication type and authentication key must match exactly to authenticate and then process the RIP message.

The **no** form of this command removes the authentication password from the configuration and disables authentication.

Default

no authentication-key

Parameters

authentication-key

The authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

The hash key. The key can be any combination of ASCII characters up to 33 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key {*authentication-key* | *hash-key*} [**hash** | **hash2** | **custom**]

no authentication-key

Context

[Tree] (config>router>ldp>tcp-session-params authentication-key)

[Tree] (config>router>ldp>tcp-session-params>peer-transport authentication-key)

Full Context

configure router ldp tcp-session-parameters authentication-key

configure router ldp tcp-session-parameters peer-transport authentication-key

Description

This command specifies the authentication key used to establish a session between LDP peers. Authentication uses the MD5 message-based digest. The peer address used in authentication must be the TCP session transport address. If one or more transport addresses used in the Hello adjacencies to the same peer LSR are different from the LSR-ID value, the user must add each transport address to the authentication-key configuration as a separate peer. As a result, when the TCP connection is bootstrapped by a specific Hello adjacency, the authentication can operate over that TCP connection by using its specific transport address. The per peer authentication configuration takes precedence over global authentication configuration, and authentication keychain configuration takes precedence over authentication key configuration.

The **no** form of this command disables authentication.

Parameters

authentication-key

Specifies the authentication key. The key can be any combination of ASCII characters, up to 255 characters (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

Specifies the hash key. The key can be any combination of up to 33 alphanumeric characters. If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2

Specifies the key is entered in a more complex, encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to the management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key {*authentication-key* | *hash-key*} [**hash** | **hash2** | **custom**]

no authentication-key

Context

[Tree] (config>router>rsvp>interface authentication-key)

Full Context

configure router rsvp interface authentication-key

Description

This command specifies the authentication key for use between RSVP neighbors to authenticate RSVP messages. Authentication uses the MD5 message-based digest.

When enabled on an RSVP interface, authentication of RSVP messages operates in both directions of the interface. A router maintains a security association using one authentication key for each interface to an RSVP neighbor.

An RSVP neighbor transmits an authenticating digest of the RSVP message that is computed using the shared authentication key and a keyed-hash algorithm. The message digest is included in an INTEGRITY object, which also contains a flags field, a key identifier field, and a sequence number field. An RSVP neighbor uses the key together with the authentication algorithm to process received RSVP messages. The RSVP MD5 authentication complies to the procedures for RSVP message generation in RFC 2747, *RSVP Cryptographic Authentication*.

The MD5 implementation does not support the authentication challenge procedures in RFC 2747.

The **no** form of this command disables authentication.

Default

no authentication-key - The authentication key value is the null string.

Parameters

authentication-key

Specifies the authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

hash-key

Specifies the hash key. The key can be any combination of up 33 alphanumeric characters. If spaces are used in the string, enclose the entire string in quotation marks (" ”)

This is useful when a user must configure the parameter, but for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be

copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key {*authentication-key* | *hash-key*} [{**hash** | **hash2** | **custom**}]

no authentication-key

Context

[Tree] (config>router>if>vrrp authentication-key)

Full Context

configure router interface vrrp authentication-key

Description

This command sets the simple text authentication key used to generate master VRRP advertisement messages and validates VRRP advertisements.

If simple text password authentication is not required, the **authentication-key** command is not required.

The command is configurable in both non-owner and owner **vrrp** nodal contexts.

The *key* parameter identifies the simple text password to be used when VRRP Authentication Type 1 is enabled on the virtual router instance. Type 1 uses an eight octet long string that is inserted into all transmitted VRRP advertisement messages and is compared against all received VRRP advertisement messages. The authentication data fields are used to transmit the *key*.

The *key* string is case sensitive and is left justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet (starting with IETF RFC bit position 0) containing the first character. The second field similarly holds the fifth through eighth characters. Any unspecified portion of the authentication data field is padded with a 0 value in the corresponding octet.

If the command is re-executed with a different password key defined, the new key is used immediately.

The **authentication-key** command can be executed at anytime.

To change the current in-use password key on multiple virtual router instances:

Identify the current master.

1. Shutdown the virtual router instance on all backups.
2. Execute the **authentication-key** command on the master to change the password key.
3. Execute the **authentication-key** command and **no shutdown** command on each backup.

The **no** form of the command reverts to the default value.

Default

no authentication-key — The authentication key value is the null string.

Parameters

authentication-key

The authentication key. Allowed values are any string up to 8 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

hash-key

The hash key. The key can be any combination of ASCII characters up to 22 (hash-key1) or 121 (hash-key2) characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key *key-id* **key** *key* [**hash** | **hash2** | **custom**] **type** {**des** | **message-digest**}

no authentication-key *key-id*

Context

[Tree] (config>system>time>ntp authentication-key)

Full Context

configure system time ntp authentication-key

Description

This command sets the authentication key-id, type and key used to authenticate NTP PDUs sent to or received by other network elements participating in the NTP protocol. For authentication to work, the authentication key-id, type and key value must match.

The **no** form of the command removes the authentication key.

Parameters

key-id

Configures the authentication key-id that will be used by the node when transmitting or receiving Network Time Protocol packets

Entering the authentication-key command with a key-id value that matches an existing configuration key will result in overriding the existing entry.

Recipients of the NTP packets must have the same authentication key-id, type, and key value in order to use the data transmitted by this node. This is an optional parameter.

Values 1 to 255

key

Specifies the authentication key associated with the configured key-id, the value configured in this parameter is the actual value used by other network elements to authenticate the NTP packet.

The key can be any combination of ASCII characters up to 32 characters for message-digest (md5) or 8 characters for des (length limits are unencrypted lengths). If spaces are used in the string, enclose the entire string in quotation marks (".").

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

type

Determines if DES or message-digest authentication is used.

This is a required parameter; either DES or message-digest must be configured.

des

Specifies that DES authentication is used for this key.

message-digest

Specifies that MD5 authentication in accordance with RFC 2104 is used for this key.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2** | **custom**]

no authentication-key

Context

[Tree] (config>router>bgp>group authentication-key)

[Tree] (config>router>bgp>group>neighbor authentication-key)

[Tree] (config>router>bgp authentication-key)

Full Context

configure router bgp group authentication-key

configure router bgp group neighbor authentication-key

configure router bgp authentication-key

Description

This command configures the BGP authentication key.

Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD5 message based digest.

The **no** form of this command reverts to the default value.

Default

no authentication-key

Parameters

authentication-key

Specifies an authentication key. The key can be up to 255 characters (unencrypted).

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2** | **custom**]

no authentication-key

Context

[Tree] (config>router>isis authentication-key)

[Tree] (config>router>isis>level authentication-key)

Full Context

configure router isis authentication-key

configure router isis level authentication-key

Description

This command sets the authentication key used to verify PDUs sent by neighboring routers on the interface.

Neighboring routers use passwords to authenticate PDUs sent from an interface. For authentication to work, both the authentication *key* and the authentication *type* on a segment must match. The **authentication-type** command must also be included.

To configure authentication on the global level, configure this command in the **config>router>isis** context. When this parameter is configured on the global level, all PDUs are authenticated, including the hello PDU.

To override the global setting for a specific level, configure the **authentication-key** command in the **config>router>isis>level** context. When configured within the specific level, hello PDUs are not authenticated.

The **no** form of this command removes the authentication key.

Default

no authentication-key

Parameters

authentication-key

Specifies the authentication key. The key can be any combination of ASCII characters up to 255 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2** | **custom**]

no authentication-key

Context

[Tree] (config>router>ospf>area>interface authentication-key)

[Tree] (config>router>ospf>area>virtual-link authentication-key)

Full Context

configure router ospf area interface authentication-key

configure router ospf area virtual-link authentication-key

Description

This command configures the password used by the OSPF interface or virtual link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.

All neighboring routers must use the same type of authentication and password for proper protocol communication. If **authentication-type password** is configured, this key must be configured.

By default, no authentication key is configured.

The **no** form of this command removes the authentication key.

Default

no authentication-key

Parameters

authentication-key

Specifies the authentication key. The key can be any combination of ASCII characters up to 8 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 22 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

authentication-key

Syntax

authentication-key {*authentication-key* | *hash-key*} [{**hash** | **hash2** | **custom**}]

no authentication-key

Context

[Tree] (config>router>rip>group>neighbor authentication-key)

[Tree] (config>router>rip>group authentication-key)

[Tree] (config>router>rip authentication-key)

Full Context

configure router rip group neighbor authentication-key

configure router rip group authentication-key

configure router rip authentication-key

Description

This command sets the authentication password to be passed between RIP neighbors.

The authentication type and authentication key must match exactly for the RIP message to be considered authentic and processed.

The **no** form of the command removes the authentication password from the configuration and disables authentication.

Default

no authentication-key

Parameters

authentication-key

Specifies the authentication key. Allowed values are any string up to 16 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 33 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

5.91 authentication-keychain

authentication-keychain

Syntax

authentication-keychain *keychain-name*

no authentication-keychain

Context

[Tree] (config>system>time>ntp authentication-keychain)

Full Context

configure system time ntp authentication-keychain

Description

This command configures the authentication keychain used to handle unsolicited NTP requests.

If the system receives a request with a key ID that matches both the configured key and the keychain, the system checks the MAC ID using the key information first. If the key authentication fails, the system then checks the MAC ID using the information from the keychain.

The **no** form of the command removes the authentication keychain.

Parameters

keychain-name

Specifies the keychain name, up to 32 characters.

Platforms

7705 SAR Gen 2

authentication-keychain

Syntax

authentication-keychain *keychain-name*

no authentication-keychain

Context

[Tree] (config>service>vprn>ntp authentication-keychain)

Full Context

configure service vprn ntp authentication-keychain

Description

This command configures the authentication keychain used to handle unsolicited NTP requests.

If the system receives a request with a key ID that matches both the configured key and the keychain, the system checks the MAC ID using the key information first. If the key authentication fails, the system then checks the MAC ID using the information from the keychain.

The **no** form of the command removes the authentication keychain.

Parameters

keychain-name

Specifies the keychain name, up to 32 characters.

Platforms

7705 SAR Gen 2

5.92 authentication-method

authentication-method

Syntax

authentication-method

Context

[Tree] (config>system>security>ssh authentication-method)

Full Context

configure system security ssh authentication-method

Description

Commands in this context configure at the system level the SSH authentication method.

Platforms

7705 SAR Gen 2

5.93 authentication-order

authentication-order

Syntax

authentication-order [*method-1*] [*method-2*] [*method-3*] [*method-4*] [**exit-on-reject**]

no authentication-order

Context

[Tree] (config>system>security>password authentication-order)

Full Context

configure system security password authentication-order

Description

This command configures the sequence in which the system attempts authentication and authorization among the local user database, RADIUS servers, TACACS+ servers, and LDAP servers.

Configure the order from the most preferred method to the least preferred. The presence of all methods in the command line does not guarantee they are all operational. Specifying options that are not available delays user authentication.

If all operational methods are attempted and no authentication for a particular login has been granted, an entry in the security log records the failed attempt. Both the attempted login identification and originating IP address are logged with a timestamp.

The **no** form of this command reverts to the default order.

The order is not applicable to SNMPv3. SNMPv3 messages ignore the configured order and are authorized using the locally configured users only. TACACS+, RADIUS, and LDAP are not supported for SNMPv3 authentication.

Note:
This command applies to a local user, in addition to users on RADIUS, TACACS+, and LDAP.

Default

authentication-order radius tacplus ldap local

Parameters

method-1

Specifies the first password authentication method to attempt.

Values local, radius, tacplus, ldap

method-2

Specifies the second password authentication method to attempt.

Values local, radius, tacplus, ldap

method-3

Specifies the third password authentication method to attempt.

Values local, radius, tacplus, ldap

method-4

Specifies the fourth password authentication method to attempt.

Values local, radius, tacplus, ldap

local

Specifies the password authentication based on the local password database.

radius

Specifies RADIUS authentication.

tacplus

Specifies TACACS+ authentication.

ldap

Specifies LDAP authentication.

exit-on-reject

When this parameter is configured, the router stops authentication and authorization if one of the AAA methods configured in the order sends a rejection.

When this parameter is not configured, the router attempts the next AAA method if a AAA method sends a rejection. If all AAA methods are exhausted, authentication and authorization are rejected.

If the order specifies **local** as the first method, the following actions apply:

- If this parameter is configured and the user does not exist, the user is not authenticated.
- If the user can be authenticated locally, other methods, if configured, are used for authorization and accounting.
- If the user is configured locally but without console access, login is denied.

Platforms

7705 SAR Gen 2

5.94 authentication-over-bypass

authentication-over-bypass

Syntax

authentication-over-bypass [**enable** | **disable**]

Context

[Tree] (config>router>rsvp authentication-over-bypass)

Full Context

configure router rsvp authentication-over-bypass

Description

This command configures the MD5 authentication over the bypass LSP of all Point of Local Repairs (PLRs) and Merge Points (MPs) on the router. Only enable this command when the TE interfaces in the RSVP-TE network use the same MD5 authentication parameters.

When a Point of Local Repair (PLR) activates a bypass LSP towards a Merge Point (MP), by default, the INTEGRITY object corresponding to the bypass LSP interface is not added to a transmitted RSVP message except for packets of routed RSVP messages (Resv, Srefresh, and ACK), and only when the packet is intended for a bypass LSP endpoint (PLR or MP) that is a directly connected neighbor.

When this command is enabled, the INTEGRITY object of the interface corresponding to the bypass LSP is added to a transmitted RSVP message regardless of whether the bypass LSP endpoint (PLR or MP) is a directly connected RSVP neighbor. The INTEGRITY object is included with the following RSVP messages: Path, PathTear, PathErr, Resv, ResvTear, ResvErr, Srefresh, and ACK.

In all cases, an RSVP message received from a PLR or a MP (sender address in the SenderTemplate/ FilterSpec is different from an Extended Tunnel Id in a Session Object), and which includes the INTEGRITY object is authenticated against the bypass LSP interface. An RSVP message received from a PLR or MP without the INTEGRITY object is also accepted.

Default

authentication-over-bypass disable

Parameters

enable

Enables the MD5 authentication over the bypass LSP of all PLRs on the node.

disable

Disables the MD5 authentication over the bypass LSP of all PLRs on the node.

Platforms

7705 SAR Gen 2

5.95 authentication-type

authentication-type

Syntax

authentication-type {**none** | **password** | **message-digest** | **message-digest-20**}

no authentication-type

Context

[Tree] (config>subscr-mgmt>rip-plcy authentication-type)

Full Context

configure subscriber-mgmt rip-policy authentication-type

Description

This command sets the type of authentication to be used between RIP neighbors. The type and password must match exactly for the RIP message to be considered authentic and processed.

The **no** form of this command removes the authentication type from the configuration and effectively disables authentication.

Parameters

none

Disables authentication at a given level (global, group, neighbor). If the command does not exist in the configuration, the parameter is inherited.

password

Specifies enable simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple password authentication is enabled.

message-digest

Configures 16 byte message digest for MD5 authentication. If this option is configured, then at least one message-digest-key must be configured.

message-digest-20

Configures 20 byte message digest for MD5 authentication in accordance with RFC 2082, RIP-2 MD5 Authentication. If this option is configured, then at least one message-digest-key must be configured.

Platforms

7705 SAR Gen 2

authentication-type

Syntax

authentication-type {**password** | **message-digest**}

no authentication

Context

[Tree] (config>service>vprn>isis authentication-type)

[Tree] (config>service>vprn>isis>level authentication-type)

Full Context

configure service vprn isis authentication-type

configure service vprn isis level authentication-type

Description

This command enables either simple password or message digest authentication or must go in either the global IS-IS or IS-IS level context.

Both the authentication key and the authentication type on a segment must match. The **authentication-key** statement must also be included.

Configure the authentication type on the global level in the **config>router>isis** context.

Configure or override the global setting by configuring the authentication type in the **config>router>isis>level** context.

The **no** form of this command disables authentication.

Default

no authentication-type — No authentication type is configured and authentication is disabled.

Parameters

password

Specifies that simple password (plain text) authentication is required.

message-digest

Specifies that MD5 authentication in accordance with RFC 2104 is required.

Platforms

7705 SAR Gen 2

authentication-type

Syntax

authentication-type {**password** | **message-digest**}

no authentication-type

Context

[Tree] (config>service>vprn>ospf>area>sham-link authentication-type)

[Tree] (config>service>vprn>ospf>area>virtual-link authentication-type)

[Tree] (config>service>vprn>ospf>area>if authentication-type)

Full Context

configure service vprn ospf area sham-link authentication-type

configure service vprn ospf area virtual-link authentication-type

configure service vprn ospf area interface authentication-type

Description

This command enables authentication and specifies the type of authentication to be used on the OSPF interface, virtual-link, and sham-link.

This command is not valid in the OSPF3 context.

Both simple **password** and **message-digest** authentication are supported.

The **no** form of this command disables authentication on the interface.

Default

no authentication-type — No authentication is enabled on an interface.

Parameters

password

This keyword enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple **password** authentication is enabled.

message-digest

This keyword enables message digest MD5 authentication in accordance with RFC 1321. If this option is configured, then at least one message-digest-key must be configured.

Platforms

7705 SAR Gen 2

authentication-type

Syntax

authentication-type {**none** | **password** | **message-digest** | **message-digest-20**}

no authentication-type

Context

[Tree] (config>service>vprn>rip>group authentication-type)

[Tree] (config>service>vprn>rip authentication-type)

[Tree] (config>service>vprn>rip>group>neighbor authentication-type)

Full Context

configure service vprn rip group authentication-type

configure service vprn rip authentication-type

configure service vprn rip group neighbor authentication-type

Description

This command defines the type of authentication used between RIP neighbors. The type and password must match exactly to authenticate and then process the RIP message.

The **no** form of this command removes the authentication type from the configuration and effectively disables authentication.

Default

no authentication-type

Parameters

none

No authentication is used.

password

A simple cleartext password is sent.

message-digest

MD5 authentication is used.

message-digest-20

MD20 authentication is used.

Platforms

7705 SAR Gen 2

authentication-type

Syntax

authentication-type {**password** | **message-digest**}

no authentication

Context

[Tree] (config>router>isis>level authentication-type)

[Tree] (config>router>isis authentication-type)

Full Context

configure router isis level authentication-type

configure router isis authentication-type

Description

This command enables either simple password or message digest authentication or must go in either the global IS-IS or IS-IS level context.

Both the authentication key and the authentication type on a segment must match. The **authentication-key** statement must also be included.

Configure the authentication type on the global level in the **config>router>isis** context.

Configure or override the global setting by configuring the authentication type in the **config>router>isis>level** context.

The **no** form of this command disables authentication.

Parameters

password

Specifies that simple password (plain text) authentication is required.

message-digest

Specifies that MD5 authentication in accordance with RFC 2104 is required.

Platforms

7705 SAR Gen 2

authentication-type

Syntax

authentication-type {**password** | **message-digest**}

no authentication-type

Context

[Tree] (config>router>ospf>area>virtual-link authentication-type)

[Tree] (config>router>ospf>area>interface authentication-type)

Full Context

configure router ospf area virtual-link authentication-type

configure router ospf area interface authentication-type

Description

This command enables authentication and specifies the type of authentication to be used on the OSPF interface.

Both simple **password** and **message-digest** authentication are supported.

By default, authentication is not enabled on an interface.

The **no** form of this command disables authentication on the interface.

Default

no authentication-type

Parameters

password

Enables the simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple **password** authentication is enabled.

message-digest

Enables message digest MD5 authentication in accordance with RFC 1321. If this option is configured, then at least one message-digest-key must be configured.

Platforms

7705 SAR Gen 2

authentication-type

Syntax

authentication-type {**none** | **password** | **message-digest** | **message-digest-20**}

no authentication-type

Context

[Tree] (config>router>rip>group>neighbor authentication-type)

[Tree] (config>router>rip>group authentication-type)

[Tree] (config>router>rip authentication-type)

Full Context

configure router rip group neighbor authentication-type

configure router rip group authentication-type

configure router rip authentication-type

Description

This command sets the type of authentication to be used between RIP neighbors.

The type and password must match exactly for the RIP message to be considered authentic and processed.

The **no** form of the command removes the authentication type from the configuration and effectively disables authentication.

Default

no authentication-type

Parameters

none

The **none** parameter explicitly disables authentication at a given level (global, group, neighbor). If the command does not exist in the configuration, the parameter is inherited.

password

Specifies that the password enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple **password** authentication is enabled.

message-digest

Configures 16 byte message digest for MD5 authentication. If this option is configured, then at least one **message-digest-key must** be configured.

message-digest-20

Configures 20 byte message digest for MD5 authentication in accordance with RFC 2082, *RIP-2 MD5 Authentication*. If this option is configured, then at least one **message-digest-key** must be configured.

Platforms

7705 SAR Gen 2

5.96 authenticator-init

authenticator-init

Syntax

[**no**] **authenticator-init**

Context

[Tree] (config>port>ethernet>dot1x>per-host-authentication authenticator-init)

Full Context

configure port ethernet dot1x per-host-authentication authenticator-init

Description

This command configures the authenticator-initiated mode of the host.

The **no** form of this command disables the authenticator-initiated mode of the host.

Default

authenticator-init

Platforms
7705 SAR Gen 2

5.97 authorization

authorization

Syntax
authorization

Context
[Tree] (config>system>security>cli-script authorization)

Full Context
configure system security cli-script authorization

Description
Commands in this context authorize CLI script execution.

Platforms
7705 SAR Gen 2

authorization

Syntax
[**no**] **authorization**

Context
[Tree] (config>service>vprn>aaa>rmt-srv>radius authorization)
[Tree] (config>system>security>radius authorization)

Full Context
configure service vprn aaa remote-servers radius authorization
configure system security radius authorization

Description
This command configures RADIUS authorization parameters.
The **no** form of this command removes RADIUS authorization parameters from the configuration.

Default

no authorization

Platforms

7705 SAR Gen 2

authorization

Syntax

authorization [**use-priv-lvl**]

no authorization

Context

[Tree] (config>service>vprn>aaa>rmt-srv>tacplus authorization)

[Tree] (config>system>security>tacplus authorization)

Full Context

configure service vprn aaa remote-servers tacplus authorization

configure system security tacplus authorization

Description

This command configures TACACS+ command authorization parameters.

If this command is enabled without the **use-priv-lvl** option, each command is sent to the TACACS+ server for authorization (this is true whether the **tacplus use-default-template** setting is enabled or not).

If the **tacplus authorization** command is disabled, and the **tacplus use-default-template** setting is enabled, the local profile in the **user-template tacplus_default** is used for command authorization.

The **no** form of this command removes authorization parameters from the configuration.

Default

no authorization

Parameters

use-priv-lvl

Specifies to automatically perform a single authorization request to the TACACS+ server for cmd* (all commands) immediately after login, and then use the local profile associated (via the **priv-lvl-map** command) with the privilege level returned by the TACACS+ server for all subsequent authorization (except **enable-admin**). After the initial authorization for cmd*, no further authorization requests are sent to the TACACS+ server (except **enable-admin**). If the TACACS+ server does not return a privilege level for a user, the profile from the **user-template tacplus_default** is used for command authorization (as long as **tacplus use-default-template** is enabled, otherwise all commands are rejected).

Platforms

7705 SAR Gen 2

5.98 auto-bind-tunnel

auto-bind-tunnel

Syntax

auto-bind-tunnel

Context

[Tree] (config>service>vpls>bgp-evpn>mpls auto-bind-tunnel)

[Tree] (config>service>vprn>bgp-evpn>mpls auto-bind-tunnel)

[Tree] (config>service>epipe>bgp-evpn>mpls auto-bind-tunnel)

[Tree] (config>service>vprn>bgp-ipvpn>mpls auto-bind-tunnel)

Full Context

configure service vpls bgp-evpn mpls auto-bind-tunnel

configure service vprn bgp-evpn mpls auto-bind-tunnel

configure service epipe bgp-evpn mpls auto-bind-tunnel

configure service vprn bgp-ipvpn mpls auto-bind-tunnel

Description

Commands in this context configure automatic binding of a VPRN service using tunnels to MP-BGP peers.

The **auto-bind-tunnel** node is simply a context to configure the binding of BGP IPVPN or EVPN routes to tunnels. The user must configure the **resolution** option to enable auto-bind resolution to tunnels in TTM. If the **resolution** option is explicitly set to **disabled**, the auto-binding to tunnel is removed.

If resolution is set to **any**, any supported tunnel type in the Epipe/VPRN/VPLS context is selected following TTM preference. If one or more explicit tunnel types are specified using the **resolution-filter** option, then only these tunnel types are selected again following the TTM preference.

The user must set **resolution** to **filter** in order to activate the list of tunnel-types configured under resolution-filter.

In VPRN services and for BGP-IPVPN, when an explicit SDP to a BGP next hop is configured (**config**>**service**>v**prn**>**spoke-sdp**), it overrides the auto-bind-tunnel selection for that BGP next hop only. There is no support for reverting automatically to the auto-bind-tunnel selection if the explicit SDP goes down. The user must delete the explicit spoke-sdp in the VPRN service context to resume using the auto-bind-tunnel selection for the BGP next hop.

Platforms

7705 SAR Gen 2

auto-bind-tunnel

Syntax

auto-bind-tunnel

Context

[Tree] (config>service>vprn auto-bind-tunnel)

Full Context

configure service vprn auto-bind-tunnel

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

5.99 auto-boot

auto-boot

Syntax

auto-boot [**management-port**] [**inband** [**vlan** *vlan-id* | **vlan-discovery**]] [**ipv4**] [**ipv6**] [**client-identifier** {**string** *ascii-string* | **hex** *hex-string* | **chassis-mac**}] [**include-user-class**] [**timeout** *minutes*]

auto-boot ospf [**neid** *neid-hex-string*] [**vendor-id** *vendor-id*] [**neip-ipv4** *ip-address*] [**neip-ipv6** *ipv6-address*] [**port-mtu** *mtu-bytes*] [**ospf-mtu** *ip-mtu-bytes*] [**vlan** *vlan-id*] [**timeout** *minutes*]

no auto-boot

Context

[Tree] (bof auto-boot)

Full Context

bof auto-boot

Description

This command enables the **auto-boot** flag in the BOF and configures the **auto-boot** options for ZTP. When modifying **auto-boot** options using CLI, all required options must be explicitly configured, as the default cases will no longer be used.

The **no** form of this command disables the **auto-boot** flag.

Default

no auto-boot

Parameters

management-port

Specifies that the out-of-band management port (Mgmt port) should be used for ZTP.

inband

Specifies that in-band management through an Ethernet port should be used for ZTP. Unless the **vlan-discovery** flag is used, the **inband** option disables VLAN discovery.

vlan-id

Specifies an in-band VLAN to use for the auto-boot process.

Values 1 to 4094

vlan-discovery

Floods all VLANs (1 to 4094) with DHCP discovery messages and is supported only on **inband** ports. The first offer received on a specific VLAN is processed.

ipv4

Enables IPv4 DHCP discovery. This parameter is mandatory if the **ipv6** parameter is not specified.

ipv6

Enables IPv6 DHCP solicitation. This parameter is mandatory if the **ipv4** parameter is not specified.

ascii-string

Specifies a DHCP client identification string, up to 58 ASCII characters, to be used for Option 61 (IPv4) or Option 1 (IPv6).

hex-string

Specifies a DHCP client identification string, up to 116 hexadecimal nibbles, to be used for Option 61 (IPv4) or Option 1 (IPv6).

Values 0x0 to 0xFFFFFFFF

chassis-mac

Specifies that the chassis MAC address should be used as the DHCP client identification string for Option 61 (IPv4) or Option 1 (IPv6).

include-user-class

Specifies that Option 77 should be included in DHCP messages.

client-identifier

Specifies that a custom client ID should be used in network discovery requests.

minutes

Specifies the time interval after which, if the auto-boot process is unsuccessful (in the case of auto-boot using OSPF, if no OSPF adjacency is found), the node is rebooted and the auto-boot process is retried.

Values 30 to 1440

Default 30

ospf

Specifies that OSPF auto-discovery should be used.

neid-hex-string

Specifies a hexadecimal network element identification string.

Values 0x10101to 0xFEFEFE

ip-address

Specifies the IPv4 address for the network element.

Values a.b.c.d

Default *vendor-id.neid-hex-string*

ipv6-address

Specifies the IPv6 address for the network element.

Values

x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x [0 to FFFF]H

d [0 to 255]D

Default The IPv6 version of *vendor-id.neid-hex-string*

vendor-id

Specifies the vendor identification number. The number 140 corresponds to "Nokia”.

Values 1 to 254

Default 140

ip-mtu-bytes

Specifies the OSPF MTU in bytes.

Values 512 to 9786

Default 1500

mtu-bytes

Specifies the port MTU in bytes.

Values 512 to 9800

Default The default MTU of the port type.

Platforms

7705 SAR Gen 2

5.100 auto-config

auto-config

Syntax

[**no**] **auto-config**

Context

[Tree] (config>service>epipe>spoke-sdp-fec auto-config)

Full Context

configure service epipe spoke-sdp-fec auto-config

Description

This command enables single sided automatic endpoint configuration of the spoke SDP. The router acts as the passive T-PE for signaling this MS-PW.

Automatic Endpoint Configuration allows the configuration of a spoke SDP endpoint without specifying the TAII associated with that spoke SDP. It allows a single-sided provisioning model where an incoming label mapping message with a TAII that matches the SAII of that spoke SDP to be automatically bound to that endpoint. In this mode, the far end T-PE actively initiates MS-PW signaling and will send the initial label mapping message using T-LDP, while the router T-PE for which auto-config is specified will act as the passive T-PE.

The **auto-config** command is blocked in CLI if signaling active has been enabled for this spoke SDP. It is only applicable to spoke SDPs configured under the Epipe, IES and VPRN interface context.

The **no** form of this command means that the router T-PE either acts as the active T-PE (if signaling active is configured) or automatically determines which router will initiate MS-PW signaling based on the prefix values configured in the SAII and TAII of the spoke SDP. If the SAII has the greater prefix value, then the router will initiate MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAII has the greater value prefix, then the router will assume that the far end T-PE will initiate MS-PW signaling and will wait for that label mapping message before responding with a T-LDP label mapping message for the MS-PW in the reverse direction.

Default

no auto-config

Platforms

7705 SAR Gen 2

5.101 auto-config-save

auto-config-save

Syntax

[**no**] **auto-config-save**

Context

[Tree] (config>system>management-interface>cli>md-cli auto-config-save)

Full Context

configure system management-interface cli md-cli auto-config-save

Description

This command enables the functionality to automatically write the running configuration to the saved configuration file as part of a successful MD-CLI commit operation.

The **no** form of this command disables this functionality.

Default

auto-config-save

Platforms

7705 SAR Gen 2

auto-config-save

Syntax

[**no**] **auto-config-save**

Context

[Tree] (config>system>netconf auto-config-save)

Full Context

configure system netconf auto-config-save

Description

This command enables the functionality to automatically write the running configuration to the saved configuration file as part of a successful NETCONF or pySROS commit operation.

The **no** form of this command disables this functionality.

Default

auto-config-save

Platforms

7705 SAR Gen 2

auto-config-save

Syntax

[**no**] **auto-config-save**

Context

[Tree] (config>system>grpc>gnmi auto-config-save)

Full Context

configure system grpc gnmi auto-config-save

Description

This command enables the functionality to automatically write the running configuration to the saved configuration file as part of a successful gNMI commit operation.

The **no** form of this command disables this functionality.

Default

auto-config-save

Platforms

7705 SAR Gen 2

5.102 auto-crl-update

auto-crl-update

Syntax

auto-crl-update [**create**]

no auto-crl-update

Context

[Tree] (config>system>security>pki>ca-prof auto-crl-update)

Full Context

configure system security pki ca-profile auto-crl-update

Description

This command creates an auto CRL update configuration context with the **create** parameter, or enters the auto-crl-update configuration context without the **create** parameter.

This mechanism auto downloads a CRL file from a list of configured HTTP URLs either periodically or before existing CRL expires. If the downloaded CRL is more recent than the existing one, then the existing one will be replaced.

Note:
The configured URL must point to a DER encoded CRL file.

Parameters

create

Creates an auto CRL update for the ca-profile.

Platforms

7705 SAR Gen 2

auto-crl-update

Syntax

[**no**] **auto-crl-update**

Context

[Tree] (debug>certificate auto-crl-update)

Full Context

debug certificate auto-crl-update

Description

This command enables trace for automated and manual CRL updates.

Platforms

7705 SAR Gen 2

5.103 auto-eap-method

auto-eap-method

Syntax

auto-eap-method {**psk** | **cert** | **psk-or-cert**}

Context

[Tree] (config>ipsec>ike-policy auto-eap-method)

Full Context

configure ipsec ike-policy auto-eap-method

Description

This command enables following behavior for IKEv2 remote-access tunnel when auth-method is configured as auto-eap-radius:

- If there is no AUTH payload in IKE_AUTH request, then system use EAP to authenticate client and also will own-auth-method to generate AUTH payload.
- If there is AUTH payload in IKE_AUTH request:
  - if auto-eap-method is psk, then system proceed as auth-method:psk-radius
  - if auto-eap-method is cert, then system proceed as auth-method:cert-radius
  - if auto-eap-method is psk-or-cert, then:
    - if the "Auth Method" field of AUTH payload is PSK, then system proceed as auth-method:psk-radius
    - if the "Auth Method" field of AUTH payload is RSA or DSS, then system proceed as auth-method:cert-radius
  - The system will use **auto-eap-own-method** to generate AUTH payload.

This command only applies when **auth-method** is configured as **auto-eap-radius**.

Default

auto-eap-method cert

Parameters

psk

Uses the pre-shared-key as the authentication method.

cert

Uses the certificate as the authentication method.

psk-or-cert

Uses either the pre-shared-key or certificate based on the "Auth Method" field of the received AUTH payload.

Platforms

7705 SAR Gen 2

5.104 auto-eap-own-method

auto-eap-own-method

Syntax

auto-eap-own-method {**psk** | **cert**}

Context

[Tree] (config>ipsec>ike-policy auto-eap-own-method)

Full Context

configure ipsec ike-policy auto-eap-own-method

Description

This command enables following behavior for IKEv2 remote-access tunnel when auth-method is configured as auto-eap-radius:

- If there is no AUTH payload in IKE_AUTH request, then system use EAP to authenticate client and also will own-auth-method to generate AUTH payload.
- If there is AUTH payload in IKE_AUTH request:
  - if auto-eap-method is psk, then system proceed as auth-method:psk-radius.
  - if auto-eap-method is cert, then system proceed as auth-method:cert-radius.
  - if auto-eap-method is psk-or-cert, then:
    - if the "Auth Method" field of AUTH payload is PSK, then system proceed as auth-method:psk-radius.
    - if the "Auth Method" field of AUTH payload is RSA or DSS, then system proceed as auth-method:cert-radius.
  - The system will use **auto-eap-own-method** to generate AUTH payload.

This command only applies when **auth-method** is configured as **auto-eap-radius**.

Default

auto-eap-own-method cert

Parameters

psk

Uses a pre-shared-key to generate AUTH payload.

cert

Uses a public/private key to generate AUTH payload.

Platforms

7705 SAR Gen 2

5.105 auto-edge

auto-edge

Syntax

[**no**] **auto-edge**

Context

[Tree] (config>service>template>vpls-sap-template>stp auto-edge)

[Tree] (config>service>vpls>sap>stp auto-edge)

[Tree] (config>service>vpls>spoke-sdp>stp auto-edge)

Full Context

configure service template vpls-sap-template stp auto-edge

configure service vpls sap stp auto-edge

configure service vpls spoke-sdp stp auto-edge

Description

This command configures automatic detection of the edge port characteristics of the SAP or spoke SDP.

If **auto-edge** is enabled, and STP concludes there is no bridge behind the spoke SDP, the OPER_EDGE variable is dynamically set to true. If **auto-edge** is enabled, and a BPDU is received, the OPER_EDGE variable is dynamically set to false.

The **no** form of this command disables automatic detection.

Default

auto-edge

Platforms

7705 SAR Gen 2

auto-edge

Syntax

[**no**] **auto-edge**

Context

[Tree] (config>service>pw-template>stp auto-edge)

Full Context

configure service pw-template stp auto-edge

Description

This command configures automatic detection of the edge port characteristics of the SAP or spoke SDP.

If **auto-edge** is enabled, and STP concludes there is no bridge behind the spoke SDP, the OPER_EDGE variable is dynamically set to true. If **auto-edge** is enabled, and a BPDU is received, the OPER_EDGE variable is dynamically set to false.

The **no** form of this command disables automatic detection.

Default

auto-edge

Platforms

7705 SAR Gen 2

5.106 auto-establish

auto-establish

Syntax

[**no**] **auto-establish**

Context

[Tree] (config>ipsec>trans-mode-prof>dyn auto-establish)

[Tree] (config>router>if>ipsec>ipsec-tunnel>dyn auto-establish)

[Tree] (config>service>vprn>if>sap>ipsec-tun>dyn auto-establish)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>dyn auto-establish)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>dyn auto-establish)

Full Context

configure ipsec ipsec-transport-mode-profile dynamic-keying auto-establish

configure router interface ipsec ipsec-tunnel dynamic-keying auto-establish

configure service vprn interface sap ipsec-tunnel dynamic-keying auto-establish

configure service vprn interface ipsec ipsec-tunnel dynamic-keying auto-establish

configure service ies interface ipsec ipsec-tunnel dynamic-keying auto-establish

Description

This command enables automatic attempts to establish a phase 1 exchange.

The system automatically establishes a phase 1 SA as soon as the tunnel is provisioned and enabled (**no shutdown**). This option should only be configured on one side of the tunnel.

Any associated static routes remains up as long as the tunnel is up, even though it may actually be operationally down according to the CLI.

The **no** form of this command disables the automatic attempts to establish a phase 1 exchange.

Default

no auto-establish

Platforms

7705 SAR Gen 2

5.107 auto-learn-mac-protect

auto-learn-mac-protect

Syntax

[**no**] **auto-learn-mac-protect**

Context

[Tree] (config>service>vpls>endpoint auto-learn-mac-protect)

[Tree] (config>service>pw-template>split-horizon-group auto-learn-mac-protect)

Full Context

configure service vpls endpoint auto-learn-mac-protect

configure service pw-template split-horizon-group auto-learn-mac-protect

Description

This command enables the automatic protection of source MAC addresses learned on the associated object. MAC protection is used in conjunction with the **restrict-protected-src**, **restrict-unprotected-dst**,

and **mac-protect** commands. When **auto-learn-mac-protect** command is applied or removed, the MAC addresses are cleared from the related object.

When the **auto-learn-mac-protect** is enabled on an SHG the action only applies to the associated SAPs (no action is taken by default for spoke SDPs in the SHG). To enable this function for spoke SDPs within a SHG, the **auto-learn-mac-protect** command must be enabled explicitly under the spoke SDP. If required, the **auto-learn-mac-protect** command can also be enabled explicitly under specific SAPs within the SHG.

The **no** form of the command reverts to the default.

Default

no auto-learn-mac-protect

Platforms

7705 SAR Gen 2

auto-learn-mac-protect

Syntax

auto-learn-mac-protect [**exclude-list** *name*]

no auto-learn-mac-protect

Context

[Tree] (config>service>vpls>mesh-sdp auto-learn-mac-protect)

[Tree] (config>service>vpls>spoke-sdp auto-learn-mac-protect)

[Tree] (config>service>vpls>sap auto-learn-mac-protect)

[Tree] (config>service>vpls>split-horizon-group auto-learn-mac-protect)

[Tree] (config>service>pw-template auto-learn-mac-protect)

Full Context

configure service vpls mesh-sdp auto-learn-mac-protect

configure service vpls spoke-sdp auto-learn-mac-protect

configure service vpls sap auto-learn-mac-protect

configure service vpls split-horizon-group auto-learn-mac-protect

configure service pw-template auto-learn-mac-protect

Description

This command specifies whether to enable automatic population of the MAC protect list with source MAC addresses learned on the associated object under which the command is configured.

When configured, dynamically learned MAC Source Addresses (SA) are protected only if they are learned on an object with ALMP configured and there is no exclude list associated to the same object or if there is an exclude list but the MAC does not match any entry.

The same list can be used in multiple objects of the same or different service. If the list is empty, ALMP does not exclude any learned MAC from protection on the object.

The **no** form of the command disables the automatic population of the MAC protect list.

Default

auto-learn-mac-protect

Parameters

name

Specifies the name of the exclude list, up to 32 characters.

Platforms

7705 SAR Gen 2

5.108 auto-lsp

auto-lsp

Syntax

auto-lsp lsp-template *template-name* {**policy** *peer-prefix-policy* [**peer-prefix-policy**] | **one-hop**}

no auto-lsp lsp-template *template-name*

Context

[Tree] (config>router>mpls auto-lsp)

Full Context

configure router mpls auto-lsp

Description

This command enables the automatic creation of an RSVP point-to-point LSP to a destination node whose router ID matches a prefix in the specified peer prefix policy. This LSP type is referred to as auto-LSP of type mesh.

The user can associate multiple templates with same or different peer prefix policies. Each application of an LSP template with a given prefix in the prefix list results in the instantiation of a single CSPF computed LSP primary path using the LSP template parameters as long as the prefix corresponds to a router ID for a node in the TE database. This command does not support the automatic signaling of a secondary path for an LSP. If the signaling of multiple LSPs to the same destination node is required, the user must apply a separate LSP template to the same or different prefix list that contains the same destination node. Each instantiated LSP will have a unique LSP ID and a unique tunnel ID. This command also does not support the signaling of a non-CSPF LSP. The selection of the **no cspf** option in the LSP template is blocked.

Up to five peer prefix policies can be associated with a given LSP template at all times. Each time the user runs the **auto-lsp** command with the same or different prefix policy associations, or the user changes a

prefix policy associated with an LSP template, the system re-evaluates the prefix policy. The outcome of the re-evaluation tells MPLS if an existing LSP needs to be torn down or if a new LSP needs to be signaled to a destination address that is already in the TE database.

If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with an LSP template, the preceding prefix policy re-evaluation is performed.

The user must perform a **no shutdown** of the template before the template takes effect. After a template is in use, the user must shut down the template before effecting any changes to the parameters, except for those LSP parameters for which the change can be handled with the Make-Before-Break (MBB) procedures. These parameters are **bandwidth** and enabling **fast-reroute** with or without the **hop-limit** or **node-protect** options. For all other parameters, the user must shut down the template, makes the change, and perform a **no shutdown**. This results in the existing instances of the LSP using this template to be torn down and re-signaled.

When a router with a router ID that matches a prefix in the prefix list appears in the TE database, it is a trigger to signal the LSP. The signaled LSP is installed in the Tunnel Table Manager (TTM) and is available to applications such as LDP-over-RSVP, resolution of BGP label routes, resolution of BGP, IGP, and static routes. It is, however, not available for use as a provisioned SDP for explicit binding or auto-binding by services.

Except for the MBB limitations to the configuration parameter change in the LSP template, MBB procedures for manual and timer based re-signaling of the LSP, for TE Graceful Shutdown and for soft preemption are supported.

The **one-to-one** option under **fast-reroute**, the LSP Diff-Serv **class-type** and **backup-class-type** parameters are not supported. If **diffserv-te** is enabled under RSVP, the auto-created LSP is still signaled but with the default LSP class type.

If the **one-hop** option is specified instead of a prefix list, this command enables the automatic signaling of one-hop point-to-point LSPs using the specified template to all directly connected neighbors. This LSP type is referred to as auto-LSP of type one-hop. Although the provisioning model and CLI syntax differ from that of a mesh LSP only by the absence of a prefix list, the actual behavior is quite different. When this command is executed, the TE database keeps track of each TE link that comes up to a directly connected IGP neighbor whose router ID is discovered. It then instructs MPLS to signals an LSP with a destination address matching the router ID of the neighbor and with a strict hop consisting of the address of the interface used by the TE link. Thus, the **auto-lsp** command with the **one-hop** option results in one or more LSPs signaled to the neighboring router.

An auto-created mesh or one-hop LSP can collect egress statistics at the ingress LER by adding the **egress-statistics** node configuration into the LSP template. The user can also collect ingress statistics at the egress LER by using the same **ingress-statistics** node configuration. The user must specify the full LSP name as signaled by the ingress LER in the RSVP session name field of the Session Attribute object in the received Path message.

This feature also provides for the auto-creation of an SR-TE mesh LSP and for an SR-TE one-hop LSP.

The SR-TE mesh LSP feature specifically binds a **mesh-p2p-srte** LSP template with one or more prefix lists. When the TE database discovers a router that has a router ID matching an entry in the prefix list, it triggers MPLS to instantiate an SR-TE LSP to that router using the LSP parameters in the LSP template.

The SR-TE one-hop LSP feature specifically activates a **one-hop-p2p-srte** LSP template. In this case, the TE database keeps track of each TE link that comes up to a directly connected IGP neighbor. It then instructs MPLS to instantiate a SR-TE LSP with the following parameters:

- the source address of the local router
- an outgoing interface matching the interface index of the TE-link

- a destination address matching the router ID of the neighbor on the TE link

In both types of SR-TE auto-LSP, the router's hop-to-label translation computes the label stack required to instantiate the LSP.

Note:
An SR-TE auto-LSP can be reported to a PCE but cannot be delegated or have its paths computed by PCE.

The **no** form of this command deletes all LSPs signaled using the specified template and prefix policy. When the **one-hop** option is used, it deletes all one-hop LSPs signaled using the specified template to all directly-connected neighbors.

Parameters

lsp-template *template-name*

Specifies an LSP template name, up to 32 characters in length.

policy *peer-prefix-policy*

Specifies an peer prefix policy name, up to 32 characters in length.

one-hop

Enables the automatic signaling of one-hop point-to-point LSPs.

Platforms

7705 SAR Gen 2

5.109 auto-rp

auto-rp

Syntax

auto-rp [**detail**]

no auto-rp

Context

[Tree] (debug>router>pim auto-rp)

Full Context

debug router pim auto-rp

Description

This command enables debugging for PIM auto-RP.

The **no** form of this command disables PIM auto-RP debugging.

Parameters

detail

Debugs detailed information on the PIM auto-RP mechanism.

Platforms

7705 SAR Gen 2

5.110 auto-rp-discovery

auto-rp-discovery

Syntax

auto-rp-discovery [**candidate**] [**mapping-agent**]

no auto-rp-discovery

Context

[Tree] (config>service>vprn>pim>rp auto-rp-discovery)

Full Context

configure service vprn pim rp auto-rp-discovery

Description

This command enables the auto-RP protocol in discovery mode. In discovery mode, RP-mapping and RP-candidate messages are received and forwarded to downstream nodes. RP-mapping messages are received locally to learn the availability of RP nodes present in the network. In a VPRN configuration, Nokia recommends that a local loopback interface should be created with the same IP address as the system IP address.

The following configuration guidelines apply.

- Either **bsr-candidate** for IPv4 or **auto-rp-discovery** can be configured; the two mechanisms cannot be enabled together.
- **bsr-candidate** for IPv6 and **auto-rp-discovery** for IPv4 can be enabled together.
- **auto-rp-discovery** cannot be enabled together with **mdt-type sender-only** or **mdt-type receiver-only**, or **wildcard-spmsi** configurations.

This command also enables the auto-RP listener functionality. The auto-RP listener forwards the candidate 224.0.1.39 and mapping 224.0.1.40 messages over the PIM interfaces.

The **no** form of this command disables auto-RP discovery, auto-RP listener, candidate, and mapping-agent.

Default

no auto-rp-discovery

Parameters

candidate

Specifies that the RP is a candidate RP. The auto-RP C-RP announces the candidate RP messages on the 224.0.1.39 multicast address. This functionality is in addition to the listener functionality enabled by the auto RP discovery.

The default value is **no candidate**.

mapping agent

Specifies the mapping agent on the node. The auto-RP MA observes the **auto-rp-announcement** messages, selects the RP, and generates the RP discovery 224.0.1.40 messages. This functionality is in addition to the auto RP discovery functionality.

The default value is **no mapping-agent**.

Platforms

7705 SAR Gen 2

auto-rp-discovery

Syntax

auto-rp-discovery [**candidate**] [**mapping-agent**]

no auto-rp-discovery

Context

[Tree] (config>router>pim>rp auto-rp-discovery)

Full Context

configure router pim rp auto-rp-discovery

Description

This command enables the auto-RP protocol in discovery mode. In discovery mode, RP-mapping and RP candidate messages are received and forwarded to downstream nodes. RP-mapping messages are received locally to learn the availability of RP nodes present in the network.

The following configuration guidelines apply.

- Either **bsr-candidate** for IPv4 or **auto-rp-discovery** can be configured; the two mechanisms cannot be enabled together.
- **bsr-candidate** for IPv6 and **auto-rp-discovery** for IPv4 can be enabled together.

This command also enables the auto-RP listener functionality. The auto-RP listener forwards the candidate 224.0.1.39 and mapping 224.0.1.40 messages over the PIM interfaces.

The **no** form of this command disables auto-RP discovery, auto-RP listener, candidate, and mapping-agent.

Default

no auto-rp-discovery

Parameters

candidate

Specifies that the RP is a candidate RP. The auto-RP C-RP announces the candidate RP messages on the 224.0.1.39 multicast address. This functionality is in addition to the listener functionality enabled by the auto RP discovery.

The default value is **no candidate**.

mapping agent

Specifies the mapping agent on the node. The auto-RP MA observes the **auto-rp-announcement** messages, selects the RP, and generates the RP discovery 224.0.1.40 messages. This functionality is in addition to the auto RP discovery functionality.

The default value is **no mapping-agent**.

Platforms

7705 SAR Gen 2

5.111 auto-rx

auto-rx

Syntax

auto-rx

Context

[Tree] (config>router>ldp>targeted-session auto-rx)

Full Context

configure router ldp targeted-session auto-rx

Description

Commands in this context configure an automatic targeted LDP session and accept targeted Hello messages from any peer.

Platforms

7705 SAR Gen 2

5.112 auto-tx

auto-tx

Syntax

auto-tx

Context

[Tree] (config>router>ldp>targeted-session auto-tx)

Full Context

configure router ldp targeted-session auto-tx

Description

Commands in this context configure an automatic targeted LDP session and send targeted Hello messages towards PQ nodes determined by the rLFA algorithm.

Platforms

7705 SAR Gen 2

5.113 autoconfigure

autoconfigure

Syntax

autoconfigure

Context

[Tree] (bof autoconfigure)

Full Context

bof autoconfigure

Description

Commands in this context autoconfigure the IP address for the BOF. The IPv4 DHCP client, IPv6 DHCP client, and NDP/RA can be configured on the management interface.

Default

no autoconfigure

Platforms

7705 SAR Gen 2

5.114 autonegotiate

autonegotiate

Syntax

autonegotiate [**limited**]

no autonegotiate

Context

[Tree] (config>port>ethernet autonegotiate)

Full Context

configure port ethernet autonegotiate

Description

This command enables speed and duplex autonegotiation on Fast Ethernet ports and enables far-end fault indicator support on Gb ports.

There are three possible settings for autonegotiation:

- "on" or enabled with full port capabilities advertised
- "off" or disabled where there are no autonegotiation advertisements
- "limited" where a single speed/duplex is advertised.

When autonegotiation is enabled on a port, the link attempts to automatically negotiate the link speed and duplex parameters. If autonegotiation is enabled, the configured duplex and speed parameters are ignored.

When autonegotiation is disabled on a port, the port does not attempt to autonegotiate and will only operate at the **speed** and **duplex** settings configured for the port. Note that disabling autonegotiation on Gb ports is not allowed as the IEEE 802.3 specification for Gb Ethernet requires autonegotiation be enabled for far end fault indication.

If the **autonegotiate limited** keyword option is specified the port will auto-negotiate but will only advertise a specific speed and duplex. The speed and duplex advertised are the **speed** and **duplex** settings configured for the port. One use for limited mode is for multi-speed Gb ports to force Gb operation while keeping autonegotiation enabled for compliance with IEEE 802.3.

Router requires that autonegotiation be disabled or limited for ports in a Link Aggregation Group to guarantee a specific port speed.

The **no** form of this command disables autonegotiation on this port.

Default

autonegotiate

Parameters

limited

The Ethernet interface will automatically negotiate link parameters with the far end, but will only advertise the speed and duplex mode specified by the Ethernet **config>port>ethernet speed** and **config>port>ethernet duplex** commands.

Platforms

7705 SAR Gen 2

autonegotiate

Syntax

[**no**] **autonegotiate**

Context

[Tree] (bof autonegotiate)

Full Context

bof autonegotiate

Description

This command enables speed and duplex autonegotiation on the management Ethernet port in the running configuration and the Boot Option File (BOF).

When **autonegotiation** is enabled, the link attempts to automatically negotiate the link speed and duplex parameters. If **autonegotiation** is enabled, then the configured duplex and speed parameters are ignored.

The **no** form of this command disables the autonegotiate feature on this port.

Platforms

7705 SAR Gen 2

5.115 autonomous

autonomous

Syntax

[**no**] **autonomous**

Context

[Tree] (config>service>vprn>router-advert>if>prefix autonomous)

Full Context
configure service vprn router-advertisement interface prefix autonomous

Description
This command specifies whether the prefix can be used for stateless address autoconfiguration.

Default
autonomous

Platforms
7705 SAR Gen 2

autonomous

Syntax
[**no**] **autonomous**

Context
[Tree] (config>router>router-advert>if>prefix autonomous)

Full Context
configure router router-advertisement interface prefix autonomous

Description
This command specifies whether the prefix can be used for stateless address autoconfiguration.

Default
autonomous

Platforms
7705 SAR Gen 2

5.116 autonomous-system

autonomous-system

Syntax
autonomous-system *as-number*
no autonomous-system

Context

[Tree] (config>service>vprn autonomous-system)

Full Context

configure service vprn autonomous-system

Description

This command defines the autonomous system (AS) to be used by this VPN routing/forwarding (VRF). This command defines the autonomous system to be used by this VPN routing

The **no** form of this command removes the defined AS from this VPRN context.

Default

no autonomous-system

Parameters

as-number

Specifies the AS number for the VPRN service.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

autonomous-system

Syntax

autonomous-system *autonomous-system*

no autonomous-system

Context

[Tree] (config>router autonomous-system)

Full Context

configure router autonomous-system

Description

This command configures the autonomous system (AS) number for the router. A router can only belong to one AS. An AS number is a globally unique number with an AS. This number is used to exchange exterior routing information with neighboring ASs and as an identifier of the AS itself.

If the AS number is changed on a router with an active BGP instance, the new AS number is not used until the BGP instance is restarted either by administratively disabling/enabling (**shutdown**/**no shutdown**) the BGP instance or rebooting the system with the new configuration.

Default

no autonomous-system

Parameters

autonomous-system

Specifies the autonomous system number expressed as a decimal integer.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

5.117 avg-flr-event

avg-flr-event

Syntax

avg-flr-event {**forward** | **backward**} **threshold** *raise-threshold-percentage* [**clear** *clear-threshold-percentage*]

no avg-flr-event {**forward** | **backward**}

Context

[Tree] (config>oam-pm>session>ip>twamp-light>loss-events avg-flr-event)

Full Context

configure oam-pm session ip twamp-light loss-events avg-flr-event

Description

This command sets the frame loss ratio threshold configuration to be applied and checked at the end of the measurement interval for the specified direction. This is a percentage based on average frame loss ratio over the entire measurement interval. If the *clear-threshold-percent* value is not specified, the traffic crossing alarm is stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and without regard to any previous window. Each unique event can only be raised once within measurement interval. If the optional *clear-threshold-percent* value is specified, the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to following measurement intervals. If a threshold crossing event is raised another is not raised until a measurement interval completes and the clear threshold has not been exceeded. A clear event is raised under that condition.

The **no** form of this command removes the event threshold for frame loss ratio. The direction must be included with the **no** command.

Default

no avg-flr-event forward

no avg-flr-event backward

Parameters

forward

Specifies the threshold is applied to the forward direction value.

backward

Specifies the threshold is applied to the backward direction value.

raise-threshold-percentage

Specifies the rising percentage that determines when the event is to be generated.

Values 0.001 to 100.000

clear-threshold-percentage

Specifies an optional value used for stateful behavior that allows the operator to configure a percentage of loss value lower than the rising percentage to indicate when the clear event should be generated.

Values 0.000 to 99.999

A value 0.000 means that the FLR must be 0.000.

Platforms

7705 SAR Gen 2

5.118 avg-frame-overhead

avg-frame-overhead

Syntax

avg-frame-overhead *percent*

no avg-frame-overhead

Context

[Tree] (config>qos>sap-egress>queue avg-frame-overhead)

[Tree] (config>qos>network-queue>queue avg-frame-overhead)

Full Context

configure qos sap-egress queue avg-frame-overhead

configure qos network-queue queue avg-frame-overhead

Description

This command configures the average frame overhead, expressed as a percentage, at which the offered load expands on the physical medium (wire) at egress. This is important for accurate "on-the-wire"

rate calculations at various levels of H-QoS that do not inherently account for the physical medium characteristics. For example, without considering this overhead, a port scheduler in H-QoS might inaccurately estimate the available bandwidth on the wire, potentially leading to congestion issues and unexpected packet loss.

The rates impacted by the average frame overhead encompass the rates set on port schedulers and aggregate rate limits for subscribers and Vports. This impact is evident in the configured values, which represent on-the-wire (OTW) rates. Queue-configured rates, however, remain unaffected by this adjustment and continue to reflect Layer 2 rates.

This average frame overhead should be configured in networks with physical mediums that have constant sizes of transmission units (packets or cells) or in scenarios where the average packet size is known.

For Ethernet ports, the effect of this command depends on the setting of the **avg-frame-overhead-mode** command in advanced QoS configuration policy associated with the queue. If the **avg-frame-overhead-mode** is set to **auto**, the packet encapsulation overhead calculation is based on a fixed 20 bytes (7 bytes for preamble, 1 byte for start of frame delimiter, and 12 bytes for Inter-Frame Gap (IFG)) that the Ethernet medium adds to every packet during transmission. In other words, the configured rates for port-scheduler and aggregate rate limits for subscribers and Vports represent OTW rates.

The average frame overhead only affects rate and weight calculations and does not impact collected statistics for accounting purposes.

The **no** form of this command disables the average frame overhead.

Default

no avg-frame-overhead

Parameters

percent

Specifies the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0.00 to 100.00

Platforms

7705 SAR Gen 2

avg-frame-overhead

Syntax

avg-frame-overhead *percent*

no avg-frame-overhead

Context

[Tree] (config>service>ies>if>sap>egress>queue-override>queue avg-frame-overhead)

[Tree] (config>service>vpls>sap>egress>queue-override>queue avg-frame-overhead)

[Tree] (config>service>vprn>if>sap>egress>queue-override>queue avg-frame-overhead)

[Tree] (config>service>epipe>sap>egress>queue-override>queue avg-frame-overhead)

Full Context

configure service ies interface sap egress queue-override queue avg-frame-overhead

configure service vpls sap egress queue-override queue avg-frame-overhead

configure service vprn interface sap egress queue-override queue avg-frame-overhead

configure service epipe sap egress queue-override queue avg-frame-overhead

Description

This command configures overrides that supersede the average frame overhead configuration under the queue.

For a full description of this command, see the command description under the following contexts:

configure qos network-queue queue avg-frame-overhead

configure qos sap-egress queue avg-frame-overhead

The **no** form of this command disables overrides for the queue.

Default

no avg-frame-overhead

Parameters

percent

Specifies the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0.00 to 100.00, default

Platforms

7705 SAR Gen 2

6 b Commands

6.1 back

back

Syntax

back

Context

[Tree] (back)

Full Context

back

Description

This command moves the context back one level of the command hierarchy. For example, if the current level is the **config router ospf** context, the **back** command moves the cursor to the **config router** context level.

Platforms

7705 SAR Gen 2

6.2 backup

backup

Syntax

[**no**] **backup** *ip-address*

Context

[Tree] (config>service>ies>if>ipv6>vrrp backup)

Full Context

configure service ies interface ipv6 vrrp backup

Description
This command configures virtual router IP addresses for the interface.

Platforms
7705 SAR Gen 2

backup

Syntax
[**no**] **backup** *ip-address*

Context
[Tree] (config>service>ies>if>vrrp backup)

Full Context
configure service ies interface vrrp backup

Description
This command configures virtual router IP addresses for the interface.

Platforms
7705 SAR Gen 2

backup

Syntax
[**no**] **backup** *ip-address*

Context
[Tree] (config>service>vprn>if>vrrp backup)

Full Context
configure service vprn interface vrrp backup

Description
This command configures virtual router IP addresses for the interface.

Platforms
7705 SAR Gen 2

backup

Syntax

[**no**] **backup** *ipv6-address*

Context

[Tree] (config>service>vprn>if>ipv6>vrrp backup)

Full Context

configure service vprn interface ipv6 vrrp backup

Description

This command configures virtual router IP addresses for the interface.

Platforms

7705 SAR Gen 2

backup

Syntax

backup *mda-id*

no backup

Context

[Tree] (config>isa>tunnel-grp backup)

Full Context

configure isa tunnel-group backup

Description

This command assigns a tunnel ISA configured in the specified slot to this IPsec group. The backup module provides the IPsec group with warm redundancy when the primary module in the group is configured. An IPsec group must always have a primary configured.

Primary and backup modules have equal operational status and when both modules are coming up, the one that becomes operational first becomes the active module. An IPsec module can serve as a backup for multiple IPsec groups but the backup can become active for only one ISA IPsec group at a time.

All configuration information is pushed down to the backup MDA from the CPM once the CPM gets notice that the primary module has gone down. This allows multiple IPsec groups to use the same backup module. Any statistics not yet spooled are lost. Auto-switching from the backup to primary, after the primary becomes available again, is supported.

The user is notified through SNMP events when:

- When the ISA IPsec service goes down (all modules in the group are down) or comes back up (a module in the group becomes active).
- When ISA IPsec redundancy fails (one of the modules in the group is down) or recovers (the failed module comes back up).
- When an ISA IPsec activity switch took place.

The **no** form of this command removes the specified module from the IPsec group.

Default

no backup

Parameters

mda-id

Specifies the card/slot identifying a provisioned module to be used as a backup module.

Values mda-id: *slot*/*mda* slot 1 to up to 10 depending on chassis model mda 1 to 2

Platforms

7705 SAR Gen 2

backup

Syntax

[**no**] **backup** *ip-address*

Context

[Tree] (config>router>if>vrrp backup)

Full Context

configure router interface vrrp backup

Description

This command associates router IP addresses with the parental IP interface IP addresses.

The **backup** command has two distinct functions when used in an **owner** or a **non-owner** context of the virtual router instance.

Non-owner virtual router instances actually create a routable IP interface address that is operationally dependent on the virtual router instance mode (master or backup). The **backup** command in **owner** virtual router instances does not create a routable IP interface address; it simply defines the existing parental IP interface IP addresses that are advertised by the virtual router instance.

For **owner** virtual router instances, the **backup** command defines the IP addresses that are advertised within VRRP advertisement messages. This communicates the IP addresses that the master is representing to backup virtual routers receiving the messages. Advertising a correct list is important. The

specified *ip-address* must be equal to one of the existing parental IP interface IP addresses (primary or secondary) or the **backup** command fails.

For non-owner virtual router instances, the **backup** command actually creates an IP interface IP address used for routing IP packets and communicating with the system when the access commands are defined (**ntp-reply, ping-reply**, **telnet-reply**, and **ssh-reply**). The specified *ip-address* must be an IP address that is within one of the parental IP interface local subnets created with the **address** or **secondary** commands. If a local subnet does not exist that includes the specified *ip-address* or if *ip-address* is the same IP address as the parental IP interface IP address, the **backup** command fails.

The new interface IP address created with the **backup** command assumes the mask and parameters of the corresponding parent IP interface IP address. The *ip-address* is only active when the virtual router instance is operating in the master state. When not operating as master, the virtual router instance acts as if it is operationally down. It does not respond to ARP requests to *ip-addr*ess, nor does it route packets received with its *vrid* derived source MAC address. A non-master virtual router instance always silently discards packets destined to *ip-address*. A single virtual router instance may only have a single virtual router IP address from a given parental local subnet. Multiple virtual router instances can define a virtual router IP address from the same local subnet as long as each is a different IP address.

In IPv4, up to sixteen **backup** *ip-address* commands can be executed within the same virtual router instance. Executing **backup** multiple times with the same *ip-addr*ess results in no operation performed and no error generated. At least one successful **backup** *ip-address* command must be executed before the virtual router instance can enter the operational state.

When operating as (non-owner) master, the default functionality associated with *ip-address* is ARP response to ARP requests to *ip-address*, routing of packets destined to the virtual router instance source MAC address and silently discarding packets destined to *ip-address*. Enabling the non-owner-access parameters selectively allows ping, Telnet and SSH connectivity to *ip-address* when the virtual router instance is operating as master.

The **no** form of the command removes the specified virtual router IP address from the virtual router instance. For non-owner virtual router instances, this causes all routing and local access associated with the *ip-address* to cease. For **owner** virtual router instances, the **no backup** command only removes *ip-addr*ess from the list of advertised IP addresses. If the last *ip-address* is removed from the virtual router instance, the virtual router instance will enter the operationally down state

Default

no backup — No virtual router IP address is assigned.

Parameters

ip-address

The virtual router IP address expressed in dotted decimal notation. The IP virtual router IP address must be in the same subnet of the parental IP interface IP address or equal to one of the primary or secondary IP addresses for **owner** virtual router instances.

Values 1.0.0.1 to 223.255.255.254

Platforms

7705 SAR Gen 2

backup

Syntax

[**no**] **backup** *ipv6-address*

Context

[Tree] (config>router>if>ipv6>vrrp backup)

Full Context

configure router interface ipv6 vrrp backup

Description

This command associates router IPv6 addresses with the parental IP interface IP addresses.

The **backup** command has two distinct functions when used in an **owner** or a **non-owner** context of the virtual router instance.

Non-owner virtual router instances actually create a routable IP interface address that is operationally dependent on the virtual router instance mode (master or backup). The **backup** command in **owner** virtual router instances does not create a routable IP interface address; it simply defines the existing parental IP interface IP addresses that are advertised by the virtual router instance.

For **owner** virtual router instances, the **backup** command defines the IP addresses that are advertised within VRRP advertisement messages. This communicates the IP addresses that the master is representing to backup virtual routers receiving the messages. Advertising a correct list is important. The specified *ipv6-addr* must be equal to one of the existing parental IP interface IP addresses (link-local or global) or the **backup** command will fail.

For non-owner virtual router instances, the **backup** command actually creates an IP interface IP address used for routing IP packets and communicating with the system when the access commands are defined (**ntp-reply, ping-reply**, **telnet-reply**, and **ssh-reply**). The specified *ipv6-addr* must be an IP address that is within one of the parental IP interface local subnets created with the **link-local-address or address** commands. If a local subnet does not exist that includes the specified *ipv6-addr* or if *ipv6-addr* is the same IP address as the parental IP interface IP address, the **backup** command will fail.

The new interface IP address created with the **backup** command assumes the mask and parameters of the corresponding parent IP interface IP address. The *ipv6-addr* is only active when the virtual router instance is operating in the master state. For IPv6 VRRP, the parental interface's IP address that is in the same subnet as the backup address must be manually-configured, non EUI-64 and configured to be in the preferred state.

When not operating as master, the virtual router instance acts as if it is operationally down. It will not respond to Neighbor Solicitation (NS) requests to *ipv6-addr*, nor will it route packets received with its *vrid* derived source MAC address. A non-master virtual router instance always silently discards packets destined to *ipv6-addr*.

IPv6 allows the configuration of a link-local IPv6 address and multiple global IPv6 addresses on an interface. For each of these configured subnets, a virtual router IP address can be configured. Each IPv6 enabled device on a particular IPv6 subnet dynamically learns the connected IPv6 routers and correlated subnets in addition to the IPv6 default gateway using IPv6 neighbor discovery protocol (RFC 4861). This protocol behavior is revised from IPv4 where the default gateway is manually configured or derived from supporting protocols (for example, DHCP). During the IPv6 neighbor discovery process, VRRP enabled

routers will use backup IPv6 addresses and correlated derived virtual MAC addresses. Multiple virtual router instances can define a virtual router IP address from the same local subnet as long as each is a different IP address.

Executing **backup** multiple times with the same *ipv6-addr* results in no operation performed and no error generated. At least one successful **backup** *ipv6-addr* command must be executed before the virtual router instance can enter the operational state.

When operating as (non-owner) master, the default functionality associated with *ipv6-addr* results in the IPv6 Neighbor Advertisement response to IPv6 Neighbor Solicitation requests to *ip-addr*, routing of packets destined to the virtual router instance source MAC address, and silently discarding packets destined to *ipv6-addr*. An IPv6 virtual router instance can enter the operational state only if one of the configured backup addresses is a link-local address and the router advertisement of the interface is configured to use the virtual MAC address. Enabling the non-owner-access parameters selectively allows ping, Telnet, and traceroute connectivity to *ipv6-addr* when the virtual router instance is operating as master.

The **no** form of the command removes the specified virtual router IP address from the virtual router instance. For non-owner virtual router instances, this causes all routing and local access associated with the *ipv6-addr* to cease. For **owner** virtual router instances, the **no backup** command only removes *ipv6-addr* from the list of advertised IP addresses. If the last *ipv6-addr or the link-local address* is removed from the virtual router instance, the virtual router instance will enter the operationally down state

Default

no backup — No virtual router IP address is assigned.

Parameters

ipv6-address

The virtual router IP address is represented as eight groups of four hexadecimal digits, each group representing 16 bits (two octets). The groups are separated by colons (:). The IP virtual router IP address must be in the same subnet of the parental IP interface IP address or equal to one of the parent interface addresses for **owner** virtual router instances.

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x::d.d.d.d
	x: [0..FFFF]H
	d: [0..255]D

Platforms

7705 SAR Gen 2

6.3 backup-class-type

backup-class-type

Syntax

backup-class-type ct-number

no backup-class-type

Context

[Tree] (config>router>mpls>lsp>primary backup-class-type)

Full Context

configure router mpls lsp primary backup-class-type

Description

This command enables the use of the Diff-Serv backup Class-Type (CT), instead of the Diff-Serv main CT, to signal the LSP primary path when it fails and goes into retry. The Diff-Serv main CT is configured at the LSP level or at the primary path level using the following commands:

config>router>mpls>lsp>class-type *ct-number*

config>router>mpls>lsp>primary>class-type *ct-numbe*r

When an LSP primary path retries due a failure, for example, it fails after being in the UP state, or undergoes any type of Make-Before-Break (MBB), MPLS will retry a new path for the LSP using the main CT. If the first attempt failed, the head-end node performs subsequent retries using the backup CT. This procedure must be followed regardless if the currently used CT by this path is the main or backup CT. This applies to both CSPF and non-CSPF LSPs.

The triggers for using the backup CT after the first retry attempt are:

1. A local interface failure or a control plane failure (hello timeout and so on).
2. Receipt of a PathErr message with a notification of a FRR protection becoming active downstream and/or Receipt of a Resv message with a 'Local-Protection-In-Use' flag set. This invokes the FRR Global Revertive MBB.
3. Receipt of a PathErr message with error code=25 ("Notify") and sub-code=7 ("Local link maintenance required") or a sub-code=8 ("Local node maintenance required"). This invokes the TE Graceful Shutdown MBB.
4. Receipt of a Resv refresh message with the 'Preemption pending' flag set or a PathErr message with error code=34 ("Reroute") and a value=1 ("Reroute request soft preemption"). This invokes the soft preemption MBB.
5. Receipt of a ResvTear message.
6. A configuration change MBB.
7. The user executing the **clear**>**router**>**mpls**>**lsp** command.

When an unmapped LSP primary path goes into retry, it uses the main CT until the number of retries reaches the value of the new **main-ct-retry-limit** parameter. If the path did not come up, it must start using the backup CT at that point in time. By default, this parameter is set to infinite value. The new main-ct-retry-limit parameter has no effect on an LSP primary path which retries due to a failure event.

An unmapped LSP primary path is a path which has never received a Resv in response to the first Path message sent. This can occur when performing a 'shut/no-shut' on the LSP or LSP primary path or when the node reboots. An unmapped LSP primary path goes into retry if the retry timer expired or the head-end node received a PathErr message before the retry timer expired.

When the re-signal timer expires, CSPF will try to find a path with the main CT. The head-end node must re-signal the LSP even if the new path found by CSPF is identical to the existing one since the idea is to restore the main CT for the primary path. A path with main CT is not found, the LSP remains on its current primary path using the backup CT.

When the user performs a manual re-signal of the primary path, CSPF will try to find a path with the main CT. The head-end node must re-signal the LSP as in current implementation.

The **no** form of this command disables the use of the Diff-Serv backup CT.

Default

no backup-class-type

Parameters

ct-number

Specifies the Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.

Values 0 to 7, integer

Platforms

7705 SAR Gen 2

6.4 backup-next-hop

backup-next-hop

Syntax

[**no**] **backup-next-hop**

Context

[Tree] (config>router>mpls>fwd-policies>fwd-policy>nh-grp backup-next-hop)

Full Context

configure router mpls forwarding-policies forwarding-policy next-hop-group backup-next-hop

Description

Commands in this context configure the backup next hop of an NHG entry in a forwarding policy.

The **no** form of this command removes the backup next hop context from an NHG entry in a forwarding policy.

Platforms

7705 SAR Gen 2

6.5 backup-node-sid

backup-node-sid

Syntax

backup-node-sid *ip-prefix/prefix-length* **index** *index*

backup-node-sid *ip-prefix/prefix-length* **label** *label*

no backup-node-sid

Context

[Tree] (config>router>ospf>segm-rtng backup-node-sid)

Full Context

configure router ospf segment-routing backup-node-sid

Description

This command enables LFA Protection using segment routing backup node SID.

The objective of this feature is to reduce the label stack pushed in a LFA tunnel next hop of inter-area and inter-domain prefixes. This is applicable in MPLS deployments across multiple IGP areas or domains such in seamless MPLS design.

The user enables the feature by configuring a backup node SID at an ABR/ASBR that is acting as a backup to the primary exit ABR/ASBR of inter-area or inter-as routes learned as BGP labeled routes. The user can enter either a label or an index for the backup node SID.

When a node in a IGP domain resolves a BGP label route for an inter-area or inter-domain prefix via the primary ABR exit router, it will use the backup node SID of this router, which is advertised by the backup ABR/ABR, as the LFA backup instead of the SID to the remote LFA PQ node to save on the pushed label stack.

This feature only allows the configuration of a single backup node SID per IGP instance and per ABR/ASBR. In other words, only a pair of ABR/ASBR nodes can back up each other in an IGP domain. Each time the user invokes the above command within the same IGP instance, it will override any previous configuration of the backup node SID. The same ABR/ASBR can, however, participate in multiple IGP instances and provide backup support within each instance.

Default

no backup-node-sid

Parameters

ip-prefix/prefix-length

Specifies the IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.

Values ip-prefix/mask:

- ip-prefix a.b.c.d (host bits must be 0)

ipv6-prefix:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

prefix-length: 0 to 128

index

Specifies the index for this backup node SID.

Values 0 to 4294967295

label

Specifies the SID value for this backup node SID.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

6.6 backup-remote-ip

backup-remote-ip

Syntax

backup-remote-ip *ip-address*

no backup-remote-ip

Context

[Tree] (config>service>ies>if>sap>ip-tunnel backup-remote-ip)

[Tree] (config>service>vprn>if>sap>ip-tunnel backup-remote-ip)

Full Context

configure service ies interface sap ip-tunnel backup-remote-ip

configure service vprn interface sap ip-tunnel backup-remote-ip

Description

This command configures the alternate destination IPv4 or IPv6 address to use for an IP tunnel. This destination address is used only if the primary destination configured with the **remote-ip** command is unreachable in the delivery service. The **source** address, **remote-ip** address and **backup-remote-ip** address of a tunnel must all belong to the same address family (IPv4 or IPv6). When the backup-remote-ip address contains an IPv6 address it must be a global unicast address.

The **no** form of this command deletes the backup-destination address from the tunnel configuration.

Default

no backup-remote-ip

Parameters

ip-address

Specifies the destination IPv4 address or IPv6 address of the tunnel.

Values ipv4-address:

- a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

6.7 bandwidth

bandwidth

Syntax

bandwidth *bandwidth*

no bandwidth

Context

[Tree] (config>lag>access bandwidth)

[Tree] (config>port>ethernet>access bandwidth)

Full Context

configure lag access bandwidth

configure port ethernet access bandwidth

Description

This command configures the administrator bandwidth assigned and available to ports and LAGs for use by SAP bandwidth Connection Admission Control (CAC). The administrator bandwidth on a port or LAG can be overbooked or underbooked using the **booking-factor** command.

Port or LAG: Increasing the port or LAG admin bandwidth will increase the available admin bandwidth on that port or LAG. Reducing the port or LAG admin bandwidth will reduce the available admin bandwidth on that port or LAG, however, if the reduction of available admin bandwidth would cause it to be insufficient to cover the sum of the current SAP admin bandwidth on the port or LAG then the command will fail.

The **no** form of this command reverts to the default value.

Default

no bandwidth

Parameters

bandwidth

Specifies the administrator bandwidth, in kb/s, that is assigned to the port or LAG.

Values 1 to 6400000000

Platforms

7705 SAR Gen 2

bandwidth

Syntax

bandwidth *bandwidth*

no bandwidth

Context

[Tree] (config>service>epipe>sap bandwidth)

Full Context

configure service epipe sap bandwidth

Description

This command configures the administrator bandwidth assigned and available to SAPs for use by SAP bandwidth Connection Admission Control (CAC).

Attempts to increase the SAP administrator bandwidth fail if there is insufficient available administrator bandwidth on its port or LAG, otherwise the available port or LAG administrator bandwidth is reduced by the incremental SAP administrator bandwidth. Reducing the SAP administrator bandwidth increases the available administrator bandwidth on its port or LAG.

The **no** form of this command reverts to the default value.

Default

no bandwidth

Parameters

bandwidth

Specifies the administrator bandwidth, in kb/s, that is assigned to the SAP.

Values 1 to 6400000000

Platforms

7705 SAR Gen 2

bandwidth

Syntax

bandwidth *bw-value*

bandwidth max

no bandwidth

Context

[Tree] (config>service>epipe>spoke-sdp bandwidth)

Full Context

configure service epipe spoke-sdp bandwidth

Description

This command specifies the bandwidth to be used for VLL bandwidth accounting by the VLL CAC feature.

The service manager keeps track of the available bandwidth for each SDP. The maximum value is the sum of the bandwidths of all constituent LSPs in the SDP. The SDP available bandwidth is adjusted by the user configured booking factor.

If an LSP consists of a primary and many secondary standby LSPs, then the bandwidth used in the maximum SDP available bandwidth is that of the active path. Any change to and LSP active path bandwidth will update the maximum SDP available bandwidth. Note however that a change to any constituent LSP bandwidth due to re-signaling of the primary LSP path or the activation of a secondary path which causes overbooking of the maximum SDP available bandwidth causes a warning and a trap to be issued but no further action is taken. The activation of a bypass or detour LSP in the path of the primary LSP does not change the maximum SDP available bandwidth.

When the user binds a VLL service to this SDP, an amount of bandwidth equal to bandwidth is subtracted from the SDP available bandwidth adjusted by the booking factor. When the user deletes this VLL service binding from this SDP, an amount of bandwidth equal to bandwidth is added back into the SDP available bandwidth.

If the total SDP available bandwidth when adding this VLL service is about to overbook, a warning is issued and the binding is rejected. This means that the spoke SDP bandwidth does not update the maximum SDP available bandwidth. In this case, the spoke SDP is put in operational down state and a status message of "pseudowire not forwarding" is sent to the remote SR OS PE node. A trap is also generated. The service manager will not put the spoke SDP into an operationally up state until the user executes a **shutdown** command and then a **no-shutdown** command of the spoke SDP and the bandwidth check succeeds. Therefore, the service manager will not automatically audit spoke SDPs subsequently to their creation to check if bandwidth is available.

If the VLL service contains an endpoint with multiple redundant spoke SDPs, each spoke SDP will have its bandwidth checked against the available bandwidth of the corresponding SDP.

If the VLL service performs a pseudowire switching (VC switching) function, each spoke SDP is separately checked for bandwidth against the corresponding SDP.

This feature does not alter the way service packets are sprayed over multiple RSVP LSPs, which are part of the same SDP. That is, by default load balancing of service packets occurs over the SDP LSPs based on service-id, or based on a hash of the packet header if ingress SAP shared queuing is enabled. In both cases, the VLL bandwidth is not checked against the available bandwidth of the selected LSPs but on the total SDP available bandwidth. Therefore, if there is a single LSP per SDP, these two matches.

If class-forwarding is enabled on the SDP, VLL service packets are forwarded to the SDP LSP which the packet forwarding class maps to, or if this is down to the default LSP. However, the VLL bandwidth is not checked against the selected LSP available bandwidth but on the total SDP available bandwidth. If there is a single LSP per SDP, these two matches.

If a non-zero bandwidth is specified for a VLL service and attempts to bind the service to an LDP or a GRE SDP, a warning is issued that CAC failed but the VLL is established. A trap is also generated.

The **no** form of this command reverts to the default value.

Parameters

bw-value

The bandwidth to be used for VLL bandwidth accounting by the VLL CAC feature, in kilobits per second.

Values 0 to 100000000

Default 0

Platforms

7705 SAR Gen 2

bandwidth

Syntax

bandwidth *bandwidth*

no bandwidth

Context

[Tree] (config>service>vpls>sap bandwidth)

Full Context

configure service vpls sap bandwidth

Description

This command specifies the admin bandwidth assigned to SAPs, ports and LAGs which is used by SAP bandwidth CAC.

SAP: Attempts to increase the SAP admin bandwidth will fail if there is insufficient available admin bandwidth on its port or LAG, otherwise the port or LAG available admin bandwidth will be reduced by the incremental SAP admin bandwidth. Reducing the SAP admin bandwidth will increase the available admin bandwidth on its port or LAG. This is not supported for PW-SAPs, Ethernet tunnels or subscriber group interface SAPs.

The **no** version of the command reverts to the default value.

Default

no bandwidth

Parameters

bandwidth

Specifies the admin bandwidth assigned to the SAP, port or LAG, in kb/s.

Values 1 to 6400000000

Platforms

7705 SAR Gen 2

bandwidth

Syntax

bandwidth *bandwidth*

no bandwidth

Context

[Tree] (config>service>ies>if>sap bandwidth)

Full Context

configure service ies interface sap bandwidth

Description

This command specifies the admin bandwidth assigned to SAPs, ports and LAGs which is used by SAP bandwidth CAC.

SAP: Attempts to increase the SAP admin bandwidth will fail if there is insufficient available admin bandwidth on its port or LAG, otherwise the port or LAG available admin bandwidth will be reduced by the incremental SAP admin bandwidth. Reducing the SAP admin bandwidth will increase the available admin bandwidth on its port or LAG. This is not supported for PW-SAPs, Ethernet tunnels or subscriber group interface SAPs.

The **no** version of the command reverts to the default value.

Default

no bandwidth

Parameters

bandwidth

Specifies the admin bandwidth assigned to the SAP, port or LAG, in kb/s.

Values 1 to 6400000000

Platforms

7705 SAR Gen 2

bandwidth

Syntax

bandwidth *bandwidth*

no bandwidth

Context

[Tree] (config>service>vprn>if>sap bandwidth)

Full Context

configure service vprn interface sap bandwidth

Description

This command specifies the admin bandwidth assigned to SAPs, ports and LAGs which is used by SAP bandwidth CAC.

SAP: Attempts to increase the SAP admin bandwidth will fail if there is insufficient available admin bandwidth on its port or LAG, otherwise the port or LAG available admin bandwidth will be reduced by the incremental SAP admin bandwidth. Reducing the SAP admin bandwidth will increase the available admin bandwidth on its port or LAG. This is not supported for PW-SAPs, Ethernet tunnels or subscriber group interface SAPs.

The **no** version of the command reverts to the default value.

Default

no bandwidth

Parameters

bandwidth

Specifies the admin bandwidth assigned to the SAP, port or LAG, in kb/s.

Values 1 to 6400000000

Platforms

7705 SAR Gen 2

bandwidth

Syntax

bandwidth *bandwidth-in-mbps*

no bandwidth

Context

[Tree] (config>router>mpls>lsp-template bandwidth)

Full Context

configure router mpls lsp-template bandwidth

Description

This command specifies the amount of bandwidth to be reserved for the P2MP instance.

Parameters

bandwidth-in-mbps

Specifies the bandwidth, in Mb/s.

Values 0 to 6400000

Platforms

7705 SAR Gen 2

bandwidth

Syntax

bandwidth *bandwidth-in-mbps*

no bandwidth

Context

[Tree] (config>router>mpls>lsp>secondary bandwidth)

[Tree] (config>router>mpls>lsp>primary bandwidth)

Full Context

configure router mpls lsp secondary bandwidth

configure router mpls lsp primary bandwidth

Description

This command specifies the amount of bandwidth to be reserved for the LSP path.

The **no** form of this command resets bandwidth parameters (no bandwidth is reserved).

Default

no bandwidth (bandwidth setting in the global LSP configuration)

Parameters

bandwidth-in-mbps

Specifies the amount of bandwidth reserved for the LSP path in Mb/s.

Values 0 to 6400000

Platforms

7705 SAR Gen 2

6.8 base-op-authorization

base-op-authorization

Syntax

base-op-authorization

Context

[Tree] (config>system>security>profile>netconf base-op-authorization)

Full Context

configure system security profile netconf base-op-authorization

Description

Commands in this context configure the permission to use NETCONF operations at the base operation level for the specified profile.The NETCONF operations are authorized by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

6.9 begin

begin

Syntax

begin

Context

[Tree] (config>router>bfd begin)

Full Context

configure router bfd begin

Description

This command switches to edit mode for a BFD template. Changes are not activated until the **commit** command is issued for the BFD template changes.

Platforms

7705 SAR Gen 2

begin

Syntax

begin

Context

[Tree] (config>router>route-next-hop-policy begin)

Full Context

configure router route-next-hop-policy begin

Description

This command switches to edit mode for route next-hop templates. Changes are not activated until the **commit** command is issued for the route next-hop templates changes.

Default

begin

Platforms

7705 SAR Gen 2

begin

Syntax

begin {**exclusive**}

Context

[Tree] (config>router>policy-options begin)

Full Context

configure router policy-options begin

Description

This command is required in order to enter the mode to create or edit route policies.

Parameters

exclusive

Specifies an exclusive lock on the policy configuration. Other CLI and SNMP users will be unable to edit the policy configuration until the lock is removed (via commit, abort, a timeout occurring, or a forced override).

Platforms

7705 SAR Gen 2

6.10 begin-time

begin-time

Syntax

begin-time *date hours-minutes* [**UTC**]

begin-time {**now** | **forever**}

no begin-time

Context

[Tree] (config>system>security>keychain>direction>uni>send>entry begin-time)

[Tree] (config>system>security>keychain>direction>bi>entry begin-time)

[Tree] (config>system>security>keychain>direction>uni>receive>entry begin-time)

Full Context

configure system security keychain direction uni send entry begin-time

configure system security keychain direction bi entry begin-time

configure system security keychain direction uni receive entry begin-time

Description

This command specifies the calendar date and time after which the key specified by the keychain authentication key is used to sign and/or authenticate the protocol stream.

If no date and time is set, the begin-time is represented by a date and time string with all NULLs and the key is not valid by default.

Default

begin-time forever

Parameters

date hours-minutes

Specifies the date and time for the key to become active.

Values date: YYYY/MM/DD hours-minutes: hh:mm[:ss]

now

Specifies the key should become active immediately.

forever

Specifies that the key is always inactive.

UTC

Indicates that time is given with reference to Coordinated Universal Time in the input.

Platforms

7705 SAR Gen 2

6.11 best-path-selection

best-path-selection

Syntax

best-path-selection

Context

[Tree] (config>service>vprn>bgp best-path-selection)

Full Context

configure service vprn bgp best-path-selection

Description

This command enables path selection configuration.

Platforms

7705 SAR Gen 2

best-path-selection

Syntax

best-path-selection

Context

[Tree] (config>router>bgp best-path-selection)

Full Context

configure router bgp best-path-selection

Description

Commands in this context configure path selection parameters.

Platforms

7705 SAR Gen 2

6.12 bfd

bfd

Syntax

bfd *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier* [**echo-receive** *echo-interval*] [**type cpm-np**]

no bfd

Context

[Tree] (config>service>vprn>if bfd)

[Tree] (config>service>vprn>nw-if bfd)

[Tree] (config>service>vprn>if>ipv6 bfd)

[Tree] (config>service>ies>if bfd)

[Tree] (config>service>ies>if>ipv6 bfd)

Full Context

configure service vprn interface bfd

configure service vprn network-interface bfd

configure service vprn interface ipv6 bfd

configure service ies interface bfd

configure service ies interface ipv6 bfd

Description

This command specifies the BFD parameters for the associated IP interface. If no parameters are defined the default value are used.

The multiplier specifies the number of consecutive BFD messages that must be missed from the peer before the BFD session state is changed to down and the upper level protocols (OSPF, IS-IS, BGP or PIM) is notified of the fault.

The **no** form of this command removes BFD from the interface.

Note:
The *transmit-interval,* **receive** *receive-interval, and* **echo-receive** *echo-interval* values can only be modified to a value less than 100 when:

1. The **type cpm-np option** is explicitly configured.
2. The service is shut down (**shutdown**)
3. The interval is specified 10 to 100000.
4. The service is re-enabled (**no shutdown**)

To remove the **type cpm-np** option, re-issue the **bfd** command without specifying the **type** parameter.

Parameters

transmit-interval

Sets the transmit interval for the BFD session.

Values 100 to 100000
10 to 100000

Default 100

receive ***receive-interval***

Sets the receive interval for the BFD session.

Values 100 to 100000
10 to 100000

Default 100

multiplier ***multiplier***

Sets the multiplier for the BFD session.

Values 3 to 20

Default 3

echo-receive ***echo-interval***

Sets the minimum echo receive interval, in milliseconds, for the BFD session.

Values 100 to 100000

10 to 100000

Default 100

type cpm-np

Specifies that BFD sessions associated with this interface is created on the CPM network processor to allow for fast timers down to 10 ms granularity.

Platforms

7705 SAR Gen 2

bfd

Syntax

bfd *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier*] [**echo-receive** *echo-interval*] [**type cpm-np**]

no bfd

Context

[Tree] (config>router>if>ipv6 bfd)

[Tree] (config>router>if bfd)

Full Context

configure router interface ipv6 bfd

configure router interface bfd

Description

This command specifies the bidirectional forwarding detection (BFD) parameters for the associated IP interface. If no parameters are defined the default values are used.

The multiplier specifies the number of consecutive BFD messages that must be missed from the peer before the BFD session state is changed to down and the upper level protocols (OSPF, IS-IS, BGP or PIM) is notified of the fault.

The **no** form of this command removes BFD from the router interface regardless of the IGP/RSVP.

Important notes: The *transmit-interval* and **receive receive-interval** values can only be modified to a value less than 100 ms when:

1. The **type cpm-np option** is explicitly configured.
2. The service is shut down (**shutdown**)
3. The interval is specified 10 to 100000.
4. The service is re-enabled (**no shutdown**)

To remove the **type cpm-np** option, re-issue the **bfd** command without specifying the **type** parameter.

Default

no bfd

Parameters

transmit-interval

Sets the transmit interval, in milliseconds, for the BFD session.

Values 10 to 100000 (see Important Notes above)

Default 100

receive-interval

Sets the receive interval, in milliseconds, for the BFD session.

Values 10 to 100000 (see Important Notes above)

Default 100

multiplier

Sets the multiplier for the BFD session. A multiplier of less than 3 should not be used in production environments.

Values 1 to 20

Default 3

echo-interval

Sets the minimum echo receive interval, in milliseconds, for the session.

Values 100 to 100000

Default 0

cpm-np

Selects the CPM network processor type as the local termination point for the BFD session. See Important Notes, above.

Platforms

7705 SAR Gen 2

6.13 bfd-designate

bfd-designate

Syntax

[**no**] **bfd-designate**

Context

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel bfd-designate)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel bfd-designate)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel bfd-designate)

[Tree] (config>router>if>ipsec>ipsec-tunnel bfd-designate)

Full Context

configure service vprn interface ipsec ipsec-tunnel bfd-designate

configure service vprn interface sap ipsec-tunnel bfd-designate

configure service ies interface ipsec ipsec-tunnel bfd-designate

configure router interface ipsec ipsec-tunnel bfd-designate

Description

This command specifies whether this IPsec tunnel is the BFD designated tunnel.

Default

no bfd-designate

Platforms

7705 SAR Gen 2

6.14 bfd-enable

bfd-enable

Syntax

[**no**] **bfd-enable**

Context

[Tree] (config>router>bgp>group>neighbor bfd-enable)

[Tree] (config>redundancy>multi-chassis>peer>mc-ep bfd-enable)

[Tree] (config>router>bgp>group bfd-enable)

[Tree] (config>router>bgp bfd-enable)

Full Context

configure router bgp group neighbor bfd-enable

configure redundancy multi-chassis peer mc-endpoint bfd-enable

configure router bgp group bfd-enable

configure router bgp bfd-enable

Description

This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

The **no** form of this command disables BFD.

Default

no bfd-enable

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable** [*service-id*] **interface** *interface-name* **dst-ip** *ip-address*

[**no**] **bfd-enable interface** *interface-name* **dst-ip** *ip-address* **name** *name*

Context

[Tree] (config>service>ies>if>ipv6>vrrp bfd-enable)

[Tree] (config>service>ies>if>vrrp bfd-enable)

Full Context

configure service ies interface ipv6 vrrp bfd-enable

configure service ies interface vrrp bfd-enable

Description

This commands assigns a bi-directional forwarding (BFD) session providing heart-beat mechanism for the given VRRP/SRRP instance. There can be only one BFD session assigned to any given VRRP/SRRP instance, but there can be multiple SRRP/VRRP sessions using the same BFD session.

BFD control the state of the associated interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface. The specified interface may not be configured with BFD; however, when it is, the virtual router will then initiate the BFD session.

The **no** form of this command removes BFD from the configuration.

Parameters

service-id

Specifies the service ID of the interface running BFD.

Values service-id: 1 to 2147483648

No service ID indicates a network interface.

interface *interface-name*

Specifies the name of the interface running BFD.

dst-ip *ip-address*

Specifies the destination address to be used for the BFD session.

name *name*

Specifies the name, up to 64 characters.

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable**

Context

[Tree] (config>service>vprn>static-route-entry>next-hop bfd-enable)

Full Context

configure service vprn static-route-entry next-hop bfd-enable

Description

This command associates the static route state to a BFD session between the local system and the configured nexthop.

The remote end of the BFD session must also be configured to originate or accept the BFD session controlling the static route state.

The **no** form of this command removes the association of the static route state to that of the BFD session.

Default

no bfd-enable

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable interface** *interface-name* **dst-ip** *ip-address*

[**no**] **bfd-enable** *service-id* **interface** *interface-name* **dst-ip** *ip-address*

[**no**] **bfd-enable interface** *interface-name* **dst-ip** *ip-address* **name** *service-name*

Context

[Tree] (config>service>vprn>if>ipv6>vrrp bfd-enable)

[Tree] (config>service>vprn>if>vrrp bfd-enable)

Full Context

configure service vprn interface ipv6 vrrp bfd-enable

configure service vprn interface vrrp bfd-enable

Description

This commands assigns a bi-directional forwarding (BFD) session providing heart-beat mechanism for the given VRRP/SRRP instance. There can be only one BFD session assigned to any given VRRP/SRRP instance, but there can be multiple SRRP/VRRP sessions using the same BFD session. If the interface used is configured with centralized BFD, the BFD transmit and receive intervals need to be set to at least 300 ms.

BFD control the state of the associated interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface. The specified interface may not be configured with BFD; when it is, the virtual router will then initiate the BFD session.

The **no** form of this command removes BFD from the configuration.

Parameters

svc-id

Specifies the service ID of the interface running BFD. If no *svc-id* is specified then it indicates that the interface is a network interface in the Base router instance.

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **bfd-enable interface** *interface-name* **dst-ip** *ip-address* **name** *name* variant can be used in all configuration modes.

Values {*id* | *svc-name*}

id: 1 to 2147483647

svc-name: Specifies an existing service name up to 64 characters (*svc-name* is an alias for input only. The *svc-name* gets replaced with an id automatically by SR OS in the configuration)

interface *interface-name*

Specifies the name of the interface running BFD, up to 32 characters.

dst-ip *ip-address*

Specifies the destination address to be used for the BFD session.

name *name*

Specifies a service name, up to 64 characters.

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

bfd-enable {**ipv4** | **ipv6**} [**include-bfd-tlv**]

no bfd-enable {**ipv4** | **ipv6**}

Context

[Tree] (config>service>vprn>isis>if bfd-enable)

Full Context

configure service vprn isis interface bfd-enable

Description

This command enables the use of bi-directional forwarding (BFD) to control IPv4 or adjacencies. By enabling BFD on an IPv4 or IPv6 protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set by the BFD command under the IP interface. This command must be given separately to enable or disable BFD for IPv4 and IPv6.

The **no** form of this command removes BFD from the associated adjacency.

Default

no bfd-enable ipv4

no bfd-enable ipv6

Parameters

ipv4

Keyword to enable BFD to control IPv4 adjacencies.

ipv6

Keyword to enable BFD to control IPv6 adjacencies.

include-bfd-tlv

Enables support for the IS-IS BFD TLV options in accordance with RFC 6213, which specifies that a BFD session must be established before an IS-IS adjacency can transition to the established state. This option must be enabled on all IS-IS neighbors on a shared interface.

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

bfd-enable [**remain-down-on-failure**]

bfd-enable [**remain-down-on-failure**] **strict** [**strict-mode-holddown** *number*]

no bfd-enable

Context

[Tree] (config>service>vprn>ospf3>area>if bfd-enable)

[Tree] (config>service>vprn>ospf>area>if bfd-enable)

Full Context

configure service vprn ospf3 area interface bfd-enable

configure service vprn ospf area interface bfd-enable

Description

This command configures Bidirectional Forwarding Detection (BFD) to control the state of the associated protocol interface. By enabling BFD on a protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD session are set using the **bfd** command in the associated IP interface context.

The **no** form of this command removes BFD from the associated IGP protocol adjacency.

Default

no bfd-enable

Parameters

remain-down-on-failure

Keyword to force adjacency down on BFD failure.

strict

Keyword to specify that the system uses BFD strict-mode, which requires that an active BFD session exists between the OSPF neighbors before establishing a full adjacency. When this keyword is configured, the router uses Link-Local Signaling (LLS) with the B-flag set to instruct OSPF neighbors that BFD must be enabled on the link. BFD strict-mode requires that both sides have the B-flag set.

During OSPFv3 BFD strict-mode operations, the router advertises the Local Interface IPv4 Address TLV using LLS, but the SR OS router continues to use IPv6-based BFD sessions for both the IPv4 and IPv6 address families.

strict-mode-holddown ***number***

Specifies a delay in bringing up the OSPF adjacency after the BFD session is established. Holddown helps mitigate potential routing churn when BFD sessions are unstable. The holddown timer is reset on an adjacency when a BFD session operationally toggles.

Values 1 to 600

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable** [**ipv4** | **ipv6**]

Context

[Tree] (config>service>vprn>pim>if bfd-enable)

Full Context

configure service vprn pim interface bfd-enable

Description

This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

The **no** form of this command removes BFD from the associated IGP protocol adjacency.

Default

no bfd-enable

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable**

Context

[Tree] (config>service>vprn>bgp>group>neighbor bfd-enable)

[Tree] (config>service>vprn>bgp>group bfd-enable)

[Tree] (config>service>vprn>bgp bfd-enable)

Full Context

configure service vprn bgp group neighbor bfd-enable

configure service vprn bgp group bfd-enable

configure service vprn bgp bfd-enable

Description

This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. The parameters used for the BFD are set with the BFD command under the IP interface.

The **no** form of this command disables bfd-enable on the VPRN service.

Default

no bfd-enable

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable**

Context

[Tree] (config>service>vprn>ripng>group>neighbor bfd-enable)

[Tree] (config>service>vprn>ripng>group bfd-enable)

[Tree] (config>service>vprn>ripng bfd-enable)

[Tree] (config>service>vprn>rip bfd-enable)

[Tree] (config>service>vprn>rip>group>neighbor bfd-enable)

[Tree] (config>service>vprn>rip>group bfd-enable)

Full Context

configure service vprn ripng group neighbor bfd-enable

configure service vprn ripng group bfd-enable

configure service vprn ripng bfd-enable

configure service vprn rip bfd-enable

configure service vprn rip group neighbor bfd-enable

configure service vprn rip group bfd-enable

Description

This command enables bi-directional forwarding (BFD) to control the state of the associated protocol adjacency. By enabling BFD on a given protocol interface, the state of the RIP neighbor is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set using the **bfd** command under the IP interface configuration context.

The **no** form of this command removes BFD from the associated protocol adjacency.

Default

no bfd-enable

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

bfd-enable [**ipv4**][**ipv6**]

no bfd-enable

Context

[Tree] (config>router>ldp>if-params>if bfd-enable)

Full Context

configure router ldp interface-parameters interface bfd-enable

Description

This command enables tracking of the Hello adjacency to an LDP peer using BFD.

When this command is enabled on an LDP interface, LDP registers with BFD and starts tracking the LSR-id of all peers it formed Hello adjacencies with over that LDP interface. The LDP hello mechanism is used to determine the remote address to be used for the BFD session. The parameters used for the BFD session, that is, transmit-interval, receive-interval, and multiplier are those configured under the IP interface in existing implementation: **config>router>if>bfd**.

The operation of BFD over an LDP interface tracks the next-hop of the IPv4 and IPv6 prefixes in addition to tracking the LDP peer address of the Hello adjacency over that link. This is required since LDP can resolve

both IPv4 and IPv6 prefix FECs over a single IPv4 or IPv6 LDP session and as such the next-hop of a prefix will not necessarily match the LDP peer source address of the Hello adjacency.

The failure of either or both of the BFD session tracking the FEC next-hop and the one tracking the Hello adjacency will cause the LFA backup NHLFE for the FEC to be activated or the FEC to be re-resolved if there is no FRR backup.

When multiple links exist to the same LDP peer, a Hello adjacency is established over each link and a separate BFD session is enabled on each LDP interface. If a BFD session times out on a specific link, LDP will immediately associate the LDP session with one of the remaining Hello adjacencies and trigger the LDP FRR procedures. As soon as the last Hello adjacency goes down due to BFD timing out, the LDP session goes down and the LDP FRR procedures will be triggered.

The **no** form of this command disables BFD on the LDP interface.

Default

no bfd-enable

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable**

Context

[Tree] (config>router>ldp>targ-session>peer bfd-enable)

[Tree] (config>router>ldp>targ-session>peer-template bfd-enable)

Full Context

configure router ldp targeted-session peer bfd-enable

configure router ldp targeted-session peer-template bfd-enable

Description

This command enables the bidirectional forwarding detection (BFD) session for the selected TLDP session. By enabling BFD for a selected targeted session, the state of that session is tied to the state of the underneath BFD session between the two nodes.

The parameters used for the BFD are set via the BFD command under the IP interface.

The **no** form of this command removes the TLDP session operational state binding to the central BFD session one.

Default

no bfd-enable

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable**

Context

[Tree] (config>router>rsvp>interface bfd-enable)

Full Context

configure router rsvp interface bfd-enable

Description

This command enables the use of bi-directional forwarding (BFD) to control the state of the associated RSVP interface. This causes RSVP to register the interface with the BFD session on that interface.

The user configures the BFD session parameters, such as, **transmit-interval**, **receive-interval**, and **multiplier**, under the IP interface in the **config>router> if>bfd** context.

Note:
It is possible that the BFD session on the interface was started because of a prior registration with another protocol, for example, OSPF or IS-IS.

The registration of an RSVP interface with BFD is performed at the time of neighbor gets its first session. This means when this node sends or receives a new Path message over the interface. If however the session did not come up, due to not receiving a Resv for a new path message sent after the maximum number of re-tries, the LSP is shutdown and the node de-registers with BFD. In general, the registration of RSVP with BFD is removed as soon as the last RSVP session is cleared.

The registration of an RSVP interface with BFD is performed independent of whether RSVP hello is enabled on the interface or not. However, hello timeout will clear all sessions towards the neighbor and RSVP de-registers with BFD at clearing of the last session.

An RSVP session is associated with a neighbor based on the interface address the path message is sent to. If multiple interfaces exist to the same node, each interface is treated as a separate RSVP neighbor. The user will have to enable BFD on each interface and RSVP will register with the BFD session running with each of those neighbors independently

Similarly the disabling of BFD on the interface results in removing registration of the interface with BFD.

When a BFD session transitions to DOWN state, the following actions are triggered. For RSVP signaled LSPs, this triggers activation of FRR bypass/detour backup (PLR role), global revertive (head-end role), and switchover to secondary if any (head-end role) for affected LSPs with FRR enabled. It triggers switchover to secondary if any and scheduling of re-tries for signaling the primary path of the non-FRR affected LSPs (head-end role).

The **no** form of this command removes BFD from the associated RSVP protocol adjacency.

Default

no bfd-enable

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

bfd-enable service-name *service-name* **interface-name** *interface-name* **dst-ip** *ip-address*

no bfd-enable

Context

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel bfd-enable)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel bfd-enable)

[Tree] (config>router>if>ipsec>ipsec-tunnel bfd-enable)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel bfd-enable)

Full Context

configure service ies interface ipsec ipsec-tunnel bfd-enable

configure service vprn interface sap ipsec-tunnel bfd-enable

configure router interface ipsec ipsec-tunnel bfd-enable

configure service vprn interface ipsec ipsec-tunnel bfd-enable

Description

This command assigns a BFD session to provide a heart-beat mechanism for a given IPsec tunnel. There can be only one BFD session assigned to any given IPsec tunnel, but there can be multiple IPsec tunnels using same BFD session. BFD controls the state of the associated tunnel. If the BFD session goes down, the system will also bring down the associated non-designated IPsec tunnel.

Parameters

service-name

Specifies the service name, up to 64 characters, on which the BFD session resides.

interface-name

Specifies the name, up to 32 characters, of the interface used by the BFD session.

ip-address

Specifies the destination address to be used for the BFD session.

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable**

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ipsec bfd-enable)

Full Context

configure redundancy multi-chassis peer mc-ipsec bfd-enable

Description

This command enables tracking a central BFD session, if the BFD session goes down, then system consider the peer is down and change the mc-ipsec status of configured tunnel-group accordingly.

The BFD session uses specified the loopback interface (in the specified service) address as the source address and uses specified dst-ip as the destination address. Other BFD parameters are configured with the **bfd** command on the specified interface.

Default

no bfd-enable

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable** [**ipv4** | **ipv6**]

Context

[Tree] (config>router>pim>interface bfd-enable)

Full Context

configure router pim interface bfd-enable

Description

This command enables the use of IPv4 or IPv6 bidirectional forwarding detection (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

The **no** form of this command removes BFD from the associated IGP protocol adjacency.

Default

no bfd-enable

Parameters

ipv4

Enables the use of IPv4 BFD.

ipv6

Enables the use of IPv6 BFD.

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable**

Context

[Tree] (config>router>static-route-entry>next-hop bfd-enable)

Full Context

configure router static-route-entry next-hop bfd-enable

Description

This command associates the static route state to a BFD session between the local system and the configured nexthop.

The remote end of the BFD session must also be configured to originate or accept the BFD session controlling the static route state.

The **no** form of this command removes the association of the static route state to that of the BFD session.

Default

no bfd-enable

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable interface** *interface-name* **dst-ip** *ip-address*

[**no**] **bfd-enable interface** *interface-name* **dst-ip** *ip-address* **name** *name*

[**no**] **bfd-enable** *svc-id* **interface** *interface-name* **dst-ip** *ip-address*

Context

[Tree] (config>router>if>vrrp bfd-enable)

[Tree] (config>router>if>ipv6>vrrp bfd-enable)

Full Context

configure router interface vrrp bfd-enable

configure router interface ipv6 vrrp bfd-enable

Description

This commands assigns a bidirectional forwarding detect (BFD) session to a specific VRRP/SRRP instance. This BFD sessions provided a heartbeat mechanism that can be used to speed up the transition of the standby VRRP router to an active state. If the associated BFD session fails, the VRRP routers will immediately send a VRRP Advertisement message. In addition, the standby VRRP router(s) will transition to a Master state to speed convergence. The normal VRRP election process will then take place based on the Advertisement messages sent by all VRRP routers.

There can be only one BFD session assigned to any given VRRP/SRRP instance, but there can be multiple SRRP/VRRP sessions using the same BFD session.

The parameters used for the BFD sessions are set by the BFD command under the IP interface.

The **no** form of this command removes BFD from the configuration.

Parameters

interface-name

Specifies the name of the interface running BFD. The specified interface may not yet be configured with BFD. However, when it is, this virtual router will then initiate the BFD session.

ip-address

Specifies the destination address to be used for the BFD session.

svc-id

Specifies the service ID of the interface running BFD.

Values *service-id*:1 to 2147483647

svc-name: 64 characters maximum

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

bfd-enable interface *interface-name* **dest-ip** *ipv4-address* [**service** *service-id*]

no bfd-enable

Context

[Tree] (config>service>oper-group bfd-enable)

Full Context

configure service oper-group bfd-enable

Description

This command associates a BFD sessions with the named oper-group so that if the BFD session fails then the oper-group is changed to operationally down and all monitoring interfaces should also be brought operationally down.

Parameters

interface-name

Specifies the source interface, up to 32 characters in length, for the BFD sessions to be monitored for the associated oper-group.

ipv4-address

Specifies the destination IPv4 address for the BFD sessions to be monitored for the associated oper-group.

service-id

Specifies the service ID, up to 64 characters in length, in which the BFD session exists if it is not in the base routing context.

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

bfd-enable {**ipv4** | **ipv6**} [**include-bfd-tlv**]

no bfd-enable {**ipv4** | **ipv6**}

Context

[Tree] (config>router>isis>if bfd-enable)

Full Context

configure router isis interface bfd-enable

Description

This command enables the use of bidirectional forwarding detection (BFD) to control IPv4 or IPv6 adjacencies. By enabling BFD on an IPv4 or IPv6 protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for

the BFD are set by the BFD command under the IP interface. This command must be given separately to enable or disable BFD for both IPv4 and IPv6.

The **no** form of this command removes BFD from the associated adjacency.

Default

no bfd-enable ipv4

no bfd-enable ipv6

Parameters

ipv4

Keyword to enable BFD to control IPv4 adjacencies.

ipv6

Keyword to enable BFD to control IPv6 adjacencies.

include-bfd-tlv

Enables support for the IS-IS BFD TLV options in accordance with RFC 6213, which specifies that a BFD session must be established before an IS-IS adjacency can transition to the established state. This option must be enabled on all IS-IS neighbors on a shared interface.

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

bfd-enable [**remain-down-on-failure**]

bfd-enable [**remain-down-on-failure**] **strict** [**strict-mode-holddown** *number*]

no bfd-enable

Context

[Tree] (config>router>ospf3>area>interface bfd-enable)

[Tree] (config>router>ospf>area>interface bfd-enable)

Full Context

configure router ospf3 area interface bfd-enable

configure router ospf area interface bfd-enable

Description

This command configures BFD to control the state of the associated protocol interface. By enabling BFD on a protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD session are set through the **bfd** command under the IP interface.

The **no** form of this command removes BFD from the associated OSPF protocol adjacency.

Default

no bfd-enable

Parameters

remain-down-on-failure

Keyword to specify that OSPF brings down the adjacency and waits on BFD again if the BFD session does not come back up within 10 seconds. This can cause OSPF neighbors to flap, because OSPF will form the adjacency and then bring it down if the BFD session is still down. If this parameter is not configured, the OSPF adjacency will form even if the BFD adjacency does not come back up after a failure.

strict

Keyword to specify that the system uses BFD strict-mode, which requires that an active BFD session exists between the OSPF neighbors before establishing a full adjacency. When this keyword is configured, the router uses Link-Local Signaling (LLS) with the B-flag set to instruct OSPF neighbors that BFD must be enabled on the link. BFD strict-mode requires that both sides have the B-flag set.

During OSPFv3 BFD strict-mode operations, the router advertises the Local Interface IPv4 Address TLV using LLS, but the SR OS router continues to use IPv6-based BFD sessions for both the IPv4 and IPv6 address families.

strict-mode-holddown *number*

Keyword to specify a delay in bringing up the OSPF adjacency after the BFD session is established. Holddown helps mitigate potential routing churn when BFD sessions are unstable. The holddown timer is reset on an adjacency when a BFD session operationally toggles.

Values 1 to 600

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable**

Context

[Tree] (config>router>rip>group bfd-enable)

[Tree] (config>router>ripng>group bfd-enable)

[Tree] (config>router>rip bfd-enable)

[Tree] (config>router>ripng>group>neighbor bfd-enable)

[Tree] (config>router>ripng bfd-enable)

[Tree] (config>router>rip>group>neighbor bfd-enable)

Full Context

configure router rip group bfd-enable

configure router ripng group bfd-enable

configure router rip bfd-enable

configure router ripng group neighbor bfd-enable

configure router ripng bfd-enable

configure router rip group neighbor bfd-enable

Description

This command enables bidirectional forwarding detection (BFD) to control the state of the associated protocol adjacency. By enabling BFD on a given protocol interface, the state of the RIP neighbor is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set using the **bfd** command under the IP interface configuration context.

The **no** form of this command removes BFD from the associated protocol adjacency.

Platforms

7705 SAR Gen 2

bfd-enable

Syntax

[**no**] **bfd-enable**

Context

[Tree] (config>router>segment-routing>maintenance-policy bfd-enable)

Full Context

configure router segment-routing maintenance-policy bfd-enable

Description

This command enables seamless BFD on every programmed segment list of an SR policy candidate path to which the maintenance policy is applied. BFD session parameters are taken from the BFD template that is configured for the maintenance policy.

The **no** form of this command disables seamless BFD on every segment list of an SR policy.

Default

no bfd-enable

Platforms

7705 SAR Gen 2

6.15 bfd-strict-mode

bfd-strict-mode

Syntax

bfd-strict-mode

Context

[Tree] (config>router>bgp bfd-strict-mode)

[Tree] (config>service>vprn>bgp>group>neighbor bfd-strict-mode)

[Tree] (config>router>bgp>group bfd-strict-mode)

[Tree] (config>router>bgp>group>neighbor bfd-strict-mode)

[Tree] (config>service>vprn>bgp bfd-strict-mode)

[Tree] (config>service>vprn>bgp>group bfd-strict-mode)

Full Context

configure router bgp bfd-strict-mode

configure service vprn bgp group neighbor bfd-strict-mode

configure router bgp group bfd-strict-mode

configure router bgp group neighbor bfd-strict-mode

configure service vprn bgp bfd-strict-mode

configure service vprn bgp group bfd-strict-mode

Description

Commands in this context configure the BFD Strict-Mode feature.

Platforms

7705 SAR Gen 2

6.16 bfd-template

bfd-template

Syntax

[**no**] **bfd-template** *name*

Context

[Tree] (config>router>bfd bfd-template)

Full Context

configure router bfd bfd-template

Description

This command configures a BFD template. A BFD template defines the set of configurable parameters used by a BFD session. These include the transmit and receive timer intervals used for BFD CC packets, the transmit timer interval used when the session is providing a CV function, the multiplier value, the echo-receive interval, and whether the BFD session terminates in the CPM network processor.

The **no** form of this command reverts to the default value.

Default

no bfd-template

Parameters

name

Specifies a text string name for the template, up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

Platforms

7705 SAR Gen 2

bfd-template

Syntax

bfd-template *bfd-template*

no bfd-template

Context

[Tree] (config>router>segment-routing>maintenance-policy bfd-template)

Full Context

configure router segment-routing maintenance-policy bfd-template

Description

This command references a named BFD template that is used by seamless BFD. The template specifies parameters, such as the minimum transmit and receive control packet timer intervals, that are used by the BFD session. Templates are configured under the **config**>**router**>**bfd** context.

A BFD template must exist on the system before being referenced from a maintenance policy.

The **no** form of this command removes the configured template.

Parameters

bfd-template

Specifies the name of the BFD template, up to 32 characters.

Platforms

7705 SAR Gen 2

6.17 bgp

bgp

Syntax

[**no**] **bgp** [*bgp-instance*]

Context

[Tree] (config>service>epipe bgp)

Full Context

configure service epipe bgp

Description

Commands in this context configure the BGP-related parameters BGP uses for multihoming and BGP VPWS.

The **no** form of this command removes this string from the configuration.

Default

bgp 1

Parameters

bgp-instance

The BGP instance.

Values 1 to 2

Platforms

7705 SAR Gen 2

bgp

Syntax

[**no**] **bgp**

Context

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter bgp)

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel>resolution-filter bgp)

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter bgp)

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter bgp)

Full Context

configure service vprn bgp-evpn mpls auto-bind-tunnel resolution-filter bgp

configure service vprn bgp-ipvpn mpls auto-bind-tunnel resolution-filter bgp

configure service epipe bgp-evpn mpls auto-bind-tunnel resolution-filter bgp

configure service vpls bgp-evpn mpls auto-bind-tunnel resolution-filter bgp

Description

This command selects the BGP tunnel type.

This command instructs BGP EVPN to search for a BGP LSP to the address of the BGP next hop. If the user does not enable the BGP tunnel type, inter-area or inter-as prefixes are not resolved.

The **no** form of this command removes the BGP tunnel type configuration.

Default

bgp

Platforms

7705 SAR Gen 2

bgp

Syntax

bgp *bgp-instance*

no bgp *bgp-instance*

Context

[Tree] (config>service>vpls bgp)

Full Context

configure service vpls bgp

Description

Commands in this context configure the BGP related parameters for BGP VPLS.

A maximum of two BGP instances can be configured in a VPLS service. The *bgp-instance* parameter value can be configured as 1 or 2. If it is not specified, the parameter value is configured as 1 by default.

The **route-distinguisher** configured in BGP instance 1 and 2 must be different. However, the route-target value may be configured the same or different for the two instances.

Only BGP-EVPN MPLS is allowed to be assigned to instance 2. Instance 1 must be used for the VXLAN and L2VPN address families.

BGP-EVPN VXLAN and BGP-EVPN MPLS can only be configured as **no shutdown** in the same service if they are associated with different instances (When the two BGP instances are created, the **bgp-instance** command must be configured in the **bgp-evpn mpls** context).

The **evi** value in **bgp-evpn** can be used to auto-derive the route distinguisher in instance 1 only. However, the **evi** value can be used to auto-derive the **route-target** in both instances.

The **no** version of the command removes the BGP instance.

Parameters

bgp-instance

Specifies the value associated with the BGP instance.

Values 1 to 2

Platforms

7705 SAR Gen 2

bgp

Syntax

[**no**] **bgp**

Context

[Tree] (config>router bgp)

Full Context

configure router bgp

Description

This command creates the BGP protocol instance and BGP configuration context. BGP is administratively enabled upon creation.

The **no** form of this command deletes the BGP protocol instance and removes all configuration parameters for the BGP instance. BGP must be **shutdown** before deleting the BGP instance. An error occurs if BGP is not **shutdown** first.

Platforms

7705 SAR Gen 2

bgp

Syntax

[**no**] **bgp**

Context

[Tree] (config>service>vprn bgp)

Full Context

configure service vprn bgp

Description

This command enables the BGP protocol with the VPRN service.

The **no** form of this command disables the BGP protocol from the given VPRN service.

Default

no bgp

Platforms

7705 SAR Gen 2

bgp

Syntax

bgp [**source** *src-Addr*] [**group** *grpAddr*] [**peer** *peerAddr*]

no bgp

Context

[Tree] (debug>router>pim bgp)

Full Context

debug router pim bgp

Description

This command enables debugging for PIM/BGP-specific interoperation.

The **no** form of this command disables debugging for PIM/BGP-specific interoperation.

Parameters

src-Addr

Debugs BGP information associated with the specified source.

Values source address (ipv4, ipv6)

grp-Addr

Debugs BGP information associated with the specified group.

Values group address (ipv4, ipv6)

PeerAddr

Debugs BGP information associated with the specified peer.

Values peer address (ipv4, ipv6)

Platforms

7705 SAR Gen 2

bgp

Syntax

[**no**] **bgp**

Context

[Tree] (config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter bgp)

[Tree] (config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter bgp)

Full Context

configure router bgp next-hop-resolution labeled-routes transport-tunnel family resolution-filter bgp

configure router bgp next-hop-resolution shortcut-tunnel family resolution-filter bgp

Description

This command selects BGP tunneling for next-hop resolution and specifies the IPv4 tunnels created by receiving BGP label-unicast IPv4 routes for /32.

The **no** form of this command disables the selection of BGP tunneling for next-hop resolution.

Platforms

7705 SAR Gen 2

bgp

Syntax

bgp

Context

[Tree] (config>service>vprn>auto-bind-tunnel>res-filter bgp)

Full Context

configure service vprn auto-bind-tunnel resolution-filter bgp

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

6.18 bgp-ad

bgp-ad

Syntax

[**no**] **bgp-ad**

Context

[Tree] (config>service>vpls bgp-ad)

Full Context

configure service vpls bgp-ad

Description

This command configures BGP auto-discovery.

Platforms

7705 SAR Gen 2

6.19 bgp-auto-rd-range

bgp-auto-rd-range

Syntax

bgp-auto-rd-range *ip-address* **comm-val** *comm-val* **to** *comm-val*

no bgp-auto-rd-range

Context

[Tree] (config>service>system bgp-auto-rd-range)

Full Context

configure service system bgp-auto-rd-range

Description

This command defines the type-1 route-distinguisher IPv4 address and community value range within which the system will select a route-distinguisher for the **bgp-enabled** services using **auto-rd**.

Interactions:

This command is used along with the **route-distinguisher auto-rd** command supported in VPLS, VPRN and Epipe services. The system forces the user to create a **bgp-auto-range** before the **auto-rd** option can be used in the services.

The system will keep allocating values for services configured with **route-distinguisher auto-rd** as long as there are available community values within the configured range. After the command is added, the following changes are allowed:

- The ip-address can be changed without modifying the *comm-val* range, even if services using **auto-rd** are present. The affected routes will be withdrawn and re-advertised with the new route-distinguishers.
- The *comm-val* range can be modified as long as no conflicting values are present in the new range. For example, the user may expand the range as long as the new range does not overlap with existing manual route-distinguishers. The user may also reduce the range as long as the new range can accommodate the already allocated auto-RDs.

Parameters

ip-address

Specifies the IPv4 address used in the first 4 octets of all the type-1 auto route-distinguishers selected by the system.

comm-val

Specifies the community value of the type-1 auto route-distinguisher.

Values 1 to 65535

Platforms
7705 SAR Gen 2

6.20 bgp-evpn

bgp-evpn

Syntax
[**no**] **bgp-evpn**

Context
[Tree] (config>service>system bgp-evpn)
[Tree] (config>service>epipe bgp-evpn)
[Tree] (config>service>vpls bgp-evpn)

Full Context
configure service system bgp-evpn
configure service epipe bgp-evpn
configure service vpls bgp-evpn

Description
Commands in this context configure the BGP EVPN parameters in the base instance.

Platforms
7705 SAR Gen 2

bgp-evpn

Syntax
bgp-evpn

Context
[Tree] (config>service>vprn bgp-evpn)

Full Context
configure service vprn bgp-evpn

Description
Commands in this context configure the BGP EVPN parameters.

Platforms

7705 SAR Gen 2

6.21 bgp-high-priority

bgp-high-priority

Syntax

[**no**] **bgp-high-priority**

Context

[Tree] (config>router>policy-options>policy-statement>entry>action bgp-high-priority)

[Tree] (config>router>policy-options>policy-statement>default-action bgp-high-priority)

Full Context

configure router policy-options policy-statement entry action bgp-high-priority

configure router policy-options policy-statement default-action bgp-high-priority

Description

This command enables eligible BGP routes matched by the policy entry or policy default-action that are tagged for faster route table updates.

This action applies only when the policy is applied as a BGP import policy to a base router BGP peer or VPRN BGP peer and applies only to the following route types:

- IPv4
- label-IPv4
- IPv6
- label-IPv6

This command is useful when the BGP RIB contains a large number of routes and quick routing table updates are needed for a small subset of these routes. The effectiveness of this command decreases as the subset becomes a larger proportion of the total RIB.

The **no** form of this command disables the routes that are tagged for faster route table updates.

Default

no bgp-high-priority

Platforms

7705 SAR Gen 2

6.22 bgp-ipvpn

bgp-ipvpn

Syntax

bgp-ipvpn

Context

[Tree] (config>service>vprn bgp-ipvpn)

Full Context

configure service vprn bgp-ipvpn

Description

Commands in this context configure the BGP IPVPN parameters.

Platforms

7705 SAR Gen 2

6.23 bgp-labels-hold-timer

bgp-labels-hold-timer

Syntax

bgp-labels-hold-timer *seconds*

no bgp-labels-hold-timer

Context

[Tree] (config>router>mpls-labels bgp-labels-hold-timer)

Full Context

configure router mpls-labels bgp-labels-hold-timer

Description

This command configures the time to delay before the label-forwarding entries programmed by BGP are removed from the datapath. A non-zero delay is useful in the following situations:

- label-unicast route is readvertised by an ABR/ASBR operating in label-per-next-hop mode to choose a new primary path

- IP VPN route is readvertised by an ABR/ASBR operating in label-per-next-hop mode to choose a new primary path
- IP VPN best-external route is readvertised by a VPRN to choose a new backup path
- IP VPN route is readvertised by a VPRN in label-per-next-hop mode to choose a new primary path

In the preceding situations, configure the hold timer to be large enough to account for the propagation delay of the route withdrawal to all ingress routers.

Default

bgp-labels-hold-timer 0

Parameters

seconds

Specifies the time delay, in seconds.

Values 0 to 255

Platforms

7705 SAR Gen 2

6.24 bgp-leak

bgp-leak

Syntax

[**no**] **bgp-leak**

Context

[Tree] (config>router>policy-options>policy-statement>entry>action bgp-leak)

[Tree] (config>router>policy-options>policy-statement>default-action bgp-leak)

Full Context

configure router policy-options policy-statement entry action bgp-leak

configure router policy-options policy-statement default-action bgp-leak

Description

This command causes qualifying matched BGP routes to be marked as leakable, meaning they are candidates to be leaked into other routing instances (copied with their complete set of path attributes). A BGP route is a qualifying route if it is an IPv4 route (unlabeled), IPv6 route (unlabeled) or a label-IPv4 route.

Note:
A leakable BGP route is not actually leaked into another routing instance unless it is accepted by a leak-import policy of that other routing instance.

The **bgp-leak** command has an effect only when the policy is applied as a BGP import policy in the base router or a VPRN context.

Default

no bgp-leak

Platforms

7705 SAR Gen 2

6.25 bgp-med

bgp-med

Syntax

bgp-med adjust *expression*

bgp-med set {**igp** | **min-igp**}

bgp-med set *med-value*

no bgp-med

Context

[Tree] (config>router>policy-options>policy-statement>entry>action bgp-med)

[Tree] (config>router>policy-options>policy-statement>default-action bgp-med)

Full Context

configure router policy-options policy-statement entry action bgp-med

configure router policy-options policy-statement default-action bgp-med

Description

This command changes the BGP MED attribute value in BGP routes matched by the route policy entry (or the policy default action).

If the matched route already has a MED attribute, this command overwrites the existing value. If the matched route does not have a MED attribute, then one is added and the value is set based on the parameters of this command.

This command has no effect on non-BGP routes. The default, **no bgp-med**, does not modify MED values.

Default

no bgp-med

Parameters

expression

Specifies a logical expression parsed as a string. The string can contain:

- parentheses () to change the order of operations
- mathematical operators: + (addition), - (subtraction) and * (multiplication)
- directly entered decimal values that act as operands of the mathematical operators. Each decimal value supports up to three decimal places precision in the range of 0.000 to 4294967295.000
- decimal values represented by parameter names (using the usual @parameter-name@ syntax) that act as operands of the mathematical operators. Each parameterized decimal value supports up to three decimal places precision in the range of 0.000 to 4294967295.000

igp

Instructs the policy to set the MED based on the current route table or tunnel table cost to resolve the BGP next-hop address.

min-igp

Instructs the policy to set the MED based on the minimum route table or tunnel table cost to resolve the BGP next-hop of the route, over its lifetime in the local RIB.

med-value

Specifies a new MED value (or parameter name to use for the new MED value) to use with the route.

Values *value*

- 0 to 4294967295

param-name

- up to 32 characters
- Must start and end with an at-sign (@)

Platforms

7705 SAR Gen 2

6.26 bgp-multi-homing

bgp-multi-homing

Syntax

bgp-multi-homing

Context

[Tree] (config>redundancy bgp-multi-homing)

Full Context
configure redundancy bgp-multi-homing

Description
This command configures BGP multi-homing parameters.

Platforms
7705 SAR Gen 2

6.27 bgp-path-selection

bgp-path-selection

Syntax
[**no**] **bgp-path-selection**

Context
[Tree] (config>service>system>bgp-evpn>ad-per-evi-routes bgp-path-selection)

Full Context
configure service system bgp-evpn ad-per-evi-routes bgp-path-selection

Description
This command compares the received EVPN VPWS AD per-EVI routes based on BGP path attributes.

Attribute propagation must be configured before configuring this command.

The **no** form of this command disables the comparison of the routes.

Default
no bgp-path-selection

Platforms
7705 SAR Gen 2

bgp-path-selection

Syntax
bgp-path-selection [**d-path-length-ignore**]

no bgp-path-selection

Context

[Tree] (config>service>system>bgp-evpn>ip-prefix-routes>iff bgp-path-selection)

Full Context

configure service system bgp-evpn ip-prefix-routes interface-ful bgp-path-selection

Description

This command enables BGP path selection for EVPN-IFF (Interface-ful) routes.

Once the command is enabled, the EVPN-IFF routes are ordered and selected in a similar manner as IPVPN or EVPN-IFL routes, that is, based on the regular BGP path selection process.

The **no** form of this command causes the system to order EVPN-IFF routes based on their {R-VPLS Ifindex, RD, Ethernet Tag}. For example, if two EVPN-IFF routes with different Route Distinguishers (RDs) are received for the same prefix on the same R-VPLS, the route with the lowest RD is selected.

Default

no bgp-path-selection

Parameters

d-path-length-ignore

Keyword used to make EVPN ignore the D-PATH length when **iff-bgp-path-selection** is enabled.

Platforms

7705 SAR Gen 2

6.28 bgp-peers

bgp-peers

Syntax

bgp-peers *criterion-index* **group** *reg-exp* **neighbor** *reg-exp*

bgp-peers *criterion-index* **router** *router-instance* **group** *reg-exp* **neighbor** *reg-exp*

bgp-peers *criterion-index* **router service-name** *service-name* **group** *reg-exp* **neighbor** *reg-exp*

no bgp-peers *criterion-index*

Context

[Tree] (config>filter>match-list>ipv6-prefix-list>apply-path bgp-peers)

[Tree] (config>filter>match-list>ip-prefix-list>apply-path bgp-peers)

Full Context

configure filter match-list ipv6-prefix-list apply-path bgp-peers

configure filter match-list ip-prefix-list apply-path bgp-peers

Description

This command configures auto-generation of IPv4 or IPv6 address prefixes (as required by the context that the command is executed within) based on the base router BGP instance configuration.

The **no** form of this command removes the bgp-peers configuration for auto-generation of address prefixes for the specified index value.

Parameters

service-name

Specifies the service name, up to 64 characters in length.

group

Configures a match against the base router BGP instance group configuration.

Regex match (.*) can be used to match against any group.

neighbor

Configures a match against the base router BGP instance neighbor configuration.

Regex match (.*) can be used to match against any neighbor.

criterion-index

Specifies an integer from 1 to 255 enumerating BGP peers auto-generation configuration within this list.

router-instance

Specifies the router name or service ID.

Values router-instance: *router-name* or *vprn-svc-id*

router-name: "Base"

vprn-svc-id: 1 to 2147483647

service-name: Specifies the service name, up to 64 characters in length.

router

Configures a match against the base router BGP instance configuration.

reg-exp

Specifies a regular expression that defines a match string, up to 255 characters in length, to be used to auto-generate address prefixes. Matching is performed from the least-significant digit. For example, a string **10.0** matches all neighbors with addresses starting with **10**, such as **10.0.x.x** or **10.0xx.x.x**.

Platforms

7705 SAR Gen 2

6.29 bgp-shared-queue

bgp-shared-queue

Syntax

bgp-shared-queue [**cir** *rate*] [**pir** *rate*]

no bgp-shared-queue

Context

[Tree] (config>service>vprn bgp-shared-queue)

Full Context

configure service vprn bgp-shared-queue

Description

This command enables all BGP peers within a VPRN instance to share a single CPM queue. This command takes effect on new BGP connections established; already established BGP peers continue to use their own CPM queue. Any changes to PIR/CIR of the shared queue takes effect only after BGP connections are re-established.

Parameters

cir *rate*

Specifies the CIR rate for the shared queue.

pir *rate*

Specifies the PIR rate for the shared queue.

Platforms

7705 SAR Gen 2

6.30 bgp-shortcut

bgp-shortcut

Syntax

[**no**] **bgp-shortcut**

Context

[Tree] (config>router>mpls>lsp bgp-shortcut)

[Tree] (config>router>mpls>lsp-template bgp-shortcut)

Full Context

configure router mpls lsp bgp-shortcut

configure router mpls lsp-template bgp-shortcut

Description

This command enables the use of RSVP LSP for IPv4 BGP routes.

Platforms

7705 SAR Gen 2

6.31 bgp-transport-tunnel

bgp-transport-tunnel

Syntax

bgp-transport-tunnel [**include** | **exclude**]

Context

[Tree] (config>router>mpls>lsp bgp-transport-tunnel)

[Tree] (config>router>mpls>lsp-template bgp-transport-tunnel)

Full Context

configure router mpls lsp bgp-transport-tunnel

configure router mpls lsp-template bgp-transport-tunnel

Description

This command allows or blocks RSVP-TE LSP to be used as a transport LSP for BGP tunnel routes.

Default

bgp-transport-tunnel include

Parameters

include

Allows RSVP-TE LSP to be used as transport LSP from the ASBR to local PE router, from ingress PE to ASBR in the local AS or between multi-hop External Border Gateway Protocol (EBGP) peers with ASBR to ASBR adjacency.

exclude

Blocks RSVP-TE LSP to be used as transport LSP from the ASBR to local PE router, from ingress PE to ASBR in the local AS or between multi-hop EBGP peers with ASBR to ASBR adjacency.

Platforms

7705 SAR Gen 2

6.32 bgp-tunnel

bgp-tunnel

Syntax

[**no**] **bgp-tunnel**

Context

[Tree] (config>service>sdp bgp-tunnel)

Full Context

configure service sdp bgp-tunnel

Description

This command allows the use of BGP route tunnels available in the tunnel table to reach SDP far-end nodes. Use of BGP route tunnels are only available with MPLS-SDP. Only one of the transport methods is allowed per SDP - LDP, RSVP-LSP BGP, SR-ISIS, or SR-OSPF. This restriction is relaxed for some combinations of the transport methods when the mixed-lsp-mode option is enabled within the SDP.

The **no** form of the command disables resolving BGP route tunnel LSP for SDP far-end.

Default

no bgp-tunnel (BGP tunnel route to SDP far-end is disabled)

Platforms

7705 SAR Gen 2

6.33 bgp-tunnel-metric

bgp-tunnel-metric

Syntax

bgp-tunnel-metric [*value*] [**prefer-med**]

bgp-tunnel-metric [*value*] **prefer-aigp**

bgp-tunnel-metric [*value*] **prefer-aigp prefer-med**

bgp-tunnel-metric [*value*] [**prefer-aigp**]

no bgp-tunnel-metric

Context

[Tree] (config>router>bgp bgp-tunnel-metric)

Full Context

configure router bgp bgp-tunnel-metric

Description

This command sets the TTM metric of all BGP tunnels to a fixed value or a value derived from the AIGP or the MED metric of the BGP-LU route, if the BGP-LU route has an AIGP or MED path attribute. Otherwise, the TTM metric is set to the number specified using the *value* parameter. BGP import policies override the configuration of this command.

By default, BGP tunnels are installed with a fixed cost of 1000 in the tunnel table. This can overstate or understate their true cost when compared to other tunnels with IGP-derived costs.

The **no** form of the command configures the router to use the default value.

Default

no bgp-tunnel-metric

Parameters

value

Specifies the BGP tunnel metric.

Values 0 to 4294967295

prefer-aigp

Specifies that the TTM metric is based on the AIGP metric value of the BGP-LU route. When a BGP-LU route is selected for installation in TTM and is not matched by a BGP import policy entry that overrides the BGP tunnel metric action, the TTM metric of the tunnel is set to the AIGP metric value of the BGP-LU route with the resolved cost to the BGP next hop of the route added to it. Otherwise, the metric is set to the value configured using the *value* parameter.

prefer-med

Specifies that the TTM metric is based on the MED metric value of the BGP-LU route. When a BGP-LU route is selected for installation in TTM and is not matched by a BGP import policy entry that overrides the BGP tunnel metric action, the TTM metric of the tunnel is set to the MED metric value of the BGP-LU route with the resolved cost to the BGP next hop of the route added to it. Otherwise, the metric is set to the value configured using the *value* parameter.

Note: prefer-aigp takes precedence over this parameter if the received BGP-LU has both attributes.

Platforms

7705 SAR Gen 2

bgp-tunnel-metric

Syntax

bgp-tunnel-metric [*value* | *param-name*] [**prefer-aigp**] [**prefer-med**]

no bgp-tunnel-metric

Context

[Tree] (config>router>policy-options>policy-statement>default-action bgp-tunnel-metric)

[Tree] (config>router>policy-options>policy-statement>entry>action bgp-tunnel-metric)

Full Context

configure router policy-options policy-statement default-action bgp-tunnel-metric

configure router policy-options policy-statement entry action bgp-tunnel-metric

Description

This command sets the TTM metric of all BGP tunnels matched by the policy entry or the policy default action to a fixed value or a value derived from the AIGP or the MED metric of the BGP-LU route, if the BGP-LU route has an AIGP or MED path attribute. Otherwise, the TTM metric is set to the number specified using the *value* parameter.

The **no** form of this command configures the router to use the default value.

Default

no bgp-tunnel-metric

Parameters

value

Specifies the BGP tunnel metric.

Values 0 to 4294967295

param-name

Specifies the parameter name, up to 32 characters that starts and ends with an at-sign (@).

prefer-aigp

Specifies that if a BGP-LU route is selected for installation in the TTM and is matched by this action in a BGP import policy, the TTM metric of the tunnel is set to the AIGP metric value of the BGP-LU route with the IGP cost to reach the BGP next hop added to it.

prefer-med

Specifies that if a BGP-LU route is selected for installation in the TTM and is matched by this action in a BGP import policy, the TTM metric of the tunnel is set to the MED metric value of the BGP-LU route with the IGP cost to reach the BGP next hop added to it.

Platforms

7705 SAR Gen 2

6.34 bgp-tunnel-preference

bgp-tunnel-preference

Syntax

bgp-tunnel-preference [*preference*]

no bgp-tunnel-preference

Context

[Tree] (config>router>bgp bgp-tunnel-preference)

Full Context

configure router bgp bgp-tunnel-preference

Description

This command configures the tunnel table preference for BGP-LU tunnel type away from its default value.

The tunnel table preference applies to the next-hop resolution of BGP routes of the following families: EVPN, IPv4, IPv6, VPN-IPv4, VPN-IPv6, label-IPv4, and label-IPv6 in the tunnel table.

This feature does not apply to a VPRN, VPLS, or VLL service with explicit binding to an SDP which enabled the **mixed-lsp-mode** option. The tunnel preference, in such an SDP, is fixed and is controlled by the service manager. The configuration of the tunnel table preference parameter does not modify the behavior of such an SDP and the services that bind to it.

It is recommended to not set two or more tunnel types to the same preference value. In such a situation, the tunnel table prefers the tunnel type which was first introduced in SR OS implementation historically.

The **no** form of this command reverts to the default value.

Default
bgp-tunnel-preference 12

Parameters
preference
Specifies the BGP tunnel preference.

Values 1 to 255

Default 12

Platforms
7705 SAR Gen 2

6.35 bgp-vpls

bgp-vpls

Syntax
bgp-vpls

Context
[Tree] (config>service>vpls bgp-vpls)

Full Context
configure service vpls bgp-vpls

Description
Commands in this context configure the BGP-VPLS parameters and addressing.

Platforms
7705 SAR Gen 2

6.36 bgp-vpls-mh-ve-id

bgp-vpls-mh-ve-id

Syntax
bgp-vpls-mh-ve-id *number*

no bgp-vpls-mh-ve-id

Context

[Tree] (config>service>vpls>sap bgp-vpls-mh-ve-id)

Full Context

configure service vpls sap bgp-vpls-mh-ve-id

Description

This command upon the configuration of the ve-id under the SAP and if BGP-VPLS is configured and is operationally up, causes the PE to advertise a bgp-mh route for the ve-id (the route does not contain label information). The bgp-mh route contains the F and D flags properly set based on the SAP operational state. Upon switchover, the former active PE (DF in case of EVPN-MH) sends an update with a transition of the F bit from 1 to 0. This is an indication for the remote PEs to flush their MACs associated to the advertising PE.

This command is required when MC-LAG or EVPN-MH is used for multi-homing redundancy and mac-flush is required at remote BGP-VPLS PEs when there is a failure in the active PE.

The **no** form of this command withdraws the L2 VPN route.

Parameters

number

Specifies the BGP-VPLS multi-homing virtual-edge identifier.

Values 1 to 65535

Platforms

7705 SAR Gen 2

6.37 bgp-vpws

bgp-vpws

Syntax

[**no**] **bgp-vpws**

Context

[Tree] (config>service>epipe bgp-vpws)

Full Context

configure service epipe bgp-vpws

Description
Commands in this context configure BGP-VPWS parameters and addressing.

Default
no bgp-vpws

Platforms
7705 SAR Gen 2

6.38 bi

bi

Syntax
bi

Context
[Tree] (config>system>security>keychain>direction bi)

Full Context
configure system security keychain direction bi

Description
This command configures keys for both send and receive stream directions.

Platforms
7705 SAR Gen 2

6.39 bin

bin

Syntax
bin *bin-number*

Context
[Tree] (config>oam-pm>bin-group>bin-type bin)

Full Context

configure oam-pm bin-group bin-type bin

Description

Commands in this context configure the thresholds for the specified bin.

Parameters

bin-number

Specifies bin to configure.

Values 1 to 9

Platforms

7705 SAR Gen 2

6.40 bin-group

bin-group

Syntax

bin-group *bin-group-number* [**fd-bin-count** *fd-bin-count* **fdr-bin-count** *fdr-bin-count* **ifdv-bin-count** *ifdv-bin-count* **create**]

no bin-group *bin-group-number*

Context

[Tree] (config>oam-pm bin-group)

Full Context

configure oam-pm bin-group

Description

This command allows the operator to configure the parameters for a specific bin group. Bin-group 1 is a default **bin-group** and cannot be modified. If no bin group is assigned to an oam-pm session, this is assigned by default. The default values for bin-group 1 are (fd-bin-count 3 bin 1 lower-bound 5000us, bin 2 lower-bound 10000us fdr-bin-count 2 bin 1lower-bound 5000us and ifdv-bin-count 2 bin 1lower-bound 5000us)

The **no** form of this command disables the OAM Performance Monitoring bin group.

Parameters

bin-group-number

Specifies an identifier for a bin-group that is referenced by oam-pm sessions. A bin group can only shutdown and modified when all the PM Sessions referencing the bin group have been shutdown. The only exception is the description parameter.

Values 1 to 255

fd-bin-count

Specifies the number of frame delay bins that are created.

Values 2 to 10

fdr-bin-count

Specifies the number of frame delay range bins that are created.

Values 2 to 10

ifdv-bin-count

Specifies the number of inter-frame delay variation bins that are created.

Values 2 to 10

create

Keyword that creates the bin group.

Platforms

7705 SAR Gen 2

bin-group

Syntax

bin-group *bin-group-number*

no bin-group

Context

[Tree] (config>oam-pm>session bin-group)

Full Context

configure oam-pm session bin-group

Description

This command links the individual test to the group of bins that map the probe responses.

The **no** form of this command installs the default bin-group 1 as the bin-group for the session.

Parameters

bin-group-number

Specifies the number that was used to create the specific **bin-group** that is referenced for this session.

Values 1 to 255

Default 1

Platforms

7705 SAR Gen 2

6.41 bin-type

bin-type

Syntax

bin-type {**fd** | **fdr** | **ifdv**}

Context

[Tree] (config>oam-pm>bin-group bin-type)

Full Context

configure oam-pm bin-group bin-type

Description

This command is the start of the hierarchy where the specific delay metric bin structure isis defined.

Parameters

fd

Keyword to enter the frame delay bin threshold configuration.

fdr

Keyword to enter the frame delay range bin threshold configuration.

ifdv

Keyword to enter the inter-frame delay variation bin thresholds configuration.

Platforms

7705 SAR Gen 2

6.42 bind-authentication

bind-authentication

Syntax

bind-authentication *root-dn* [**password** *password*] [**hash** | **hash2** | **custom**]

no bind-authentication

Context

[Tree] (config>system>security>ldap>server bind-authentication)

Full Context

configure system security ldap server bind-authentication

Description

This command configures the LDAP binding used to log into LDAP server. A string of domain components (DC) and common names (CN) can be programmed to identify the user in addition to the password field. The password is hashed. For example, "cn=admin,dc=nokia,dc=com" indicates the user admin in domain nokia.com. Table 19: LDAP Attributes lists the LDAP attributes.

The **no** version of this command removes the bind-authentication.

Table 19: LDAP Attributes

Object Class	Naming Attribute Display Name	Naming Attribute LDAP Name
user	Common-Name	cn
organizationalUnit	Organizational-Unit-Name	ou
domain	Domain-Component	dc

Parameters

root-dn

Up to 512 characters.

password

Configures the password which enables a user to bind to the LDAP server. The maximum length is 128 characters.

hash

Specifies that the password is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the password is assumed to be in an unencrypted, clear text form. For security, all passwords are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2

Specifies the password is entered in a more complex encrypted form that involves more variables than the password value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the password is assumed to be in an unencrypted, clear text form. For security, all passwords are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

6.43 binding-label

binding-label

Syntax

binding-label *label-number*

no binding-label

Context

[Tree] (config>router>mpls>fwd-policies>fwd-policy binding-label)

Full Context

configure router mpls forwarding-policies forwarding-policy binding-label

Description

This command configures a binding label for the MPLS forwarding policy.

The policy associates an incoming label, referred to as a binding label, to an NHG in which the primary and backup direct or indirect next hops are defined. This type of MPLS forwarding policy is referred to as a label-binding policy.

The **no** form of the command removes the binding label from the MPLS forwarding policy.

Parameters

label-number

Specifies the label number.

Values 32 to 1048575

Platforms

7705 SAR Gen 2

6.44 binding-operator

binding-operator

Syntax

binding-operator {**and** | **or**}

no binding-operator

Context

[Tree] (config>filter>redirect-policy-binding binding-operator)

Full Context

configure filter redirect-policy-binding binding-operator

Description

This command configures the logical operator to use with the destinations test results to obtain the master test result (the redirect-policy binding test result). A change in this configuration results in the re-evaluation of the master test result.

The **no** version of this command sets the value to its default

Default

binding-operator and

Parameters

and | or

Keyword to specify the type of logical or boolean operation to perform between the individual destinations test results to obtain the master result.

Platforms

7705 SAR Gen 2

6.45 binding-sid

binding-sid

Syntax

binding-sid *number*

no binding-sid

Context

[Tree] (conf>router>segment-routing>sr-policies>policy binding-sid)

Full Context

configure router segment-routing sr-policies static-policy binding-sid

Description

This command associates a binding SID with a statically defined segment routing policy. This is a mandatory parameter and configuration command to enable the segment routing policy; if the binding SID label value is not configured, the execution of the **no shutdown** command on the static segment routing policy fails.The BSID label should be an available label in the **reserved-label-block** range.

The **no** form of this command removes the BSID association.

Default

no binding-sid

Parameters

number

Specifies the binding SID label value.

Values 32 to 1048575

Platforms

7705 SAR Gen 2

binding-sid

Syntax

binding-sid *label*

no binding-sid

Context

[Tree] (config>router>mpls>lsp binding-sid)

Full Context

configure router mpls lsp binding-sid

Description

This command configures a binding SID label for the LSP. The label value must belong to the reserved label block that is configured with the **configure router mpls lsp-bsid-block** command.

The **no** form of this command unbinds the label, removes the ILM entry, and triggers the appropriate PCEP messages.

Parameters

label

Specifies an MPLS label value from a specific reserved label block.

Values 32 to 1048575

Platforms

7705 SAR Gen 2

binding-sid

Syntax

[no] binding-sid

Context

[Tree] (config>router>mpls>lsp-template binding-sid)

Full Context

configure router mpls lsp-template binding-sid

Description

This command configures the system to allocate and bind a label to any LSP that is created using the template.

The **no** form of this command removes the configuration but does not affect LSPs that were already created using the template.

Default

no binding-sid

Platforms

7705 SAR Gen 2

6.46 bindings

bindings

Syntax

[**no**] **bindings**

Context

[Tree] (debug>router>ldp>peer>event bindings)

Full Context

debug router ldp peer event bindings

Description

This command displays debugging information about addresses and label bindings learned from LDP peers for LDP bindings.

The **no** form of the command disables the debugging output.

Platforms

7705 SAR Gen 2

6.47 black-hole

black-hole

Syntax

[**no**] **black-hole**

Context

[Tree] (config>service>vprn>static-route-entry black-hole)

Full Context

configure service vprn static-route-entry black-hole

Description

This command specifies that the route is a black hole route. If the destination address on a packet matches this static route, it will be silently discarded.

Default

no black-hole

Platforms

7705 SAR Gen 2

black-hole

Syntax

[**no**] **black-hole**

Context

[Tree] (config>router>static-route-entry black-hole)

Full Context

configure router static-route-entry black-hole

Description

This command specifies that the route is a black hole route. If the destination address on a packet matches this static route, it will be silently discarded.

Default

no black-hole

Platforms

7705 SAR Gen 2

6.48 black-hole-dup-mac

black-hole-dup-mac

Syntax

[**no**] **black-hole-dup-mac**

Context

[Tree] (config>service>vpls>bgp-evpn>mac-duplication black-hole-dup-mac)

Full Context

configure service vpls bgp-evpn mac-duplication black-hole-dup-mac

Description

The **black-hole-dup-mac** command is disabled by default. If enabled, a duplicated MAC detected in the network is programmed as a black-hole MAC in the FDB and displayed in the **show service id fdb detail** command as follows:

- Source-Identifier—black-hole
- Type—EvpnD:P

Because the MAC is now programmed in the FDB as a black-hole, all received frames with MAC DA matching the duplicate MAC are discarded. The duplicate black-hole MACs are installed as Protected, therefore, all received frames with MAC SA matching the duplicate MAC are discarded by default.

A BGP-EVPN (MPLS or VXLAN) shutdown is required to add or remove the **black-hole-dup-mac** command.

The **no** form of the command removes the feature, and duplicate MACs are no longer programmed as black-hole MACs.

Default

no black-hole-dup-mac

Platforms

7705 SAR Gen 2

6.49 blackhole-aggregate

blackhole-aggregate

Syntax

[**no**] **blackhole-aggregate**

Context

[Tree] (config>service>vprn>ospf>area blackhole-aggregate)

[Tree] (config>service>vprn>ospf3>area blackhole-aggregate)

Full Context

configure service vprn ospf area blackhole-aggregate

configure service vprn ospf3 area blackhole-aggregate

Description

This command installs a low priority blackhole route for the entire aggregate. Existing routes that make up the aggregate have a higher priority and only the components of the range for which no route exists are blackholed.

It is possible that when performing area aggregation, addresses may be included in the range for which no actual route exists. This can cause routing loops. To avoid this problem, configure the **blackhole-aggregate** command.

The **no** form of this command removes this configuration.

Default

blackhole-aggregate

Platforms

7705 SAR Gen 2

blackhole-aggregate

Syntax

[**no**] **blackhole-aggregate**

Context

[Tree] (config>router>ospf>area blackhole-aggregate)

[Tree] (config>router>ospf3>area blackhole-aggregate)

Full Context

configure router ospf area blackhole-aggregate

configure router ospf3 area blackhole-aggregate

Description

This command installs a low priority blackhole route for the entire aggregate. Existing routes that make up the aggregate will have a higher priority and only the components of the range for which no route exists are blackholed.

When performing area aggregation, addresses may be included in the range for which no actual route exists, which can cause routing loops. To avoid this problem, configure the **blackhole-aggregate** option.

The **no** form of this command removes this option.

Default

blackhole-aggregate

Platforms

7705 SAR Gen 2

6.50 block-limit

block-limit

Syntax

block-limit [**1..40**]

no block-limit

Context

[Tree] (config>service>nat>nat-policy block-limit)

Full Context

configure service nat nat-policy block-limit

Description

This command configures the maximum number of port blocks per subscriber.

The **no** form of the command reverts to the default.

Default

block-limit 1

Parameters

1..40

Specifies the maximum number of port-blocks per NAT subscriber.

Platforms

7705 SAR Gen 2

6.51 block-on-mesh-failure

block-on-mesh-failure

Syntax

[**no**] **block-on-mesh-failure**

Context

[Tree] (config>service>vpls>endpoint block-on-mesh-failure)

[Tree] (config>service>vpls>spoke-sdp block-on-mesh-failure)

Full Context

configure service vpls endpoint block-on-mesh-failure

configure service vpls spoke-sdp block-on-mesh-failure

Description

This command enables blocking (brings the entity to an operationally down state) after all configured SDPs or endpoints are in operationally down state. This event is signaled to corresponding T-LDP peer by withdrawing service label (status-bit-signaling non-capable peer) or by setting "PW not forwarding" status bit in T-LDP message (status-bit-signaling capable peer).

The **no** form of this command reverts to the default.

Default

no block-on-mesh-failure

Platforms

7705 SAR Gen 2

6.52 block-on-peer-fault

block-on-peer-fault

Syntax

[**no**] **block-on-peer-fault**

Context

[Tree] (config>service>epipe>spoke-sdp block-on-peer-fault)

Full Context

configure service epipe spoke-sdp block-on-peer-fault

Description

When enabled, this command blocks the transmit direction of a PW when any of the following PW status codes is received from the far end PE:

0x00000001	Pseudowire Not Forwarding
0x00000002	Local Attachment Circuit (ingress) Receive Fault
0x00000004	Local Attachment Circuit (egress) Transmit Fault
0x00000008	Local PSN-facing PW (ingress) Receive Fault
0x00000010	Local PSN-facing PW (egress) Transmit Fault

The transmit direction is unblocked when the following PW status code is received:

0x00000000	Pseudowire forwarding (clear all failures)

This command is mutually exclusive with **no pw-status-signaling**, and **standby-signaling-slave**. It is not applicable to spoke SDPs forming part of an MC-LAG or spoke SDPs in an endpoint.

Default

no block-on-peer-fault

Platforms

7705 SAR Gen 2

block-on-peer-fault

Syntax

[**no**] **block-on-peer-fault**

Context

[Tree] (config>service>pw-template block-on-peer-fault)

Full Context

configure service pw-template block-on-peer-fault

Description

When enabled, this command blocks the transmit direction of a pseudowire when any of the following pseudowire status codes is received from the far end PE:

0x00000001	Pseudowire Not Forwarding
0x00000002	Local Attachment Circuit (ingress) Receive Fault
0x00000004	Local Attachment Circuit (egress) Transmit Fault
0x00000008	Local PSN-facing PW (ingress) Receive Fault
0x00000010	Local PSN-facing PW (egress) Transmit Fault

The transmit direction is unblocked when the following pseudowire status code is received:

0x00000000	Pseudowire forwarding (clear all failures)

This command is mutually exclusive with **no pw-status-signaling**, and **standby-signaling-slave**. It is not applicable to spoke SDPs forming part of an MC-LAG or spoke SDPs in an endpoint.

Default

no block-on-peer-fault

Platforms

7705 SAR Gen 2

6.53 block-prefix-sid

block-prefix-sid

Syntax

[**no**] **block-prefix-sid**

Context

[Tree] (config>router>bgp>group>neighbor block-prefix-sid)

[Tree] (config>router>bgp block-prefix-sid)

[Tree] (config>router>bgp>group block-prefix-sid)

Full Context

configure router bgp group neighbor block-prefix-sid

configure router bgp block-prefix-sid

configure router bgp group block-prefix-sid

Description

This command specifies whether all prefix SID attributes are removed from label IPv4 and label IPv6 routes when they are exchanges with EBGP and IBGP peers covered by the scope of the command. Even locally-imposed prefix SID attributes are removed.

A change of this configuration causes the affected BGP sessions to flap.

The **no** form of this command allows prefix SID attributes associated with label IPv4 and label IPv6 routes to be propagated without restriction.

Default

no block-prefix-sid

Platforms

7705 SAR Gen 2

6.54 bof

bof

Syntax

bof

Context

[Tree] (bof)

Full Context

bof

Description

This command creates or edits the boot option file (BOF) for the specified local storage device.

A BOF file specifies where the system searches for runtime images, configuration files, and other operational parameters during system initialization.

BOF parameters can be modified. Changes can be saved to a specified compact flash. The BOF must be located in the root directory of either an internal or external compact flash local to the system and have the mandatory filename of *bof.cfg*.

When modifications are made to in-memory parameters that are currently in use or operating, the changes are effective immediately. For example, if the IP address of the management port is changed, the change takes place immediately.

Only one entry of the BOF configuration command statement can be saved once the statement has been found to be syntactically correct.

When opening an existing BOF that is not the BOF used in the most recent boot, a message is issued notifying the user that the parameters will not affect the operation of the node.

No default boot option file exists. The router boots with the factory default boot sequence and options.

Platforms

7705 SAR Gen 2

6.55 booking-factor

booking-factor

Syntax

booking-factor *factor*

no booking-factor

Context

[Tree] (config>lag>access booking-factor)

[Tree] (config>port>ethernet>access booking-factor)

Full Context

configure lag access booking-factor

configure port ethernet access booking-factor

Description

This command specifies the booking factor applied against the port or LAG administrator bandwidth by SAP administrator bandwidth CAC.

The service manager keeps track of the available administrator bandwidth for each port or LAG configured with an administrator bandwidth. The port or LAG available administrator bandwidth is adjusted by the user configured booking factor, allowing the port or LAG bandwidth to be overbooked or under booked.

If the booking factor is increased then available administrator bandwidth on the port or LAG increases. If the booking factor is decreased then available administrator bandwidth on the port or LAG decreases.

However, if the reduction of available administrator bandwidth is insufficient to cover the sum of the current SAP administrator bandwidth on the port or LAG, the command fails.

The **no** form of this command reverts to the default value.

Default

booking-factor 100

Parameters

factor

Specifies the percentage of the port or LAG admin bandwidth for SAP bandwidth CAC.

Values 1 to 1000

Platforms

7705 SAR Gen 2

booking-factor

Syntax

booking-factor *percentage*

no booking-factor

Context

[Tree] (config>service>sdp booking-factor)

Full Context

configure service sdp booking-factor

Description

This command specifies the booking factor applied against the maximum SDP available bandwidth by the VLL CAC feature.

The service manager keeps track of the available bandwidth for each SDP. The maximum value is the sum of the bandwidths of all constituent LSPs in the SDP. The SDP available bandwidth is adjusted by the user configured booking factor. A value of 0 means no VLL can be admitted into the SDP.

The **no** form of the command reverts to the default value.

Default

no booking-factor

Parameters

percentage

Specifies the percentage of the SDP maximum available bandwidth for VLL call admission. When the value of this parameter is set to zero (0), no new VLL spoke SDP bindings with non-zero bandwidth are permitted with this SDP. Overbooking, >100% is allowed.

Values 0 to 1000%

Default 100

Platforms

7705 SAR Gen 2

6.56 boot-bad-exec

boot-bad-exec

Syntax

boot-bad-exec *file-url*

no boot-bad-exec

Context

[Tree] (config>system boot-bad-exec)

Full Context

configure system boot-bad-exec

Description

Use this command to configure a URL for a CLI script to **exec** following a failure of a bootup configuration. The command specifies a URL for the CLI scripts to be run following the completion of the bootup configuration. A URL must be specified or no action is taken.

The commands are persistent between router (re)boots and are included in the configuration saves (**admin>save**).

Related Commands

exec — This command executes the contents of a text file as if they were CLI commands entered at the console.

Default

no boot-bad-exec

Parameters

file-url

Specifies the location and name of the CLI script file executed following failure of the bootup configuration file execution. When this parameter is not specified, no CLI script file is executed.

Values

file url	*local-url* | *remote-url*	255 chars max
local-url	[*cflash-id*/][*file-path*]	
remote-url	[{ftp://} *login:pswd@remote-locn*/][*file-path*]	
	remote-locn	[*hostname* | *ipv4-address* | [*ipv6- address*]]
	ipv4-address	*a.b.c.d*
	ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
		x:x:x:x:x:x:d.d.d.d[-*interface*]
		x - [0 to FFFF]H
		d - [0 to 255]D
		interface - 32 chars max, for link local addresses
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:	

Platforms

7705 SAR Gen 2

6.57 boot-file-param

boot-file-param

Syntax

boot-file-param *hex-string*

no boot-file-param

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options6 boot-file-param)

Full Context

configure subscriber-mgmt local-user-db ipoe host options6 boot-file-param

Description

This command configures a hexadecimal string that contains the value for the concatenation of all param-len n and parameter n fields of DHCPv6 option BOOTFILE_PARAM (60).

The **no** form of this command removes the configured string.

Parameters

hex-string

Specifies the hexadecimal format for this option, up to 254 hex nibbles.

Values 0x0 to 0xFFFFFFFF

Platforms

7705 SAR Gen 2

6.58 boot-file-url

boot-file-url

Syntax

boot-file-url *ascii-string*

no boot-file-url

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options6 boot-file-url)

Full Context

configure subscriber-mgmt local-user-db ipoe host options6 boot-file-url

Description

This command configures an ASCII string that contains the value for the boot-file-url field in the DHCPv6 option BOOTFILE_URL (59).

The **no** form of this command removes the configuration.

Parameters

ascii-string

Specifies the ASCII string, up to 127 characters.

Platforms

7705 SAR Gen 2

6.59 boot-good-exec

boot-good-exec

Syntax

boot-good-exec *file-url*

no boot-good-exec

Context

[Tree] (config>system boot-good-exec)

Full Context

configure system boot-good-exec

Description

Use this command to configure a URL for a CLI script to **exec** following the success of a bootup configuration.

Related Commands

exec - This command executes the contents of a text file as if they were CLI commands entered at the console.

Default

no boot-good-exec

Parameters

file-url

Specifies the location and name of the file executed following successful completion of the bootup configuration file execution. When this parameter is not specified, no CLI script file is executed.

Values

file url	*local-url* | *remote-url*	255 chars max
local-url	[*cflash-id*/][*file-path*]	
remote-url	[{ftp://} *login:pswd@remote-locn*/][*file-path*]	
	remote-locn	[*hostname* | *ipv4-address* | [*ipv6- address*]]
	ipv4-address	*a.b.c.d*

	ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*] *x:x:x:x:x:x:d.d.d.d*[-*interface*] *x* - [0 to FFFF]H *d* - [0 to 255]D *interface* - 32 chars max, for link local addresses
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:	

Platforms

7705 SAR Gen 2

6.60 boot-timer

boot-timer

Syntax

boot-timer *interval*

no boot-timer

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ep boot-timer)

Full Context

configure redundancy multi-chassis peer mc-endpoint boot-timer

Description

This command configures the boot timer interval. This command applies only when the node reboots. It specifies the time the MC-EP protocol keeps trying to establish a connection before assuming a failure of the remote peer. This is different from the keep-alive mechanism which is used just after the peer-peer communication was established. After this time interval passed all the mc-endpoints configured under services will revert to single chassis behavior, activating the best local PW.

The **no** form of this command sets the interval to default.

Default

no boot-timer

Parameters

interval

Specifies the boot timer interval.

Values 1 to 600

Platforms

7705 SAR Gen 2

boot-timer

Syntax

boot-timer *seconds*

no boot-timer

Context

[Tree] (config>service>vpls>site boot-timer)

Full Context

configure service vpls site boot-timer

Description

This command configures for how long the service manager waits after a node reboot before running the DF election algorithm. The boot-timer value should be configured to allow for the BGP sessions to come up and for the NLRI information to be refreshed/exchanged.

The **no** form of this command reverts the default.

Default

boot-timer 10

Parameters

seconds

Specifies the site boot-timer in seconds.

Values 0 to 100

Platforms

7705 SAR Gen 2

boot-timer

Syntax

boot-timer *seconds*

no boot-timer

Context

[Tree] (config>service>epipe>site boot-timer)

Full Context

configure service epipe site boot-timer

Description

This command configures for how long the service manager waits after a node reboot before running the DF election algorithm. The boot-timer value should be configured to allow for the BGP sessions to come up and for the NLRI information to be refreshed/exchanged.

The **no** form of this command reverts the default.

Default

boot-timer 10

Parameters

seconds

Specifies the site boot-timer in seconds.

Values 0 to 600

Platforms

7705 SAR Gen 2

boot-timer

Syntax

boot-timer *secs*

no boot-timer

Context

[Tree] (config>service>pw-routing boot-timer)

Full Context

configure service pw-routing boot-timer

Description

This command configures a hold-off timer for MS-PW routing advertisements and signaling and is used at boot time.

The **no** form of this command removes a previously configured timer and restores it to its default.

Default

no boot-timer

Parameters

timer-value

Specifies the value of the boot timer in seconds.

Values 0 to 600

Platforms

7705 SAR Gen 2

boot-timer

Syntax

boot-timer *seconds*

no boot-timer

Context

[Tree] (config>redundancy>bgp-multi-homing boot-timer)

Full Context

configure redundancy bgp-multi-homing boot-timer

Description

This command configures the time the service manager waits after a node reboot before running the DF election algorithm. The **boot-timer** value should be configured to allow for the BGP sessions to come up and for the NLRI information to be refreshed or exchanged.

The **no** form of the command reverts the default.

Default

no boot-timer

Parameters

seconds

Specifies the BGP multi-homing boot-timer in seconds.

Values 0 to 600

Platforms

7705 SAR Gen 2

6.61 bootstrap-export

bootstrap-export

Syntax

bootstrap-export *policy-name* [*policy-name*]

no bootstrap-export

Context

[Tree] (config>service>vprn>pim>rp bootstrap-export)

Full Context

configure service vprn pim rp bootstrap-export

Description

This command exports policies to control the flow of bootstrap messages from the RP. Up to five policies can be defined.

The **no** form of this command removes the specified policy names from the configuration.

Parameters

policy-name

Specifies up to five policy names. The policy statement must already be configured in the **config>router>policy-options** context.

Platforms

7705 SAR Gen 2

bootstrap-export

Syntax

bootstrap-export *policy-name* [*policy-name*]

no bootstrap-export

Context

[Tree] (config>router>pim>rp bootstrap-export)

Full Context

configure router pim rp bootstrap-export

Description

This command applies export policies to control the flow of bootstrap messages from the RP, and apply them to the PIM configuration.

The **no** form of this command removes the policy name from the PIM RP configuration.

Default

no bootstrap-export

Parameters

policy-name

Specifies up to five export policy names, up to 32 characters.

Platforms

7705 SAR Gen 2

6.62 bootstrap-import

bootstrap-import

Syntax

bootstrap-import *policy-name* [*policy-name* **... up to five**]

no bootstrap-import *policy-name* [*policy-name* **... up to five**]

Context

[Tree] (config>service>vprn>pim>rp bootstrap-import)

Full Context

configure service vprn pim rp bootstrap-import

Description

This command imports policies to control the flow of bootstrap messages into the RP. Up to five policies can be defined.

The **no** form of this command removes the specified policy names from the configuration.

Parameters

policy-name

Specifies the policy name. The policy statement must already be configured in the config>router>policy-options context.

Platforms

7705 SAR Gen 2

bootstrap-import

Syntax

bootstrap-import *policy-name* [*policy-name*]

no bootstrap-import

Context

[Tree] (config>router>pim>rp bootstrap-import)

Full Context

configure router pim rp bootstrap-import

Description

This command applies import policies to control the flow of bootstrap messages to the RP, and apply them to the PIM configuration.

The **no** form of this command removes the policy name from the

Default

no bootstrap-import

Parameters

policy-name

Specifies up to five import policy names, up to 32 characters.

Platforms

7705 SAR Gen 2

6.63 boundary-type

boundary-type

Syntax

boundary-type {**clock-aligned** | **test-relative**}

no boundary-type

Context

[Tree] (config>oam-pm>session>meas-interval boundary-type)

Full Context

configure oam-pm session meas-interval boundary-type

Description

This command establishes the alignment of the start of the measurement interval with either the time of day clock or the start of the test. Alignment with the time of day clock always defaults to the representative top of the hour. Clock-aligned 15-minute measurement intervals divide the hour into four equal sections 00, 15, 30, 45. Clock-aligned 1-hour measurement intervals start at 00. Clock-aligned 1-day measurement intervals start at midnight. Test relative start times launches the measurement interval when the individual test enters the active (**no shutdown**) state. It is typical for the first measurement interval of a clock-aligned test to have the suspect flag set to yes because it is unlikely the **no shutdown** exactly corresponds to the clock based measurement interval start time. Clock-aligned measurement intervals can include an additional offset.

The **no** form of this command sets the boundary to the default clock-aligned.

Default

boundary-type clock-aligned

Parameters

clock-aligned

Aligns the start of the measurement interval with the time of day clock.

test-relative

Aligns the start of the measurement interval with the start of the test.

Platforms

7705 SAR Gen 2

6.64 bpdu

bpdu

Syntax

[**no**] **bpdu**

Context

[Tree] (debug>service>id>stp bpdu)

Full Context

debug service id stp bpdu

Description

This command enables STP debugging for received and transmitted BPDUs.

Platforms

7705 SAR Gen 2

bpdu

Syntax

[**no**] **bpdu**

Context

[Tree] (debug>service>id>stp bpdu)

Full Context

debug service id stp bpdu

Description

This command enables STP debugging for received and transmitted BPDUs.

The **no** form of the command disables debugging.

Platforms

7705 SAR Gen 2

6.65 bpdu-translation

bpdu-translation

Syntax

bpdu-translation {**auto** | **auto-rw** | **pvst** | **pvst-rw** | **stp**}

no bpdu-translation

Context

[Tree] (config>service>vpls>spoke-sdp bpdu-translation)

[Tree] (config>service>vpls>sap bpdu-translation)

Full Context

configure service vpls spoke-sdp bpdu-translation

configure service vpls sap bpdu-translation

Description

This command enables the translation of BPDUs to a specified format, meaning that all BPDUs transmitted on a specified SAP or spoke-SDP will have a specified format.

The **no** form of this command reverts to the default.

Default

no bpdu-translation

Parameters

auto

Specifies that appropriate format will be detected automatically, based on type of BPDUs received on such port.

auto-rw

Specifies that appropriate format will be detected automatically and the VLAN ID will be rewritten as follows:

- BPDU sent on egress of dot1q SAP will contain the VLAN ID of the SAP in BPDU-PVID TLV
- BPDU sent on egress of default QinQ SAP will contain the outer VLAN ID of the SAP in BPDU-PVID TLV
- BPDU sent on egress of QinQ SAP will contain the inner VLAN ID of the SAP in BPDU-PVID TLV

pvst

Specifies the BPDU-format as PVST. Note: the correct VLAN tag is included in the payload (depending on encapsulation value of outgoing SAP).

pvst-rw

Specifies the BPDU-format as PVST. The VLAN ID will be rewritten as follows:

- BPDU sent on egress of dot1q SAP will contain the VLAN ID of the SAP in BPDU-PVID TLV
- BPDU sent on egress of default QinQ SAP will contain the outer VLAN ID of the SAP in BPDU-PVID TLV
- BPDU sent on egress of QinQ SAP will contain the inner VLAN ID of the SAP in BPDU-PVID TLV

stp

Specifies the BPDU-format as STP.

Platforms

7705 SAR Gen 2

6.66 breakout

breakout

Syntax

breakout *breakout*

no breakout

Context

[Tree] (config>port>connector breakout)

Full Context

configure port connector breakout

Description

This command configures the transceiver port breakout for use in the connector. Specifying the breakout type triggers the creation of accessible ports for the connector.

When a QSFP28 connector uses an SFP+ optical module with the QSFP28-to-SFP+/SFP28 adapter, the user should set the *breakout* parameter to **c1-10g**, which indicates the presence of this adapter.

The options for breakout on specific connectors depend on both the card type and level (or XMA type and level). See the applicable installation guides for more information.

For some connectors (such as QSFPDD), there can be overlap in the breakout for different host interfaces. The same port breakout can be supported on an optical module that uses a host interface of CAUI-4 as another optical module that uses 100GAUI-2. To distinguish from the CAUI-4 host interface, the "-aui2" suffix is used on some breakout options. This is only necessary where there is overlap. In other situations, SR OS sets the host interface correctly without requiring the distinction in the breakout option.

The **no** form of this command removes the ports under the connector.

Default

no breakout

Parameters

breakout

Specifies the breakout type.

Values c1-40g, c4-10g, c1-100g, c4-25g, c10-10g, c1-400g, c2-100g, c4-100g, c1-10g, c1-25g, c1-50g, c8-50g, c1-800g, c3-100g, c8-100g, c2-400g, c1-100g-aui2, c2-100g-aui2, c1-400g-aui4, c4-100g-aui1

Platforms

7705 SAR Gen 2

6.67 broadcast

broadcast

Syntax

broadcast {**interface** *ip-int-name*} [**key-id** *key-id* | **authentication-keychain** *keychain-name*] [**version** *version*] [**ttl** *ttl*]

no broadcast {**interface** *ip-int-name*}

Context

[Tree] (config>service>vprn>ntp broadcast)

Full Context

configure service vprn ntp broadcast

Description

This command configures the node to transmit NTP packets on a given interface. Broadcast and multicast messages can easily be spoofed, therefore, authentication is strongly recommended.

The **no** form of this command removes the address from the configuration.

Parameters

ip-int-name

Specifies the local interface on which to transmit NTP broadcast packets. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

Values 32 character maximum

key-id

Identifies the configured authentication key and authentication type used by this node to receive and transmit NTP packets to and from an NTP server and peers. If an NTP packet is received by this node both authentication key and authentication type must be valid otherwise the packet is rejected and an event/trap generated.

Values 1 to 255

keychain-name

Identifies the keychain name, up to 32 characters.

version

Specifies the NTP version number that is generated by this node. This parameter does not need to be configured when in client mode in which case all versions are accepted.

Values 2 to 4

Default 4

ttl

Specifies the IP Time To Live (TTL) value.

Values 1 to 255

Platforms

7705 SAR Gen 2

broadcast

Syntax

broadcast [**router** *router-name*] {**interface** *ip-int-name*} [**key-id** *key-id* | **authentication-keychain** *keychain-name*] [**version** *version*] [**ttl** *ttl*]

no broadcast [**router** *router-name*] {**interface** *ip-int-name*}

Context

[Tree] (config>system>time>ntp broadcast)

Full Context

configure system time ntp broadcast

Description

This command configures the node to transmit NTP packets on a given interface. Broadcast and multicast messages can easily be spoofed, therefore, authentication is strongly recommended.

The **no** form of this command removes the address from the configuration.

Parameters

router-name

Specifies the router name used to transmit NTP packets. Base is the default. Select management to use the management port (Ethernet port on the CPM). Note that broadcast server capability can also be enabled on an interface within a VPRN context. Refer to "NTP Within a VPRN Service" in the *7705 SAR Gen 2 Layer 3 Services Guide: IES and VPRN* for more information.

Values Base | Management

Default Base

ip-int-name

Specifies the local interface on which to transmit NTP broadcast packets, up to 32 characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

key-id

Identifies the configured authentication key and authentication type used by this node to receive and transmit NTP packets to and from an NTP server and peers. If an NTP packet is received by this node both authentication key and authentication type must be valid otherwise the packet is rejected and an event or trap generated.

Values 1 to 255

keychain-name

Identifies the keychain name, up to 32 characters.

version

Specifies the NTP version number that is generated by this node. This parameter does not need to be configured when in client mode in which case all versions are accepted.

Values 2 to 4

Default 4

ttl

Specifies the IP Time To Live (TTL) value.

Values 1 to 255

Platforms

7705 SAR Gen 2

6.68 broadcast-client

broadcast-client

Syntax

[**no**] **broadcast-client**

Context

[Tree] (config>system>time>sntp broadcast-client)

Full Context

configure system time sntp broadcast-client

Description

This command enables listening to SNTP/NTP broadcast messages on interfaces with **broadcast client** enabled at global device level.

SNTP must be shutdown prior to changing either to or from broadcast mode.

The **no** form of the command disables broadcast client mode.

Default

no broadcast-client

Platforms

7705 SAR Gen 2

6.69 broadcast-policer

broadcast-policer

Syntax

broadcast-policer *policer-id* [**fp-redirect-group**]

no broadcast-policer

Context

[Tree] (config>qos>sap-ingress>fc broadcast-policer)

Full Context

configure qos sap-ingress fc broadcast-policer

Description

Within a **sap-ingress** QoS policy forwarding class context, the **broadcast-policer** command is used to map packets that match the forwarding class and are considered broadcast in nature to the specified policer-id. The specified policer-id must already exist within the **sap-ingress** QoS policy. While the system is determining the forwarding class of a packet, it is also looking up its forwarding destination based on the ingress service type and the service instance forwarding records. If the service type is VPLS and the destination MAC address is the broadcast address (ff:ff:ff:ff:ff:ff), the packet is classified into the broadcast forwarding type.

Broadcast forwarding type packets are mapped to either an ingress multipoint queue (using the **broadcast** *queue-id* or **broadcast** *queue-id* **group** *ingress-queue-group* commands) or an ingress policer (**broadcast-policer** *policer-id*). The **broadcast** and **broadcast-policer** commands within the forwarding class context are mutually exclusive. By default, the broadcast forwarding type is mapped to the SAP ingress default multipoint queue. If the **broadcast-policer** *policer-id* command is executed, any previous policer mapping or queue mapping for the broadcast forwarding type within the forwarding class is overridden if the policer mapping is successful.

A policer defined within the **sap-ingress** policy is not actually created on an ingress SAP or a subscriber using an **sla-profile** where the policy is applied until at least one forwarding type (unicast, broadcast, unknown, or multicast) from one of the forwarding classes is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or multiservice site, or ingress policing is not supported on the port associated with the SAP or subscriber or multiservice site, the initial forwarding class forwarding type mapping will fail.

The **broadcast-policer** command is ignored for instances of the policer applied to SAPs or subscribers' multiservice site where broadcast packets are not supported.

When the broadcast forwarding type within a forwarding class is mapped to a policer, the broadcast packets classified to the subclasses within the forwarding class are also mapped to the policer.

The **no** form of this command is used to restore the mapping of the broadcast forwarding type within the forwarding class to the default multipoint queue. If all forwarding class forwarding types had been removed from the default multipoint queue, the queue will not exist on the SAPs or subscribers or multiservice site associated with the QoS policy and the **no broadcast-policer** command will cause the system to attempt to create the default multipoint queue on each object. If the system cannot create the queue on each instance, the **no broadcast-policer** command will fail and the broadcast forwarding type within the forwarding class will continue its mapping to the existing policer-id. If the **no broadcast-policer** command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscriber will be lost.

Parameters

policer-id

When the forwarding class **broadcast-policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the sap-ingress QoS policy.

Values 1 to 63

fp-redirect-group

Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

Platforms

7705 SAR Gen 2

6.70 broadcast-queue

broadcast-queue

Syntax

broadcast-queue *queue-id* [**group** *queue-group-name*]

no broadcast queue

Context

[Tree] (config>qos>sap-ingress>fc broadcast-queue)

Full Context

configure qos sap-ingress fc broadcast-queue

Description

This command overrides the default broadcast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. When the forwarding class mapping is executed, all broadcast traffic on a SAP using this policy will be forwarded using the *queue-id*.

The broadcast forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.

The **no** form of this command sets the broadcast forwarding type *queue-id* back to the default of tracking the multicast forwarding type queue mapping.

Parameters

queue-id

The *queue-id* parameter must be an existing, multipoint queue defined in the config>qos>sap-ingress context.

Values Any valid multipoint queue ID in the policy including 2 through 32.

Default 11

group *queue-group-name*

This optional parameter is used to redirect the forwarding type within the forwarding class to the specified queue-id within the queue-group-name. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the specified queue group. The *queue-group-name* are configured in the **config>qos>queue-group-templates** egress and ingress contexts.

Platforms

7705 SAR Gen 2

6.71 broadcastclient

broadcastclient

Syntax

broadcastclient [**router** *router-instance* | **service-name** *service-name*] {**interface** *ip-int-name*} [**authenticate**]

no broadcastclient [**router** *router-instance* | **service-name** *service-name*] {**interface** *ip-int-name*}

Context

[Tree] (config>system>time>ntp broadcastclient)

Full Context

configure system time ntp broadcastclient

Description

When configuring NTP, the node can be configured to receive broadcast packets on a specified subnet. This command configures a specific interface to listen for broadcast NTP messages. The interface may exist within a VPRN service.

Broadcast and multicast messages can easily be spoofed, so authentication is strongly recommended. If broadcast is not configured, then any received NTP broadcast traffic will be ignored. Use the **show** command to view the state of the configuration.

The **no** form of this command removes the interface from the configuration.

Parameters

router-instance

Specifies the routing context that contains the interface in the form of *router-name* or *service-id*.

Values *router-name* — Base | Management

service-id — 1 to 2147483647

Default Base

service name

Specifies the service name for the VPRN. The name can be up to 64 characters in length. Note that CPM routing instances are not supported.

ip-int-name

Specifies the VPRN interface on which to receive NTP broadcast packets. If the string contains special characters (such as #, $, or spaces) the entire string must be enclosed within double quotes.

authenticate

Specifies whether or not to require authentication of NTP PDUs. When enabled, NTP PDUs are authenticated upon receipt.

Platforms

7705 SAR Gen 2

6.72 bsm-check-rtr-alert

bsm-check-rtr-alert

Syntax

[**no**] **bsm-check-rtr-alert**

Context

[Tree] (config>service>vprn>pim>if bsm-check-rtr-alert)

Full Context

configure service vprn pim interface bsm-check-rtr-alert

Description

This command enables the checking of router alert option in the bootstrap messages received on this interface.

Default

no bsm-check-rtr-alert

Platforms

7705 SAR Gen 2

bsm-check-rtr-alert

Syntax

[**no**] **bsm-check-rtr-alert**

Context

[Tree] (config>router>pim>interface bsm-check-rtr-alert)

Full Context

configure router pim interface bsm-check-rtr-alert

Description

This command enables the checking of the router alert option in the bootstrap messages received on this interface.

The **no** form of this command disables accepting BSM packets without the router alert option.

Default

no bsm-check-rtr-alert

Platforms

7705 SAR Gen 2

6.73 bsr

bsr

Syntax

bsr [**detail**]

no bsr

Context

[Tree] (debug>router>pim bsr)

Full Context

debug router pim bsr

Description

This command enables/disables debugging for the PIM bootstrap mechanism.

The **no** form of the command disables debugging.

Parameters

detail

Debugs detailed information on the PIM bootstrap mechanism.

Platforms

7705 SAR Gen 2

6.74 bsr-candidate

bsr-candidate

Syntax

bsr-candidate

Context

[Tree] (config>service>vprn>pim>rp>ipv6 bsr-candidate)

[Tree] (config>service>vprn>pim>rp bsr-candidate)

Full Context

configure service vprn pim rp ipv6 bsr-candidate

configure service vprn pim rp bsr-candidate

Description

Commands in this context configure Candidate Bootstrap (BSR) parameters.

Either **bsr-candidate** for IPv4 or **auto-rp-discovery** can be configured; the two mechanisms cannot be enabled together. **bsr-candidate** for IPv6 and **auto-rp-discovery** for IPv4 can be enabled together.

The **no** form of this command disables BSR.

Default

no bsr-candidate

Platforms

7705 SAR Gen 2

bsr-candidate

Syntax

bsr-candidate

Context

[Tree] (config>router>pim>rp>ipv6 bsr-candidate)

[Tree] (config>router>pim>rp bsr-candidate)

Full Context

configure router pim rp ipv6 bsr-candidate

configure router pim rp bsr-candidate

Description

Commands in this context configure Candidate Bootstrap (BSR) parameters.

Either **bsr-candidate** for IPv4 or **auto-rp-discovery** can be configured; the two mechanisms cannot be enabled together. **bsr-candidate** for IPv6 and **auto-rp-discovery** for IPv4 can be enabled together.

Default

bsr-candidate shutdown

Platforms

7705 SAR Gen 2

6.75 buffer-unresolved-packets

buffer-unresolved-packets

Syntax

[**no**] **buffer-unresolved-packets**

Context

[Tree] (config>system>ip buffer-unresolved-packets)

Full Context

configure system ip buffer-unresolved-packets

Description

This command configures the buffering of unresolved IPv4 and IPv6 packets waiting for an address resolution process (ARP) or neighbor discovery (ND) reply.

The **no** form of this command configures the system to discard IPv4 and IPv6 traffic needing a destination resolution that is buffered while waiting for a response to avoid any potential of out-of-order delivery of packets to the resolved destination. As a result, after the ARP or ND entry is populated, the system delivers only newly received packets in order.

Default

buffer-unresolved-packets

Platforms

7705 SAR Gen 2

6.76 bundle

bundle

Syntax

bundle [**detail**]

no bundle

Context

[Tree] (debug>router>rsvp>packet bundle)

Full Context

debug router rsvp packet bundle

Description

This command debugs bundle events.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about bundle events.

Platforms

7705 SAR Gen 2

6.77 burst-limit

burst-limit

Syntax

burst-limit {**default** | *size* [**bytes** | **kilobytes**]}

no burst-limit

Context

[Tree] (config>service>vprn>if>sap>egress>queue-override>queue burst-limit)

Full Context

configure service vprn interface sap egress queue-override queue burst-limit

Description

The queue **burst-limit** command defines an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue shaping rate.

The **no** form of this command restores the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies. When specified within a queue-override queue context, any current burst limit override for the queue is removed and the queue burst limit is controlled by its defining policy.

Default

no burst-limit

Parameters

default

Reverts the queues burst limit to the system default value.

size

When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and, by default, is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the **bytes** qualifier must be added following size.

Values 0 to 13671 kilobytes

0 to 14000000 bytes

Default No default for *size*; use the **default** keyword to specify default burst limit.

bytes

Specifies that the value given for *size* must be interpreted as the burst limit in bytes.

kilobytes

Specifies that the value given for *size* must be interpreted as the burst limit in kilobytes. If neither bytes, nor kilobytes is specified, the default qualifier is **kilobytes**.

Platforms

7705 SAR Gen 2

burst-limit

Syntax

burst-limit {**default** | *size* [**bytes** | **kilobytes**]}

no burst-limit

Context

[Tree] (config>port>ethernet>access>egr>qgrp>qover>q burst-limit)

Full Context

configure port ethernet access egress queue-group queue-overrides queue burst-limit

Description

The queue **burst-limit** command overrides the shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate.

The **no** form of this command removes the current burst limit override for the queue. The queue's burst limit is controlled by its defining template.

Default

no burst-limit

Parameters

default

Reverts the queue's burst limit to the system default value.

size

When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and, by default, is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the **bytes** qualifier must be added following size.

Values 1 to 13671 kilobytes

1 to 14000000 bytes

Default No default for *size*; use the **default** keyword to specify default burst limit.

bytes

Specifies that the value given for *size* must be interpreted as the burst limit in bytes.

kilobytes

Specifies that the value given for *size* must be interpreted as the burst limit in kilobytes. If neither **bytes** nor **kilobytes** is specified, the default qualifier is **kilobytes**.

Platforms

7705 SAR Gen 2

burst-limit

Syntax

burst-limit *size* [**bytes** | **kilobytes**]

no burst-limit

Context

[Tree] (config>service>epipe>sap>egress>queue-override>queue burst-limit)

Full Context

configure service epipe sap egress queue-override queue burst-limit

Description

The queue **burst-limit** command defines an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate.

The **no** form of this command restores the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies. When specified within a queue-override queue context, any current burst limit override for the queue is removed and the queue's burst limit is controlled by its defining policy.

Default

no burst-limit

Parameters

default

Reverts the queue's burst limit to the system default value.

size

When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and, by default, is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the **bytes** qualifier must be added following size.

Values 1 to 13671 kilobytes

14000000 bytes

Default No default for *size*; use the **default** keyword to specify default burst limit.

bytes

Specifies that the value given for *size* must be interpreted as the burst limit in bytes.

kilobytes

Specifies that the value given for *size* must be interpreted as the burst limit in kilobytes. If neither bytes nor kilobytes is specified, the default qualifier is **kilobytes**.

Platforms

7705 SAR Gen 2

burst-limit

Syntax

burst-limit {**default** | *size* [**bytes** | **kilobytes**]}

no burst-limit

Context

[Tree] (config>service>vpls>sap>egress>queue-override>queue burst-limit)

Full Context

configure service vpls sap egress queue-override queue burst-limit

Description

The queue **burst-limit** command defines an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate. The **no** form of this command restores the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies. When specified within a queue-override

queue context, any current burst limit override for the queue is removed and the queue's burst limit is controlled by its defining policy.

Default

no burst-limit

Parameters

default

Reverts the queue's burst limit to the system default value.

size

When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and by default is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the **bytes** qualifier must be added following size.

Values 1 to 13671 kilobytes

1 to 14000000 bytes

Default No default for *size*; use the **default** keyword to specify default burst limit.

bytes

Specifies that the value given for size must be interpreted as the burst limit in bytes.

kilobytes

Specifies that the value given for size must be interpreted as the burst limit in kilobytes. If neither bytes nor kilobytes is specified, the default qualifier is **kilobytes**.

Platforms

7705 SAR Gen 2

burst-limit

Syntax

burst-limit {**default** | *size* [**bytes** | **kilobytes**]}

no burst-limit

Context

[Tree] (config>service>ies>if>sap>egress>queue-override>queue burst-limit)

Full Context

configure service ies interface sap egress queue-override queue burst-limit

Description

The queue **burst-limit** command defines an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queues shaping rate.

The **no** form of this command restores the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies. When specified within a queue-override queue context, any current burst limit override for the queue is removed and the queue's burst limit is controlled by its defining policy.

Default

no burst-limit

Parameters

default

Reverts the queues burst limit to the system default value.

size

When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and, by default, is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the **bytes** qualifier must be added following size.

Values 0 to 13671 kilobytes

0 to or 14000000 bytes

Default No default for *size*; use the **default** keyword to specify default burst limit.

bytes

Specifies that the value given for *size* must be interpreted as the burst limit in bytes.

kilobytes

Specifies that the value given for *size* must be interpreted as the burst limit in kilobytes. If neither bytes, nor kilobytes is specified, the default qualifier is **kilobytes**.

Platforms

7705 SAR Gen 2

burst-limit

Syntax

burst-limit *size* [**bytes** | **kilobytes**]

no burst-limit

Context

[Tree] (config>qos>sap-ingress>queue burst-limit)

Full Context

configure qos sap-ingress queue burst-limit

Description

The **queue burst-limit** command is used to define an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate.

The **burst-limit** command is supported under the sap-ingress and sap-egress QoS policy queues. The command is also supported under the ingress and egress queue-group-templates queues.

The **no** form of this command is used to restore the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies or queue group templates. When specified within a **queue-override queue** context, any current burst limit override for the queue will be removed and the queue's burst limit will be controlled by its defining policy or template.

Default

no burst-limit

Parameters

size

Specifies an explicit burst limit size. The value is expressed as an integer and, by default, is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the byte qualifier must be added following size.

Values 1 to 13,671 kbytes or 14,000,000 bytes

Default No default for size; use the default keyword to specify default burst limit.

bytes

Specifies that the value given for size must be interpreted as the burst limit in bytes.

kilobytes

Specifies that the value given for size must be interpreted as the burst limit in kilobytes. If neither bytes nor kilobytes is specified, the default qualifier is kilobytes.

Platforms

7705 SAR Gen 2

burst-limit

Syntax

burst-limit {**default** | *size* [**bytes** | **kilobytes**]}

burst-limit delay-time *microseconds*

no burst-limit

Context

[Tree] (config>qos>sap-egress>queue burst-limit)

Full Context

configure qos sap-egress queue burst-limit

Description

The **queue burst-limit** command configures an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate.

The **delay-time** command option configures the burst time as a function of the expected delay. The system automatically translates this configuration into kilobytes based on the administrative rate of the queue parent (for example, the port, scheduler, or aggregate-shaper).

The **no** form of this command restores the default burst limit to the specified queue. This is equivalent to specifying **burst-limit default** within the QoS policies or queue group templates.

Default

no burst-limit

Parameters

default

Reverts the burst limit of the queue to the system default value.

size

Specifies an explicit burst limit size. The value is expressed as an integer and, by default, is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the **bytes** keyword must be added following *size*.

Values 1 to 13671 kilobytes

1 to 14,000,000 bytes

bytes

Specifies that the configured size value must be interpreted as the burst limit in bytes.

kilobytes

Specifies that the configure size value must be interpreted as the burst limit in kilobytes. If neither bytes nor kilobytes is specified, the default qualifier is **kilobytes**.

microseconds

Specifies the burst limit as a function of delay time.

Values 0 to 1000000

Platforms

7705 SAR Gen 2

burst-limit

Syntax

burst-limit {*size* [**bytes** | **kilobytes**] | **default**}

no burst-limit

Context

[Tree] (config>qos>qgrps>egr>qgrp>queue burst-limit)

[Tree] (config>qos>qgrps>ing>qgrp>queue burst-limit)

Full Context

configure qos queue-group-templates egress queue-group queue burst-limit

configure qos queue-group-templates ingress queue-group queue burst-limit

Description

The queue burst-limit command is used to define an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate.

The burst-limit command is supported under the sap-ingress and sap-egress QoS policy queues. The command is also supported under the ingress and egress queue-group-templates queues.

The **no** form of this command is used to restore the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies or queue group templates. When specified within a queue-override queue context, any current burst limit override for the queue will be removed and the queue's burst limit will be controlled by its defining policy or template.

Parameters

size

When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and by default is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the byte qualifier must be added following size.

Values 1 to 14,000 (14,000 or 14,000,000 depending on bytes or kilobytes)

bytes

Specifies that the value given for size must be interpreted as the burst limit in bytes.

kilobytes

Specifies that the value given for size must be interpreted as the burst limit in kilobytes.

Platforms

7705 SAR Gen 2

6.78 bypass-resignal-timer

bypass-resignal-timer

Syntax

bypass-resignal-timer *minutes*

no bypass-resignal-timer

Context

[Tree] (config>router>mpls bypass-resignal-timer)

Full Context

configure router mpls bypass-resignal-timer

Description

This command triggers the periodic global re-optimization of all dynamic bypass LSP paths associated with RSVP P2P LSP. The operation is performed at each expiry of the user configurable bypass LSP resignal timer.

When this command is enabled, MPLS requests CSPF for the best path for each dynamic bypass LSP originated on this node. The constraints, hop-limit, SRLG and admin-group constraints, of the first associated LSP primary path that originally triggered the signaling of the bypass LSP must be satisfied. To do this, MPLS saves this initial Path State Block (PSB) of that LSP primary path, even if the latter is torn down.

CSPF first updates the SRLG membership of the current bypass LSP path and checks if the path violates the SRLG constraint of the initial PSB. It then attempts a new path computation for the bypass LSP using the initial PSB constraints. If CSPF returns no path or returns a new path with a cost that is lower than the current path, MPLS does not signal the new bypass path. If CSPF returns a new path with a cost that is lower than the current one, MPLS signals it. Also, if the new bypass path is SRLG strict disjoint with the primary path of the original PSB while the current path is SRLG loose disjoint, the manual bypass path is resignaled regardless of cost comparison.

Once the new path is successfully signaled, MPLS evaluates each PSB of each PLR (that is, each unique avoid-node or avoid-link constraint) associated with the current bypass LSP path to check if the corresponding LSP primary path constraints are still satisfied by the new bypass LSP path. If so, the PSB association is moved to the new bypass LSP.

Each PSB for which the constraints are not satisfied remains associated with the PLR on the current bypass LSP and is checked at the next timer or manual bypass re-optimization. Additionally, if SRLG FRR loose disjointness is configured using the **configure router mpls srlg-frr** command and the current bypass LSP is SRLG disjoint with a primary path while the new bypass LSP is not SRLG disjoint, the PSB association is not moved. When CSPF does not return a new bypass path or it returns a less optimal one, the PSBs remain associated with the current bypass path. However, it is possible that CSPF found the current bypass LSP path no longer satisfies the SRLG constraint of one or more PLRs after the update of the current path SRLG information. In that case, MPLS detaches from current bypass path the PSB associations of these PLRs. These orphaned PSBs are re-evaluated by the FRR background task which checks unprotected PSBs on a regular basis and following the same above procedure.

If a specific PLR associated with a bypass LSP is active, the corresponding PSBs remain associated with the current PLR until the Global Revertive Make-Before-Break (MBB) tears down all corresponding primary paths, which also causes the current PLR to be removed.

Note:
While it is in the preceding state, the older PLR does not get any new PSB association until the specific PLR with an active bypass LSP is removed. When the last PLR is removed, the older bypass LSP is torn down.

This feature is not supported with inter-area dynamic bypass LSP and bypass LSP protecting S2L paths of a P2MP LSP.

The **no** form of this command disables the periodic global re-optimization of dynamic bypass LSP paths.

Default

no bypass-resignal timer.

Parameters

minutes

Specifies the time, in minutes, MPLS waits before attempting to resignal dynamic bypass LSP paths originated on the system.

Values 1 to 10080

Platforms

7705 SAR Gen 2

7 c Commands

7.1 ca-name

ca-name

Syntax

ca-name *ca-name*

no ca-name

Context

[Tree] (config>port>ethernet>dot1x>macsec>sub-port ca-name)

Full Context

configure port ethernet dot1x macsec sub-port ca-name

Description

This command configures the Connectivity Association (CA) linked to this MACsec sub-port. The specified CA provides the MACsec parameter to be used or negotiated with other peers.

The **no** form of this command removes the CA from the MACsec sub-port.

Parameters

ca-name

Specifies the appropriate ca to be used under this MACsec sub-port, up to 32 characters.

Platforms

7705 SAR Gen 2

7.2 ca-profile

ca-profile

Syntax

[**no**] **ca-profile** *name*

Context

[Tree] (config>ipsec>cert-profile>entry>send-chain ca-profile)

Full Context

configure ipsec cert-profile entry send-chain ca-profile

Description

This command specifies a CA certificate in the specified **ca-profile** to be sent to the peer.

Multiple configurations (up to seven) of this command are allowed in the same entry.

Parameters

name

Specifies the profile name up to 32 characters.

Platforms

7705 SAR Gen 2

ca-profile

Syntax

ca-profile *name* [**create**]

no ca-profile *name*

Context

[Tree] (config>system>security>pki ca-profile)

Full Context

configure system security pki ca-profile

Description

This command creates a new **ca-profile** or enters the configuration context of an existing **ca-profile**. Up to 128 ca-profiles can be created in the system. A **shutdown** of the **ca-profile** will not affect the current up and running **ipsec-tunnel** or **ipsec-**gw that is associated with the **ca-profile**. However, authentication afterwards will fail with a **shutdown ca-profile**.

Executing a **no shutdown** command in this context causes the system to reload the configured cert-file and crl-file.

A **ca-profile** can be applied under the **ipsec-tunnel** or **ipsec-gw** configuration.

The **no** form of this command removes the name parameter from the configuration. A ca-profile cannot be removed until all the associated entities (ipsec-tunnel/gw) have been removed.

Parameters

name

Specifies the name of the **ca-profile** up to 32 characters.

create

Keyword used to create a new **ca-profile**. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

ca-profile

Syntax

[**no**] **ca-profile** *profile-name*

Context

[Tree] (debug>certificate>auto-crl-update ca-profile)

[Tree] (debug>certificate>cmpv2 ca-profile)

[Tree] (debug>certificate>ocsp ca-profile)

Full Context

debug certificate auto-crl-update ca-profile

debug certificate cmpv2 ca-profile

debug certificate ocsp ca-profile

Description

This command debugs output of the specified CA profile.

- Protection method of each message is logged.
- All HTTP messages are logged. Format allows offline analysis using Wireshark.
- In the event of failed transactions, saved certificates are not deleted from file system for further debug and analysis.
- The system allows CMPv2 debugging for multiple ca-profile at the same time.

Parameters

profile-name

Specifies the name of the CA profile, up to 32 characters.

Platforms

7705 SAR Gen 2

ca-profile

Syntax

[**no**] **ca-profile** *name*

Context

[Tree] (config>system>security>tls>cert-profile>entry>send-chain ca-profile)

Full Context

configure system security tls cert-profile entry send-chain ca-profile

Description

This command enables a certificate authority (CA) certificate in the specified CA profile to be sent to the peer. Up to seven configurations of this command are permitted in the same entry.

The **no** form of the command disables the transmission of a CA certificate from the specified CA profile.

Parameters

name

Specifies the name of the certificate authority profile, up to 32 characters in length.

Platforms

7705 SAR Gen 2

7.3 cacert

cacert

Syntax

cacert est-profile *name* **output** *output-cert-filename* [**force**]

Context

[Tree] (admin>certificate>est cacert)

Full Context

admin certificate est cacert

Description

This command downloads a Certificate Authority (CA) certificate from an EST server specified by the EST profile. The downloaded certificate is imported and saved with the filename specified by the *output-cert-filename*.

Parameters

name

Specifies the EST profile name, up to 32 characters

output-cert-filename

Specifies the filename of the resulting CA certificate, up to 200 characters

force

Overwrites the existing file with same filename

Platforms

7705 SAR Gen 2

7.4 cache-reset

cache-reset

Syntax

[**no**] **cache-reset**

Context

[Tree] (debug>router>rpki-session>packet cache-reset)

Full Context

debug router rpki-session packet cache-reset

Description

This command enables debugging for cache reset RPKI packets.

The **no** form of this command disables debugging for cache reset RPKI packets.

Platforms

7705 SAR Gen 2

7.5 cache-response

cache-response

Syntax

[**no**] **cache-response**

Context

[Tree] (debug>router>rpki-session>packet cache-response)

Full Context

debug router rpki-session packet cache-response

Description

This command enables debugging for cache response RPKI packets.

The **no** form of this command disables debugging for cache response RPKI packets.

Platforms

7705 SAR Gen 2

7.6 cak

cak

Syntax

cak *hex-string* [**hash** | **hash2** | **custom**]

no cak

Context

[Tree] (config>macsec>conn-assoc>static-cak>pre-shared-key cak)

Full Context

configure macsec connectivity-association static-cak pre-shared-key cak

Description

Specifies the connectivity association key (CAK) for a pre-shared key. Two values are derived from CAK.

- Key Encryption Key (KEK), this is used to encrypt the MKA and SAK (symmetric key used for data path PDUs) to be distributed between all members.
- Integrity Check Value (ICK), this is used to authenticate the MKA and SAK PDUs to be distributed between all members.

The **no** form of this command removes the value.

Parameters

hex-string

Specifies the value of the CAK.

Values up to 64 hexadecimal characters, 32 hexadecimal characters for 128-bit key and 64 hexadecimal characters for 256-bit key

hash

Keyword, specifying the hash scheme.

hash2

Keyword, specifying the hash scheme.

custom

Specifies the custom encryption for management interface.

Platforms

7705 SAR Gen 2

7.7 called-station-id

called-station-id

Syntax

[**no**] **called-station-id**

Context

[Tree] (config>ipsec>rad-auth-plcy>include called-station-id)

[Tree] (config>ipsec>rad-acct-plcy>include called-station-id)

Full Context

configure ipsec radius-authentication-policy include-radius-attribute called-station-id

configure ipsec radius-accounting-policy include-radius-attribute called-station-id

Description

This command includes called station ID attributes.

The **no** form of this command excludes called station ID attributes.

Default

no called-station-id

Platforms

7705 SAR Gen 2

7.8 calling-station-id

calling-station-id

Syntax

[**no**] **calling-station-id**

Context

[Tree] (config>ipsec>rad-auth-plcy>include calling-station-id)

[Tree] (config>ipsec>rad-acct-plcy>include calling-station-id)

Full Context

configure ipsec radius-authentication-policy include-radius-attribute calling-station-id

configure ipsec radius-accounting-policy include-radius-attribute calling-station-id

Description

This command enables the inclusion of the **calling-station-id** attribute in RADIUS authentication requests and RADIUS accounting messages.

Default

no calling-station-id

Platforms

7705 SAR Gen 2

7.9 cancel-commit

cancel-commit

Syntax

[**no**] **cancel-commit**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization cancel-commit)

Full Context

configure system security profile netconf base-op-authorization cancel-commit

Description

This command enables the NETCONF <cancel-commit> RPC.

The **no** form of this command disables the RPC.

Default

no cancel-commit

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

7.10 candidate

candidate

Syntax

candidate

Context

[Tree] (candidate)

Full Context

candidate

Description

Commands in this context edit candidate configurations.

Commands in the **candidate** CLI branch, except **candidate edit**, are available only when in edit-cfg mode.

Platforms

7705 SAR Gen 2

candidate

Syntax

[**no**] **candidate**

Context

[Tree] (config>system>netconf>capabilities candidate)

Full Context

configure system netconf capabilities candidate

Description

This command allows the SR OS NETCONF server to access the candidate configuration datastore. Configuring this command also enables using **commit** and **discard-changes**.

When **configure system management-interface configuration-mode** is set to **classic**, the candidate capability is disabled, even if this command is configured.

The **no** form of the command disables the SR OS NETCONF server from accessing the candidate datastore. If the candidate is disabled, requests that reference the candidate datastore return an error, and when a NETCONF client establishes a new session, the candidate capability is not advertised in the SR OS NETCONF Hello message.

Default

candidate

Platforms

7705 SAR Gen 2

7.11 cannot-change-password

cannot-change-password

Syntax

[**no**] **cannot-change-password**

Context

[Tree] (config>system>security>user>console cannot-change-password)

Full Context

configure system security user console cannot-change-password

Description

This command allows a user the privilege to change their password for both FTP and console login.

To disable a user's privilege to change their password, use the **cannot-change-password** form of this command.

Note:
The **cannot-change-password** flag is not replicated when a user copy is performed. A new-password-at-login flag is created instead.

Default

no cannot-change-password

Platforms

7705 SAR Gen 2

7.12 capture

capture

Syntax

capture [{**start** | **stop**}]

Context

[Tree] (debug>pcap capture)

Full Context

debug pcap capture

Description

This command starts and stops the packet capture process for the specified *session-name*.

Parameters

start

Starts the packet capture process and also start or restarts the FTP or TFTP session. If the FTP or TFTP server is unreachable, the command prompt rejects further input until the retires are timed out after 24 seconds (after four attempts of about six seconds each). If the same file name is unchanged in the **config**>**mirror**>**mirror-dest**>**pcap** context between captures, this command overwrites the file content.

stop

Stops the packet capture process and also stops the FTP or TFTP session. If the FTP or TFTP server is unreachable, the command prompt rejects further input until the retires are timed out after 24 seconds (after four attempts of about six seconds each).

Platforms

7705 SAR Gen 2

7.13 card

card

Syntax

[**no**] **card** *slot-number*

Context

[Tree] (config card)

Full Context

configure card

Description

This mandatory command enables access to the chassis and context. In SR OS cards cover IOM, IMM, and XCM.

The **no** form of this command removes the card from the configuration. All associated ports, services, and MDAs must be shutdown.

Default

no card

Parameters

slot-number

Specifies the slot number of the card in the chassis. The maximum slot number is platform dependent. Refer to the hardware installation guides.

Values 1 to 10

Platforms

7705 SAR Gen 2

7.14 card-type

card-type

Syntax

card-type *card-type* [**level** *card-level*]

no card-type

Context

[Tree] (config>card card-type)

Full Context

configure card card-type

Description

This mandatory command adds an IOM/XCM to the device configuration for the slot. The card type can be preprovisioned, meaning that the card does not need to be installed in the chassis.

A card must be provisioned before an MDA, connector, or port can be configured.

A card can only be provisioned in a slot that is vacant, meaning no other card can be provisioned (configured) for that particular slot. To reconfigure a slot position, use the **no** form of this command to remove the current information.

A card can only be provisioned in a slot if the card type is allowed in the slot. An error message is generated if an attempt is made to provision a card type that is not allowed.

If a card is inserted that does not match the configured card type for the slot, then a log event and facility alarm is raised. The alarm is cleared when the correct card type is installed or the configuration is modified.

A log event and facility alarm are is raised if an administratively enabled card is removed from the chassis. The alarm is cleared when the correct card type is installed or the configuration is modified. A log event is issued when a card is removed that is administratively disabled.

Because IMMs do not have the capability to install separate MDAs, the configuration of the MDA is automatic. This configuration only includes the default parameters such as default buffer policies. Commands to manage the MDA such as **shutdown** and so on, remain in the MDA configuration context.

Some card hardware can support two different firmware loads. One load includes the base Ethernet functionality, including 10G WAN mode, but does not include 1588 port-based timestamping. The second load includes the base Ethernet functionality and 1588 port-based timestamping, but does not include 10G WAN mode. These are identified as two card types that are the same, except for a "-ptp” suffix to indicate the second loadset; for example, *imm40-10gb-sfp* and *imm40-10gb-sfp-ptp*. A hard reset of the card occurs when switching between the two provisioned types.

An appropriate alarm is raised if a partial or complete card failure is detected. The alarm is cleared when the error condition ceases.

New generations of cards include variants controlled by hardware and software licensing. For these cards, the license level must be provisioned in addition to the card type. A card cannot become operational unless the provisioned license level matches the license level of the card installed into the slot. The set of license levels varies by card type.

The provisioned level controls aspects related to connector provisioning and the consumption of hardware egress queues and egress policers. Changes to the provisioned license level may be blocked if configuration exists that would not be permitted with the new target license level.

If the license level is not specified, the level is set to the highest license level for that card.

The **no** form of this command removes the card from the configuration.

Default

no card-type

Parameters

card-type

Specifies the type of card to be configured and installed in that slot. Values for this attribute vary by platform and release. The release notes include a listing of all supported card-types and their CLI strings. In addition, the command can be queried to check which card-types are relevant for the active platform type. Some examples include iom4-e-b and imm-2pac-fp3.

card-level

Specifies the license level of the card, up to 32 characters. Possible values vary by card type.

Platforms

7705 SAR Gen 2

7.15 carrier-carrier-vpn

carrier-carrier-vpn

Syntax

[**no**] **carrier-carrier-vpn**

Context

[Tree] (config>service>vprn carrier-carrier-vpn)

Full Context

configure service vprn carrier-carrier-vpn

Description

This command configures a VPRN service to support a Carrier Supporting Carrier model. It should be configured on a network provider's CSC-PE device.

This command cannot be applied to a VPRN unless it has no SAP or spoke-SDP interfaces. Once this command has been entered one or more MPLS-capable CSC interfaces can be created in the VPRN.

The **no** form of this command removes the Carrier Supporting Carrier capability from a VPRN.

Default

no carrier-carrier-vpn

Platforms

7705 SAR Gen 2

7.16 cbs

cbs

Syntax

cbs *size-in-kbytes*

no cbs

Context

[Tree] (config>service>vpls>sap>ingress>queue-override>queue cbs)

[Tree] (config>service>ies>if>sap>ingress>queue-override>queue cbs)

[Tree] (config>service>ies>if>sap>egress>queue-override>queue cbs)

[Tree] (config>service>vpls>sap>egress>queue-override>queue cbs)

Full Context

configure service vpls sap ingress queue-override queue cbs

configure service ies interface sap ingress queue-override queue cbs

configure service ies interface sap egress queue-override queue cbs

configure service vpls sap egress queue-override queue cbs

Description

This command overrides specific attributes of the specified queue's CBS parameters.

It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS settings into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.

If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.

The **no** form of this command returns the CBS size to the default value.

Parameters

size-in-kbytes

Specifies the size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10 kbytes is desired, enter the value 10. A value of 0 specifies

that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

Values 0 to 1048576, default

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs *size-in-kbytes*

no cbs

Context

[Tree] (config>service>vprn>if>sap>ingress>queue-override>queue cbs)

[Tree] (config>service>vprn>if>sap>egress>queue-override>queue cbs)

Full Context

configure service vprn interface sap ingress queue-override queue cbs

configure service vprn interface sap egress queue-override queue cbs

Description

This command can be used to override specific attributes of the specified queue's CBS parameters.

It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS setting into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.

If the CBS value is larger than the MBS value, an error occurs, preventing the CBS change.

The **no** form of this command returns the CBS to the default value.

Default

no cbs

Parameters

size-in-kbytes

The size parameter is an integer expression of the number of kilobytes reserved for the queue. For a value of 10 kbytes, enter the number 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimum reserved size can be applied for scheduling purposes).

Values 0 to 131072 or default

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs {*size* [**bytes** | **kilobytes**] | **default**}

no cbs

Context

[Tree] (config>card>fp>ingress>network>qgrp>policer-over>plcr cbs)

[Tree] (config>card>fp>ingress>access>qgrp>policer-over>plcr cbs)

Full Context

configure card fp ingress network queue-group policer-override policer cbs

configure card fp ingress access queue-group policer-override policer cbs

Description

This command configures the policer's CIR leaky bucket's exceed threshold. The CIR bucket's exceed threshold represents the committed burst tolerance allowed by the policer. If the policer's forwarding rate is equal to or less than the policer's defined CIR, the CIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the forwarding rate increases beyond the profiling rate, the amount of data allowed to be in-profile above the rate is capped by the threshold.

The policer's **cbs** size defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command returns the policer to its default CBS size.

Parameters

size

Specifies that the *size* parameter is required when specifying **cbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **bytes** and **kilobytes** keywords are mutually exclusive and are used to explicitly define whether *size* represents bytes or kilobytes.

Values 0 to 2683435456

bytes

When **bytes** is defined, the value given for size is interpreted as the queue's CBS value specified in bytes.

kilobytes

When **kilobytes** is defined, the value is interpreted as the queue's CBS value given in kilobytes.

Default **kilobyte**

default

Specifying the keyword **default** sets the CBS to its default value.

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs *size-in-kbytes*

no cbs

Context

[Tree] (config>port>ethernet>access>egr>qgrp>qover>q cbs)

[Tree] (config>port>ethernet>access>ing>qgrp>qover>q cbs)

Full Context

configure port ethernet access egress queue-group queue-overrides queue cbs

configure port ethernet access ingress queue-group queue-overrides queue cbs

Description

This command defines the default committed buffer size for the template queue. Overall, the CBS command follows the same behavior and provisioning characteristics as the CBS command in the queue-group or network QoS policy. The exception is the addition of the cbs-value qualifier keywords bytes or kilobytes.

The **no** form of this command restores the default CBS size to the template queue.

Default

cbs default

Parameters

size-in-kbytes

The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10 kbytes is desired, enter the value 10. A value of 0 specifies that no

reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

Values 0 to 1048576 or default

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs *size* [**bytes** | **kilobytes**]

no cbs

Context

[Tree] (config>service>epipe>sap>ingress>policer-over>plcr cbs)

[Tree] (config>service>epipe>sap>egress>policer-over>plcr cbs)

Full Context

configure service epipe sap ingress policer-override policer cbs

configure service epipe sap egress policer-override policer cbs

Description

This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured CBS parameter for the specified *policer-id*.

The **no** form of this command returns the CBS size to the default value.

Default

no cbs

Parameters

size

The size parameter is required when specifying cbs override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 2683435456, **default**

bytes

When bytes is defined, the value given for *size* is interpreted as the policer's MBS value in bytes.

kilobytes

When kilobytes is defined, the value given for *size* is interpreted as the policer's MBS value in kilobytes.

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs {*size-in-kbytes* | **default**}

no cbs

Context

[Tree] (config>service>epipe>sap>ingress>queue-override>queue cbs)

[Tree] (config>service>epipe>sap>egress>queue-override>queue cbs)

Full Context

configure service epipe sap ingress queue-override queue cbs

configure service epipe sap egress queue-override queue cbs

Description

This command can be used to override specific attributes of the specified queue's CBS parameters.

It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a specific access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS setting into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly to drop packets.

The **no** form of this command returns the CBS size to the default value.

Default

no cbs

Parameters

size-in-kbytes

The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is wanted, enter the value 10. A value of 0 specifies that no

reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

Values 0 to 131072, default

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs *size* [{**bytes** | **kilobytes**}]

no cbs

Context

[Tree] (config>service>vpls>sap>ingress>policer-override>plcr cbs)

[Tree] (config>service>vpls>sap>egress>policer-override>plcr cbs)

Full Context

configure service vpls sap ingress policer-override policer cbs

configure service vpls sap egress policer-override policer cbs

Description

This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured CBS parameter for the specified policer-id.

The **no** form of this command returns the CBS size to the default value.

Default

no cbs

Parameters

size

This parameter is required when specifying CBS override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 2683435456, **default**

Default kilobytes

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs *size* [{**bytes** | **kilobytes**}]

no cbs

Context

[Tree] (config>service>ies>if>sap>ingress>policer-over>plcr cbs)

[Tree] (config>service>ies>if>sap>egress>policer-over>plcr cbs)

Full Context

configure service ies interface sap ingress policer-override policer cbs

configure service ies interface sap egress policer-override policer cbs

Description

This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured CBS parameter for the specified policer-id.

The **no** form of this command returns the CBS size to the default value.

Default

no cbs

Parameters

size

This parameter is required when specifying CBS override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 2683435456, **default**

Default kilobytes

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs *size* [{**bytes** | **kilobytes**}]

no cbs

Context

[Tree] (config>service>vprn>if>sap>ingress>policer-over>plcr cbs)

[Tree] (config>service>vprn>if>sap>egress>policer-over>plcr cbs)

Full Context

configure service vprn interface sap ingress policer-override policer cbs

configure service vprn interface sap egress policer-override policer cbs

Description

This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured CBS parameter for the specified policer-id.

The **no** form of this command returns the CBS size to the default value.

Default

no cbs

Parameters

size

This parameter is required when specifying CBS override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 2683435456, **default**

Default kilobytes

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs *size* [**bytes** | **kilobytes**]

no cbs

Context

[Tree] (config>qos>sap-egress>policer cbs)

[Tree] (config>qos>sap-ingress>policer cbs)

Full Context

configure qos sap-egress policer cbs

configure qos sap-ingress policer cbs

Description

This command configures the policer’s CIR leaky bucket’s exceed threshold. The CIR bucket’s exceed threshold represents the committed burst tolerance allowed by the policer. If the policer’s forwarding rate is equal to or less than the policer's defined CIR, the CIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the forwarding rate increases beyond the profiling rate, the amount of data allowed to be in-profile above the rate is capped by the threshold.

The policer’s **cbs** size defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command returns the policer to its default CBS size.

By default, the CBS is 16 Mbytes when CIR equals max or is greater than or equal to the FP capacity (this overrides an explicit configured CBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max CIR capped to 3968 kbytes, with a minimum of 256 bytes.

Parameters

size [bytes | kilobytes]

Specifies an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 2683435456, **default**

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs {*size-in-kbytes*| default}

cbs delay-time *microseconds*

cbs delay-percent *percent*

no cbs

Context

[Tree] (config>qos>sap-egress>queue cbs)

Full Context

configure qos sap-egress queue cbs

Description

This command provides a mechanism to override the default reserved buffers for the queue. It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a specific access

port egress buffer pool. Oversubscription may be desirable because of the potentially large number of service queues and the economy of statistical multiplexing the CBS settings of the individual into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues use their CBS buffers and the total-in-use exceeds the defined reserved total, essentially the buffers are removed from the shared portion of the pool without the shared in-use average and total counts being decremented. This can affect the operation of the high- and low-priority RED slopes on the pool, causing them to miscalculate when to start randomly dropping packets.

If the CBS value is larger than the MBS value, the CBS is capped to the value of the MBS or the minimum CBS value. If the MBS and CBS values are configured to be equal (or nearly equal), this will result in the CBS being slightly higher than the value configured.

The **delay-time** command option configures the CBS as a function of the expected delay. The system automatically translates this configuration into kilobytes based on the administrative rate of the queue parent (for example, the port, scheduler, or aggregate-shaper).

The **delay-percent** command option configures the CBS as percentage of the SAP delay budget of the queue configured using the **latency-budget** command.

The **no** form of this command returns the CBS size to the default value.

Default

cbs default

Parameters

size-in-kbytes

The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10 kbytes is required, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes). The CBS maximum value used is constrained by the pool size in which the queue exists.

Values 0 to 1048576

Minimum configurable non-zero value: 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4

Minimum non-zero default value: maximum of 10 ms of CIR, or 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4

microseconds

Specifies the CBS as a function of delay time.

Values 0 to 1000000

percent

Specifies the CBS as a percentage of the SAP latency budget.

Values 0.00 to 100.00

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs *size-in-kbytes*

no cbs

Context

[Tree] (config>qos>sap-ingress>queue cbs)

Full Context

configure qos sap-ingress queue cbs

Description

This command provides a mechanism to override the default reserved buffers for the queue. It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potentially large number of service queues and the economy of statistical multiplexing the individual queue's CBS settings into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high- and low-priority RED slopes on the pool, causing them to miscalculate when to start randomly dropping packets.

If the CBS value is larger than the MBS value, the CBS is capped to the value of the MBS or the minimum CBS value. If the MBS and CBS values are configured to be equal (or nearly equal), this will result in the CBS being slightly higher than the value configured.

The **no** form of this command returns the CBS size to the default value.

Default

cbs default

Parameters

size-in-kbytes

The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10 kbytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes) The CBS maximum value used is constrained by the pool size in which the queue exists.

Values 0 to 1048576 or default

Minimum configurable non-zero value: 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4

Minimum non-zero default value: maximum of 10 ms of CIR, or 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs *percent*

no cbs

Context

[Tree] (config>qos>network-queue>queue cbs)

Full Context

configure qos network-queue queue cbs

Description

The Committed Burst Size (**cbs**) command specifies the relative number of reserved buffers for a specific ingress network FP forwarding class queue or egress network port forwarding class queue. The value is entered as a percentage.

The CBS for a queue is used to determine whether it has exhausted its reserved buffers while enqueuing packets. When the queue has exceeded the number of buffers considered in reserve for this queue, it must contend with other queues for the available shared buffer space within the buffer pool. Access to this shared pool space is controlled through Random Early Detection (RED) slope application.

Two RED slopes are maintained in each buffer pool. A high-priority slope is used by in-profile packets. A low-priority slope is used by out-of-profile packets. At egress, there are two additional RED slopes maintained in each buffer pool: the highplus slope is used by inplus-profile packets, and the exceed slope is used by exceed-profile packets. All network control and management packets are considered in-profile. Assured packets are handled by their in-profile and out-of-profile markings. All best-effort packets are considered out-of-profile. Premium queues should be configured such that the CBS percent is sufficient to prevent shared buffering of packets. This is generally taken care of by the CIR scheduling of premium queues and the overall small amount of traffic on the class. Premium queues in a properly designed system will drain before all others, limiting their buffer utilization.

The RED slopes will detect congestion conditions and work to discard packets and slow down random TCP session flows through the queue. The RED slope definitions can be defined, modified, or disabled through the slope policy assigned to the FP for the network ingress buffer pool or assigned to the network port for network egress buffer pools.

The resultant CBS size can be larger than the MBS. This will result in a portion of the CBS for the queue to be unused and should be avoided.

The **no** form of this command returns the CBS size for the queue to the default for the forwarding class.

Default

The cbs forwarding class defaults are listed in the Table 20: CBS Forwarding Class Defaults.

Table 20: CBS Forwarding Class Defaults

Forwarding Class	Forwarding Class Label	Default CBS
Network-Control	nc	3
High-1	h1	3
Expedited	ef	1
High-2	h2	1
Low-1	l1	3
Assured	af	1
Low-2	l2	3
Best-Effort	be	1

Parameters

percent

The percent of buffers reserved from the total buffer pool space, expressed as a decimal integer. If 10 Mbytes is the total buffer space in the buffer pool, a value of 10 would reserve 1 Mbyte (10%) of buffer space for the forwarding class queue. The value 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can be applied for scheduling purposes).

Values 0 to 100

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs {*size-in-kbytes* | **default**}

no cbs

Context

[Tree] (config>qos>qgrps>ing>qgrp>policer cbs)

[Tree] (config>qos>qgrps>egr>qgrp>policer cbs)

Full Context

configure qos queue-group-templates ingress queue-group policer cbs

configure qos queue-group-templates egress queue-group policer cbs

Description

The **cbs** command is used to define the default committed buffer size for the template queue or the CBS for the template policer. Overall, the cbs command follows the same behavior and provisioning characteristics as the cbs command in the SAP ingress and egress QoS policy.

The **no** form of this command restores the default CBS size to the template policer.

Default

default

Parameters

size-in-kbytes

For the queues, the size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10 kbytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes). For policers, the size parameter is an integer expression of the number of kilobytes for the policer CBS.

Values 0 to 2683435456, **default**

Minimum default value: 16 Mbytes when CIR equals max or is greater than or equal to the FP capacity (this overrides an explicit configured CBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max CIR capped to 3968 kbytes, with a minimum of 256 bytes.

Platforms

7705 SAR Gen 2

cbs

Syntax

cbs {*size-in-kbytes* | **default**}

no cbs

Context

[Tree] (config>qos>qgrps>ing>qgrp>queue cbs)

[Tree] (config>qos>qgrps>egr>qgrp>queue cbs)

Full Context

configure qos queue-group-templates ingress queue-group queue cbs

configure qos queue-group-templates egress queue-group queue cbs

Description

The **cbs** command is used to define the default committed buffer size for the template queue or the CBS for the template policer. Overall, the cbs command follows the same behavior and provisioning characteristics as the cbs command in the SAP ingress and egress QoS policy.

The **no** form of this command restores the default CBS size to the template policer.

Default

default

Parameters

size-in-kbytes

For the queues, the size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10 kbytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes). For policers, the size parameter is an integer expression of the number of kilobytes for the policer CBS.

Values 0 to 1048576 or **default**

Minimum configurable non-zero value: 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4

Minimum non-zero default value: maximum of 10 ms of CIR or 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4

Platforms

7705 SAR Gen 2

7.17 cd

cd

Syntax

cd [*file-url*]

Context

[Tree] (file cd)

Full Context

file cd

Description

This command displays or changes the current working directory in the local file system.

Parameters

file-url

Specifies the file URL.

Values		
	local-url	[*cflash-id*/][*file-path*] up to 200 characters, including cflash-id directory length 99 chars max each
	remote-url	[{ftp:// | tftp://}*login*:*pswd@remote-locn*/][*file-path*]
		up to 247 characters
		directory length up to 199 characters
	remote-locn	[*hostname* | *ipv4-address* | [*ipv6-address*]]
	ipv4-address	*a.b.c.d*
	ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
		x:x:x:x:x:x:d.d.d.d[-*interface*]
		x - [0 to FFFF]H
		d - [0 to 255]D
		interface - up to 32 characters, for link local addresses 255
	cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

If no *file-url* is entered, the current working directory is displayed.

..

signifies the parent directory. This can be used in place of an actual directory name in a *directory-url*.

directory-url

Specifies the destination directory.

Platforms

7705 SAR Gen 2

7.18 cert

cert

Syntax

cert *cert-filename*

no cert

Context

[Tree] (config>ipsec>cert-profile>entry cert)

Full Context

configure ipsec cert-profile entry cert

Description

This command specifies the file name of an imported certificate for the cert-profile entry.

The **no** form of this command removes the cert-file-name from the entry configuration.

Default

no cert

Platforms

7705 SAR Gen 2

cert

Syntax

cert

Context

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>dyn cert)

[Tree] (config>service>ies>if>sap>ipsec-gw cert)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>dyn cert)

[Tree] (config>service>vprn>if>sap>ipsec-gw cert)

[Tree] (config>ipsec>trans-mode-prof>dyn cert)

[Tree] (config>router>if>ipsec>ipsec-tunnel>dyn cert)

Full Context

configure service vprn interface ipsec ipsec-tunnel dynamic-keying cert

configure service ies interface sap ipsec-gw cert

configure service ies interface ipsec ipsec-tunnel dynamic-keying cert

configure service vprn interface sap ipsec-gw cert

configure ipsec ipsec-transport-mode-profile dynamic-keying cert

configure router interface ipsec ipsec-tunnel dynamic-keying cert

Description

Commands in this context configure certificate parameters.

Platforms

7705 SAR Gen 2

cert

Syntax

cert *cert-filename*

no cert

Context

[Tree] (config>system>security>tls>cert-profile>entry cert)

Full Context

configure system security tls cert-profile entry cert

Description

This command specifies the file name of an imported certificate for the **cert-profile** entry.

The **no** form of the command removes the certificate.

Default

no cert

Parameters

cert-filename

Specifies the file name of the TLS certificate, up to 95 characters in length.

Platforms

7705 SAR Gen 2

cert

Syntax

cert *cert-file-name* [**create**]

no cert

Context

[Tree] (config>system>security>pki>cert-auto-upd cert)

Full Context

configure system security pki certificate-auto-update cert

Description

This command configures the imported certificate filename for the certificate automatic update.

The **no** form of this command removes the *cert-file-name* from the configuration.

Parameters

cert-file-name

Specifies the filename of the certificate, up to 95 characters in length.

Platforms

7705 SAR Gen 2

7.19 cert-file

cert-file

Syntax

cert-file *filename*

no cert-file

Context

[Tree] (config>system>security>pki>ca-profile cert-file)

Full Context

configure system security pki ca-profile cert-file

Description

This command specifies the filename of a file in cf3:\system-pki\cert as the CA's certificate of the ca-profile.

Notes:

- The system will perform following checks against configured cert-file when a **no shutdown** command is issued:
  - Configured cert-file must be a DER formatted X.509v3 certificate file.
  - All non-optional fields defined in section 4.1 of RFC 5280 must exist and conform to the RFC 5280 defined format.
  - Check the version field to see if its value is 0x2.
  - Check The Validity field to see that if the certificate is still in validity period.
  - X509 basic constraints extension must exists, and CA Boolean must be True.
  - If Key Usage extension exists, then at least keyCertSign and cRLSign should be asserted.
  - If the certificate is not a self-signing certificate, then system will try to look for issuer's CA's certificate to verify if this certificate is signed by issuer's CA; but if there is no such CA-profile configured, then system will just proceed with a warning message.
  - If the certificate is not a self-signing certificate, then system will try to look for issuer's CA's CRL to verify that it has not been revoked; but if there is no such CA-profile configured or there is no such CRL, then system will just proceed with a warning message.

  If any of above checks fails, then the **no shutdown** command will fail.
- Changing or removing of **cert-file** is only allowed when the **ca-profile** is in a **shutdown** state.

The **no** form of this command removes the filename from the configuration.

Parameters

filename

Specifies a local CF card file URL.

Platforms

7705 SAR Gen 2

7.20 cert-profile

cert-profile

Syntax

cert-profile *profile-name* [**create**]

no cert-profile *profile-name*

Context

[Tree] (config>ipsec cert-profile)

Full Context

configure ipsec cert-profile

Description

This command creates a new cert-profile or enters the configuration context of an existing cert-profile.

The **no** form of this command removes the profile name from the cert-profile configuration.

Parameters

profile-name

Specifies the name of the certification profile up to 32 characters.

Platforms

7705 SAR Gen 2

cert-profile

Syntax

cert-profile *name*

no cert-profile

Context

[Tree] (config>router>if>ipsec>ipsec-tun>dyn>cert cert-profile)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>dyn>cert cert-profile)

[Tree] (config>service>ies>if>sap>ipsec-gw>cert cert-profile)

[Tree] (config>service>vprn>if>sap>ipsec-gw>cert cert-profile)

[Tree] (config>service>vprn>if>sap>ipsec-tun>dyn>cert cert-profile)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>dyn>cert cert-profile)

[Tree] (config>ipsec>trans-mode-prof>dyn>cert cert-profile)

Full Context

configure router interface ipsec ipsec-tunnel dynamic-keying cert cert-profile

configure service ies interface ipsec ipsec-tunnel dynamic-keying cert cert-profile

configure service ies interface sap ipsec-gw cert cert-profile

configure service vprn interface sap ipsec-gw cert cert-profile

configure service vprn interface sap ipsec-tunnel dynamic-keying cert cert-profile

configure service vprn interface ipsec ipsec-tunnel dynamic-keying cert cert-profile

configure ipsec ipsec-transport-mode-profile dynamic-keying cert cert-profile

Description

This command specifies the name of certificate profile to be used for authentication.

The **no** form of this command removes the name from the configuration.

Parameters

name

Specifies the profile name, up to 32 characters

Platforms

7705 SAR Gen 2

cert-profile

Syntax

cert-profile *profile-name* [**create**]

no cert-profile *profile-name*

Context

[Tree] (config>system>security>tls cert-profile)

Full Context

configure system security tls cert-profile

Description

This command configures TLS certificate profile information. The certificate profile contains the certificates that are sent to the TLS peer (server or client) to authenticate itself. It is mandatory for the TLS server to send this information. The TLS client may optionally send this information upon request from the TLS server.

The **no** form of the command deletes the specified TLS certificate profile.

Parameters

profile-name

Specifies the name of the TLS certificate profile, up to 32 characters in length.

create

Keyword used to create the TLS certificate profile.

Platforms

7705 SAR Gen 2

cert-profile

Syntax

cert-profile *name*

no cert-profile

Context

[Tree] (config>system>security>tls>client-tls-profile cert-profile)

Full Context

configure system security tls client-tls-profile cert-profile

Description

This command assigns a TLS certificate profile to be used by the TLS client profile. This certificate is sent to the server for authentication of the client and public key.

The **no** form of the command removes the TLS certificate profile assignment.

Parameters

name

Specifies the name of the TLS certificate profile, up to 32 characters in length.

Platforms

7705 SAR Gen 2

cert-profile

Syntax

cert-profile *name*

no cert-profile

Context

[Tree] (config>system>security>tls>server-tls-profile cert-profile)

Full Context

configure system security tls server-tls-profile cert-profile

Description

This command assigns a TLS certificate profile to be used by the TLS server profile. This certificate is sent to the client for authentication of the server and public key.

The **no** form of the command removes the TLS certificate profile assignment.

Parameters

name

Specifies the name of the TLS certificate profile, up to 32 characters in length.

Platforms

7705 SAR Gen 2

7.21 cert-request

cert-request

Syntax

cert-request ca *ca-profile-name* **current-key** *key-filename* **current-cert** *cert-filename* [**hash-alg** *hash-algorithm*] **newkey** *key-filename* **subject-dn** *subject-dn* [**domain-name** *domain-names*] [**ip-addr** *ip-address* | *ipv6-address*] **save-as** *save-path-of-result-cert*

Context

[Tree] (admin>certificate>cmpv2 cert-request)

Full Context

admin certificate cmpv2 cert-request

Description

This command requests an additional certificate after the system has obtained the initial certificate from the CA.

The request is authenticated by a signature signed by the current-key, along with the current-cert. The hash algorithm used for signature is depends on the key type:

- DSA key: SHA1
- RSA key: MD5/SHA1/SHA224 | SHA256 | SHA384 | SHA512, by default is SHA1

In some cases, the CA may not return a certificate immediately, due to reasons such as **request processing need manual intervention**. In such cases, the **admin certificate cmpv2 poll** command can be used to poll the status of the request.

Parameters

ca *ca-profile-name*

Specifies a ca-profile name which includes CMP server information up to 32 characters.

current-key *key-filename*

Specifies corresponding certificate issued by the CA up to 95 characters.

current-cert *cert-filename*

Specifies the file name of an imported certificate that is attached to the certificate request up to 95 characters.

newkey *key-filename*

Specifies the file name of the imported key up to 95 characters.

hash-alg *hash-algorithm*

Specifies the hash algorithm for RSA key.

Values md5,sha1,sha224,sha256,sha384,sha512

subject-dn *dn*

Specifies the subject of the requesting certificate up to 256 characters.

Values attr1=val1,attr2=val2 where: attrN={C | ST | O | OU | CN}

save-as *save-path-of-result-cert*

Specifies the save full path name of saving the result certificate, up to 200 characters.

domain-name *domain-names*

Specifies FQDNs for SubjectAltName of the requesting certificate, separated by commas, up to 512 characters.

ip-addr *ip-address* | *ipv6-address*

Specifies an IPv4 or IPv6 address for SubjectAltName of the requesting certificate.

Platforms

7705 SAR Gen 2

7.22 cert-sync

cert-sync

Syntax

[**no**] **cert-sync**

Context

[Tree] (config>redundancy cert-sync)

Full Context

configure redundancy cert-sync

Description

This command automatically synchronizes the certificate/CRL/key when importing or generating (for the key). If a new CF card is inserted into slot3 into the backup CPM, the system will sync the whole system-pki directory from the active CPM.

Default

enabled

Platforms

7705 SAR Gen 2

7.23 certificate

certificate

Syntax

certificate

Context

[Tree] (admin certificate)

Full Context

admin certificate

Description

Commands in this context configure X.509 certificate related operational parameters. For information about CMPv6 admin certificate commands, see the *7705 SAR Gen 2 Multiservice ISA and ESA Guide*

Platforms

7705 SAR Gen 2

certificate

Syntax

certificate

Context

[Tree] (debug certificate)

Full Context

debug certificate

Description

Commands in this context debug certificates.

Platforms

7705 SAR Gen 2

certificate

Syntax

certificate *filename*

Context

[Tree] (debug>ipsec certificate)

Full Context

debug ipsec certificate

Description

This command enables debug for certificate chain computation in cert-profile.

Parameters

filename

Displays the filename of imported certificate, up to 95 characters.

Platforms

7705 SAR Gen 2

7.24 certificate-auto-update

certificate-auto-update

Syntax

certificate-auto-update

Context

[Tree] (config>system>security>pki certificate-auto-update)

Full Context

configure system security pki certificate-auto-update

Description

This command configures automatic updates for the specified certificate. This must be an imported certificate.

Platforms

7705 SAR Gen 2

7.25 certificate-display-format

certificate-display-format

Syntax

certificate-display-format {**ascii** | **utf8**}

Context

[Tree] (config>system>security>pki certificate-display-format)

Full Context

configure system security pki certificate-display-format

Description

This command specifies the display format used for the Certificates and Certificate Revocation Lists.

Default

certificate-display-format ascii

Parameters

ascii

Specifies the ASCII format to use for the Certificates and Certificate Revocation Lists.

utf8

Specifies the UTF8 format to use for the Certificates and Certificate Revocation Lists.

Platforms

7705 SAR Gen 2

7.26 certificate-expiration-warning

certificate-expiration-warning

Syntax

certificate-expiration-warning *hours* [**repeat** *repeat-hours*]

no certificate-expiration-warning

Context

[Tree] (config>system>security>pki certificate-expiration-warning)

Full Context

configure system security pki certificate-expiration-warning

Description

With this command configured, the system issues two types of warnings related to certificate expiration:

- **BeforeExp —** A warning message issued before certificate expire
- **AfterExp —** A warning message issued when certificate expire

This command specifies when system will issue **BeforeExp** message before a certificate expires. For example, with c**ertificate-expiration-warning 5**, the system will issue a **BeforeExp** message 5 hours before a certificate expires. An optional **repeat** <*repeat-hour*> parameter will enable the system to repeat the **BeforeExp** message every hour until the certificate expires.

If the user only wants **AfterExp**, then **certificate-expiration-warning 0** can be used to achieve this.

BeforeExp and **AfterExp** warnings can be cleared in following cases:

- The certificate is reloaded by the **admin certificate reload** command. In this case, if the reloaded file is not expired, then **AfterExp** is cleared. And, if the reloaded file is outside of configured warning window, then the **BeforeExp** is also cleared.
- When the **ca-profile/ipsec-gw/ipsec-tunnel/cert-profile** is shutdown, then **BeforeExp** and **AfterExp** of corresponding certificates are cleared.
- When **no certificate-expiration-warning** command is configured, then all existing **BeforeExp** and **AfterExp** are cleared.
- Users may change the configuration of the **certificate-expiration-warning** so that certain certificates are no longer in the warning window. **BeforeExp** of corresponding certificates are cleared.
- If the system time changes so that the new time causes the certificates to no longer be in the warning window, then **BeforeExp** is cleared. If the new time causes an expired certificate to come non-expired, then **AfterExp** is cleared.

Default

no certificate-expiration-warning

Parameters

hours

Specifies the amount of time before a certificate expires when system issues BeforeExp.

Values 0 to 8760

repeat-hours

Specifies the time the system will repeat BeforeExp every repeat-hour.

Values 0 to 8760

Platforms

7705 SAR Gen 2

7.27 certificate-update-profile

certificate-update-profile

Syntax

certificate-update-profile *profile-name* [**create**]

no certificate-profile *profile-name*

Context

[Tree] (config>system>security>pki certificate-update-profile)

Full Context

configure system security pki certificate-update-profile

Description

Commands in this context configure a certificate update profile that specifies the behavior of the automatic update certificate.

The **no** form of this command removes the profile.

Parameters

profile-name

Specifies the name of the profile, up to 32 characters.

create

Mandatory keyword to create a certificate update profile.

Platforms

7705 SAR Gen 2

7.28 cflash-cap-alarm

cflash-cap-alarm

Syntax

cflash-cap-alarm *cflash-id* **rising-threshold** *threshold* [**falling-threshold** *threshold*] **interval** *seconds* [*rmon-event-type*] [**startup-alarm** *alarm-type*]

no cflash-cap-alarm *cflash-id*

Context

[Tree] (config>system>thresholds cflash-cap-alarm)

Full Context

configure system thresholds cflash-cap-alarm

Description

This command enables capacity monitoring of the compact flash specified in this command. The severity level is alarm. Both a rising and falling threshold can be specified.

The **no** form of this command removes the configured compact flash threshold alarm.

Parameters

cflash-id

Specifies the name of the cflash device to be monitored.

Values cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

rising-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is greater than or equal to this threshold and the associated **startup-alarm** is equal to **rising** or **either**.

After a rising threshold crossing event is generated, another such event will not be generated until the sampled value falls below this threshold and reaches less than or equal to the **falling-threshold** value.

The threshold value represents units of 512 bytes.

Values -2147483648 to 2147483647

Default 0

falling-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is less than or equal to this threshold and the associated **startup-alarm** is equal to **falling** or **either**.

After a falling threshold crossing event is generated, another such event will not be generated until the sampled value rises above this threshold and reaches greater than or equal to the **rising-threshold** value.

The threshold value represents units of 512 bytes.

Values -2147483648 to 2147483647

Default 0

seconds

Specifies the polling period, in seconds, over which the data is sampled and compared with the rising and falling thresholds.

Values 1 to 2147483647

rmon-event-type

Specifies the type of notification action to be taken when this event occurs.

Values log — An entry is made in the RMON-MIB log table for each event occurrence. This does not create an SR OS logger entry. The RMON-MIB log table entries can be viewed using the **show>system>thresholds** CLI command.

trap — An SR OS logger event is generated. The SR OS logger utility then distributes the notification of this event to its configured log destinations, which may be CONSOLE, telnet session, memory log, cflash file, syslog, or SNMP trap destinations logs.

both — Both an entry in the RMON-MIB logTable and an SR OS logger event are generated.

none — No action is taken.

Default both

alarm-type

Specifies the alarm that may be sent when this alarm is first created

If the first sample is greater than or equal to the rising threshold value and **startup-alarm** is equal to **rising** or **either**, a single rising threshold crossing event is generated.

If the first sample is less than or equal to the falling threshold value and **startup-alarm** is equal to **falling** or **either**, a single falling threshold crossing event is generated.

Values rising, falling, either

Default either

Configuration example

```
cflash-cap-alarm cf1-A: rising-threshold 50000000 falling-
threshold 49999900
interval 120 rmon-event-type both start-alarm rising
```

Platforms

7705 SAR Gen 2

7.29 cflash-cap-alarm-pct

cflash-cap-alarm-pct

Syntax

cflash-cap-alarm-pct *cflash-id* **rising-threshold** *percentage* [**falling-threshold** *percentage*] **interval** *seconds* [**rmon-event-type** *event-type*] [**startup-alarm** *alarm-type*]

no cflash-cap-alarm-pct *cflash-id*

Context

[Tree] (config>system>thresholds cflash-cap-alarm-pct)

Full Context

configure system thresholds cflash-cap-alarm-pct

Description

This command enables capacity monitoring of the compact flash specified in this command. The usage is monitored as a percentage of the capacity of the compact flash. The severity level is alarm. Both a rising and falling threshold can be specified.

The **no** form of this command removes the configured compact flash threshold alarm.

Parameters

cflash-id

Specifies the name of the cflash device to be monitored.

Values cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

rising-threshold ***percentage***

Specifies a threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is greater than or equal to this threshold and the associated **startup-alarm** is equal to **rising** or **either**.

After a rising threshold crossing event is generated, another such event will not be generated until the sampled value falls below this threshold and reaches less than or equal to the **falling-threshold** value.

The threshold value is the percentage of used space versus capacity for the specified compact flash.

Values 0 to 100

Default 0

falling-threshold ***percentage***

Specifies a threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is less than or equal to this threshold and the associated **startup-alarm** is equal to **falling** or **either**.

After a falling threshold crossing event is generated, another such event will not be generated until the sampled value rises above this threshold and reaches greater than or equal to the **rising-threshold** value.

The threshold value is the percentage of used space versus capacity for the specified compact flash.

Values 0 to 100

Default 0

seconds

Specifies the polling period, in seconds, over which the data is sampled and compared with the rising and falling thresholds.

Values 1 to 2147483647

event-type

Specifies the type of notification action to be taken when this event occurs.

Values log — An entry is made in the RMON-MIB log table for each event occurrence. This does not create an SR OS logger entry. The RMON-MIB log table entries can be viewed using the **show>system>thresholds** CLI command.

trap — An SR OS logger event is generated. The SR OS logger utility then distributes the notification of this event to its configured log destinations, which may be CONSOLE, telnet session, memory log, cflash file, syslog, or SNMP trap destinations logs.

both — Both an entry in the RMON-MIB logTable and an SR OS logger event are generated.

none — No action is taken.

Default both

alarm-type

Specifies the alarm that may be sent when this alarm is first created.

If the first sample is greater than or equal to the rising threshold value and **startup-alarm** is equal to **rising** or **either**, a single rising threshold crossing event is generated.

If the first sample is less than or equal to the falling threshold value and **startup-alarm** is equal to **falling** or **either**, a single falling threshold crossing event is generated.

Values rising, falling, either

Default either

Configuration example

```
cflash-cap-alarm-pct cf1-A: rising-threshold 70 falling-
threshold 60 interval 120 rmon-event-type both start-alarm
 rising
```

Platforms

7705 SAR Gen 2

7.30 cflash-cap-warn

cflash-cap-warn

Syntax

cflash-cap-warn *cflash-id* **rising-threshold** *threshold* [**falling-threshold** *threshold*] **interval** *seconds* [*rmon-event-type*] [**startup-alarm** *alarm-type*]

no cflash-cap-warn *cflash-id*

Context

[Tree] (config>system>thresholds cflash-cap-warn)

Full Context

configure system thresholds cflash-cap-warn

Description

This command enables capacity monitoring of the compact flash specified in this command.

The severity level is warning. Both a rising and falling threshold can be specified. The **no** form of this command removes the configured compact flash threshold warning.

Parameters

cflash-id

Specifies that the cflash-id specifies the name of the cflash device to be monitored.

Values cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

rising-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is greater than or equal to this threshold and the associated **startup-alarm** is equal to **rising** or **either**.

After a rising threshold crossing event is generated, another such event will not be generated until the sampled value falls below this threshold and reaches less than or equal to the **falling-threshold** value.

The threshold value represents units of 512 bytes.

Values -2147483648 to 2147483647

Default 0

falling-threshold ***threshold***

Specifies a threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is less than or equal to this threshold and the associated **startup-alarm** is equal to **falling** or **either**.

After a falling threshold crossing event is generated, another such event will not be generated until the sampled value rises above this threshold and reaches greater than or equal to the **rising-threshold** value.

The threshold value represents units of 512 bytes.

Values -2147483648 to 2147483647

Default 0

seconds

Specifies the polling period over which the data is sampled and compared with the rising and falling thresholds.

Values 1 to 2147483647

rmon-event-type

Specifies the type of notification action to be taken when this event occurs.

Values log — An entry is made in the RMON-MIB log table for each event occurrence. This does not create an SR OS logger entry. The RMON-MIB log table entries can be viewed using the **show>system>thresholds** CLI command.

trap — An SR OS logger event is generated. The SR OS logger utility then distributes the notification of this event to its configured log destinations, which may be CONSOLE, telnet session, memory log, cflash file, syslog, or SNMP trap destinations logs.

both — Both an entry in the RMON-MIB logTable and a SR OS logger event are generated.

none — No action is taken.

Default both

alarm-type

Specifies the alarm that may be sent when this alarm is first created. If the first sample is greater than or equal to the rising threshold value and **startup-alarm** is equal to **rising** or **either**, a single rising threshold crossing event is generated.

If the first sample is less than or equal to the falling threshold value and **startup-alarm** is equal to **falling** or **either**, a single falling threshold crossing event is generated.

Values rising, falling, either

Default either

Configuration example

```
cflash-cap-warn cf1-B: rising-threshold 2000000 falling-threshold 1999900
interval 240 rmon-event-type trap start-alarm either
```

Platforms

7705 SAR Gen 2

7.31 cflash-cap-warn-pct

cflash-cap-warn-pct

Syntax

cflash-cap-warn-pct *cflash-id* **rising-threshold** *percentage* [**falling-threshold** *percentage*] **interval** *seconds* [**rmon-event-type** *event-type*] [**startup-alarm** *alarm-type*]

no cflash-cap-warn-pct *cflash-id*

Context

[Tree] (config>system>thresholds cflash-cap-warn-pct)

Full Context

configure system thresholds cflash-cap-warn-pct

Description

This command enables capacity monitoring of the compact flash specified in this command. The usage is monitored as a percentage of the capacity of the compact flash.

The severity level is warning. Both a rising and falling threshold can be specified. The **no** form of this command removes the configured compact flash threshold warning.

Parameters

cflash-id

Specifies that the cflash-id specifies the name of the cflash device to be monitored.

Values cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

rising-threshold ***percentage***

Specifies a threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is greater than or equal to this threshold and the associated **startup-alarm** is equal to **rising** or **either**.

After a rising threshold crossing event is generated, another such event will not be generated until the sampled value falls below this threshold and reaches less than or equal to the **falling-threshold** value.

The threshold value is the percentage of used space versus capacity for the specified compact flash.

Values 0 to 100

Default 0

falling-threshold ***percentage***

Specifies a threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is less than or equal to this threshold and the associated **startup-alarm** is equal to **falling** or **either**.

After a falling threshold crossing event is generated, another such event will not be generated until the sampled value rises above this threshold and reaches greater than or equal to the **rising-threshold** value.

The threshold value is the percentage of used space versus capacity for the specified compact flash.

Values 0 to 100

Default 0

seconds

Specifies the polling period over which the data is sampled and compared with the rising and falling thresholds.

Values 1 to 2147483647

event-type

Specifies the type of notification action to be taken when this event occurs.

Values log — An entry is made in the RMON-MIB log table for each event occurrence. This does not create an SR OS logger entry. The RMON-MIB log table entries can be viewed using the **show>system>thresholds** CLI command.

trap — An SR OS logger event is generated. The SR OS logger utility then distributes the notification of this event to its configured log

destinations, which may be CONSOLE, telnet session, memory log, cflash file, syslog, or SNMP trap destinations logs.

both —Both an entry in the RMON-MIB logTable and an SR OS logger event are generated.

none — No action is taken.

Default both

alarm-type

Specifies the alarm that may be sent when this alarm is first created. If the first sample is greater than or equal to the rising threshold value and **startup-alarm** is equal to **rising** or **either**, a single rising threshold crossing event is generated.

If the first sample is less than or equal to the falling threshold value and **startup-alarm** is equal to **falling** or **either**, a single falling threshold crossing event is generated.

Values rising, falling, either

Default either

Configuration example

```
cflash-cap-warn-pct cf1-B: rising-threshold 70 falling-threshold 60
interval 240 rmon-event-type trap start-alarm either
```

Platforms

7705 SAR Gen 2

7.32 chain-to-system-filter

chain-to-system-filter

Syntax

[**no**] **chain-to-system-filter**

Context

[Tree] (config>filter>ipv6-filter chain-to-system-filter)

[Tree] (config>filter>ip-filter chain-to-system-filter)

Full Context

configure filter ipv6-filter chain-to-system-filter

configure filter ip-filter chain-to-system-filter

Description

This command chains this filter to a currently active system filter. When the filter is chained to the system filter, the system filter rules are executed first, and the filter rules are only evaluated if no match on the system filter was found.

The **no** form of the command detaches this filter from the system filter.

Operational note:

If no system filter is currently active, the command has no effect.

Default

no chain-to-system-filter

Platforms

7705 SAR Gen 2

7.33 check-id-kp-cmcra-only

check-id-kp-cmcra-only

Syntax

[**no**] **check-id-kp-cmcra-only**

Context

[Tree] (config>system>security>pki>est-profile check-id-kp-cmcra-only)

Full Context

configure system security pki est-profile check-id-kp-cmcra-only

Description

This command enables checking id-kp-cmcRA in the EST certificate. When enabled, instead of the subject or subject alternative name, only the id-kp-cmcRA existence in extended key usage extension of EST server certificate is checked. The id-kp-cmcRA identifies a Registration Authority.

The **no** form of this command reverts to the default value.

Default

no check-id-kp-cmcra-only

Platforms

7705 SAR Gen 2

7.34 check-zero

check-zero

Syntax

check-zero {**enable** | **disable**}

no check-zero

Context

[Tree] (config>service>vprn>rip check-zero)

[Tree] (config>service>vprn>ripng check-zero)

[Tree] (config>service>vprn>rip>group>neighbor check-zero)

[Tree] (config>service>vprn>ripng>group>neighbor check-zero)

[Tree] (config>service>vprn>rip>group check-zero)

[Tree] (config>service>vprn>ripng>group check-zero)

Full Context

configure service vprn rip check-zero

configure service vprn ripng check-zero

configure service vprn rip group neighbor check-zero

configure service vprn ripng group neighbor check-zero

configure service vprn rip group check-zero

configure service vprn ripng group check-zero

Description

This command enables checking for zero values in fields specified to be zero by the RIPv1 and RIPv2 specifications.

The **no** form of this command disables this check and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.

Default

no check-zero

Parameters

enable

Enables checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting noncompliant RIP messages.

disable

Disables the checking and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.

Platforms

7705 SAR Gen 2

check-zero

Syntax

check-zero {**enable** | **disable**}

no check-zero

Context

[Tree] (config>router>rip check-zero)

[Tree] (config>router>rip>group>neighbor check-zero)

[Tree] (config>router>ripng>group>neighbor check-zero)

[Tree] (config>router>ripng>group check-zero)

[Tree] (config>router>ripng check-zero)

[Tree] (config>router>rip>group check-zero)

Full Context

configure router rip check-zero

configure router rip group neighbor check-zero

configure router ripng group neighbor check-zero

configure router ripng group check-zero

configure router ripng check-zero

configure router rip group check-zero

Description

This command enables checking for zero values in fields specified to be zero by the RIPv1 and RIPv2 specifications.

The **check-zero enable** command enables checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting non-compliant RIP messages.

The **check-zero disable** command disables this check and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.

This configuration parameter can be set at three levels: global level (applies to all groups and neighbor interfaces), group level (applies to all neighbor interfaces in the group) or neighbor level (only applies to the specified neighbor interface). The most specific value is used. In particular if no value is set (**no check-zero**), the setting from the less specific level is inherited by the lower level.

The **no** form of the command removes the **check-zero** command from the configuration.

Parameters

enable

Specifies to reject RIP messages which do not have zero in the RIPv1 and RIPv2 mandatory fields.

disable

Specifies allows receipt of RIP messages which do not have the mandatory zero fields reset.

Platforms

7705 SAR Gen 2

7.35 checksum

checksum

Syntax

checksum {**md5** | **sha256**} *file-url*

Context

[Tree] (file checksum)

Full Context

file checksum

Description

This command computes and displays a checksum for a file.

Parameters

md5

Specifies the use of the MD5 algorithm to produce the file checksum.

sha256

Specifies the use of the SHA-256 algorithm to produce the file checksum.

file-url

Specifies the location of the file.

Values

local-url	[*cflash-id*/][*file-path*] up to 200 characters, including cflash-id directory length 99 chars max each
remote-url	[{ftp:// | tftp:// | http:// | https://}*login*:*pswd*@*remote-locn*/][*file-path*] up to 247 characters

	directory length up to 199 characters
remote-locn	[*hostname* | *ipv4-address* | [*ipv6-address*]]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[*-interface*] *x:x:x:x:x:x:d.d.d.d*[*-interface*] *x* - [0 to FFFF]H *d* - [0 to 255]D interface - up to 32 characters, for link local addresses 255
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

7.36 chli-event

chli-event

Syntax

chli-event {**forward** | **backward** | **aggregate**} **threshold** *raise-threshold* [**clear** *clear-threshold*]

no chli-event {**forward** | **backward** | **aggregate**}

Context

[Tree] (config>oam-pm>session>ip>twamp-light>loss-events chli-event)

Full Context

configure oam-pm session ip twamp-light loss-events chli-event

Description

This command sets the consecutive high loss interval (CHLI) threshold to be monitored and the associated thresholds using the counter of the specified direction. The aggregate is a function of summing forward and backward. This value is only used as a threshold mechanism and is not part of the stored statistics. If the optional **clear** *clear-threshold* parameter is not specified, the traffic crossing alarm is stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and regardless of any previous window. Each unique event can only be raised once within measurement interval. If the optional **clear** *clear-threshold* parameter is specified, the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried

forward to following measurement intervals. If a threshold crossing event is raised another is raised until a measurement interval completes and the clear threshold has not been exceeded. A clear event is raised under that condition.

The **no** form of this command removes the event threshold for frame loss ratio. The direction must be included with the **no** command.

Default

no chli-event forward

no chli-event backward

no chli-event aggregate

Parameters

forward

Specifies the threshold is applied to the forward direction count.

backward

Specifies the threshold is applied to the backward direction count.

aggregate

Specifies the threshold is applied to the aggregate count (sum of forward and backward).

raise-threshold

Specifies the numerical value compared to the CHLI counter that is the rising threshold that determines when the event is to be generated, when the percentage of loss value is reached.

Values 1 to 864000

clear-threshold

Specifies an optional numerical value compared to the CHLI counter used for stateful behavior that allows the operator to configure a value lower than the rising percentage to indicate when the clear event should be generated.

Values 0 to 863999

A value of zero means that the CHLI counter must be 0.

Platforms

7705 SAR Gen 2

7.37 cipher

cipher

Syntax

cipher *index* **name** *cipher-name*

no cipher *index*

Context

[Tree] (config>system>security>ssh>server-cipher-list cipher)

[Tree] (config>system>security>ssh>client-cipher-list cipher)

Full Context

configure system security ssh server-cipher-list cipher

configure system security ssh client-cipher-list cipher

Description

This command configures a cipher. Client-ciphers are used when the SR OS is acting as an SSH client. Server-ciphers are used when the SR OS is acting as an SSH server.

The **no** form of this command removes the index and cipher name from the configuration.

Default

no cipher *index*

Parameters

index

Specifies the index of the cipher in the list.

Values 1 to 255

cipher-name

Specifies the algorithm used when performing encryption or decryption.

Values Client ciphers: 3des-cbc, aes128-cbc, aes192-cbc, aes256-cbc, aes128-ctr, aes192-ctr, aes256-ctr.

Server ciphers: 3des-cbc, aes128-cbc, aes192-cbc, aes256-cbc, aes128-ctr, aes192-ctr, aes256-ctr.

The following table lists the default ciphers used for SSHv2.

Table 21: SSHv2 Default Ciphers

Cipher index value	Cipher name
190	aes256-ctr
192	aes192-ctr
194	aes128-ctr
200	aes128-cbc
205	3des-cbc

Cipher index value	Cipher name
225	aes192-cbc
230	aes256-cbc

Platforms

7705 SAR Gen 2

cipher

Syntax

cipher *index* **name** *cipher-suite-code*

no cipher *index*

Context

[Tree] (config>system>security>tls>server-cipher-list cipher)

[Tree] (config>system>security>tls>client-cipher-list cipher)

Full Context

configure system security tls server-cipher-list cipher

configure system security tls client-cipher-list cipher

Description

This command configures the cipher suite to be negotiated by the server and client.

Parameters

index

Specifies the index number. The index number provides the location of the cipher in the negotiation list, with the lower index numbers being higher in the negotiation list and the higher index numbers being at the bottom of the list.

Values 1 to 255

cipher-suite-code

Specifies the cipher suite code.

Values tls-rsa-with-3des-ede-cbc-sha

tls-rsa-with-aes128-cbc-sha

tls-rsa-with-aes256-cbc-sha

tls-rsa-with-aes128-cbc-sha256

tls-rsa-with-aes256-cbc-sha256

tls-rsa-with-aes128-gcm-sha256

tls-rsa-with-aes256-gcm-sha384
tls-ecdhe-rsa-aes128-gcm-sha256
tls-ecdhe-rsa-aes256-gcm-sha384

Platforms
7705 SAR Gen 2

7.38 cipher-list

cipher-list

Syntax
cipher-list *name*
no cipher-list

Context
[Tree] (config>system>security>tls>client-tls-profile cipher-list)

Full Context
configure system security tls client-tls-profile cipher-list

Description
This command assigns the cipher list to be used by the TLS client profile for negotiation in the client Hello message.

Parameters
name
Specifies the name of the cipher list.

Platforms
7705 SAR Gen 2

cipher-list

Syntax
cipher-list *name*
no cipher-list

Context

[Tree] (config>system>security>tls>server-tls-profile cipher-list)

Full Context

configure system security tls server-tls-profile cipher-list

Description

This command assigns a cipher list to be used by the TLS server profile. This cipher list is used to find matching ciphers with the cipher list that is received from the client.

The **no** form of the command removes the cipher list.

Parameters

name

Specifies the name of the cipher list, up to 32 characters in length.

Platforms

7705 SAR Gen 2

7.39 cipher-suite

cipher-suite

Syntax

cipher-suite *cipher-suite*

no cipher-suite

Context

[Tree] (config>macsec>connectivity-association cipher-suite)

Full Context

configure macsec connectivity-association cipher-suite

Description

This command configures encryption of data path PDUs. When all parties in the Connectivity Association (CA) have the SAK, they use the above algorithm in conjunction with the SAK to encrypt the data path PDUs.

The XPN 64 bit (extended packet number) can be used for higher rate ports such as 10 GigE to minimize the window rollover and renegotiation of the SAK.

The **no** form of this command disables encryption of data path PDUs.

Default

cipher-suite gcm-aes-128

Parameters

cypher-suite

Specifies the algorithm.

Values gcm-aes-128 — algorithm is used for control plain encryption

gcm-aes-256 — algorithm is used for control plain encryption

gcm-aes-xpn-128 — algorithm with extended packet number is used for control plain encryption

gcm-aes-xpn-256 — algorithm with extended packet number is used for control plain encryption

Platforms

7705 SAR Gen 2

7.40 circuit-id

circuit-id

Syntax

circuit-id string *ascii-string*

circuit-id hex *hex-string*

no circuit-id

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>host-ident circuit-id)

Full Context

configure subscriber-mgmt local-user-db ipoe host host-identification circuit-id

Description

This command specifies the circuit ID to match for a host lookup. When the LUDB is accessed using a DHCPv4 server, the circuit ID is matched against DHCP Option 82.

Note:
This command is only used when **circuit-id** is configured as one of the **match-list** parameters.

The **no** form of this command removes the circuit ID from the configuration.

Parameters

ascii-string

Specifies the circuit ID from the Option 82, up to 127 characters.

hex-string

Specifies the circuit ID in hexadecimal format from the Option 82.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

Platforms

7705 SAR Gen 2

circuit-id

Syntax

circuit-id

circuit-id {**ascii-tuple** | **if-index** | **sap-id** | **vlan-ascii-tuple**}

circuit-id hex [*hex-string*]

no circuit-id

Context

[Tree] (config>service>vprn>if>dhcp>option circuit-id)

[Tree] (config>service>vpls>sap>dhcp>option circuit-id)

[Tree] (config>service>ies>if>dhcp>option circuit-id)

Full Context

configure service vprn interface dhcp option circuit-id

configure service vpls sap dhcp option circuit-id

configure service ies interface dhcp option circuit-id

Description

When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by "|". If no keyword is configured, then the circuit-id sub-option will not be part of the information option (Option 82). When the command is configured without any parameters, it equals to circuit-id ascii-tuple.

To send a tuple in the circuit ID, the **action replace** command must be configured in the same context.

If disabled, the **circuit-id** sub-option of the DHCP packet is left empty.

The **no** form of this command specifies to leave the circuit-id option of the packet empty.

Default

circuit-id ascii-tuple

Parameters

ascii-tuple

Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used.

ifindex

Specifies that the interface index is used. The If Index of a router interface can be displayed using the command **show>router>if>detail**.

sap-id

Specifies that the SAP identifier is used.

vlan-ascii-tuple

Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the Option 82 bits are stripped, dot1p bits are copied to the Ethernet header of an outgoing packet.

hex-string

Specifies the hex value of this option.

Values 0x0 to 0xFFFFFFFF...(up to 64 hex nibbles)

Platforms

7705 SAR Gen 2

circuit-id

Syntax

circuit-id {**ascii-tuple** | **ifindex** | **if-name** | **port-id** | **vlan-ascii-tuple** | **none**}

no circuit-id

Context

[Tree] (config>router>if>dhcp>option circuit-id)

Full Context

configure router interface dhcp option circuit-id

Description

When enabled, the router sends the interface index (If Index) in the **circuit-id** suboption of the DHCP packet. The If Index of a router interface can be displayed using the command **show router interface detail**. This option specifies data that must be unique to the router that is relaying the circuit.

If disabled, the **circuit-id** suboption of the DHCP packet will be left empty.

The **no** form of this command returns the system to the default.

Default

circuit-id ascii-tuple

Parameters

ascii-tuple

Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by a pipe (|).

ifindex

Specifies that the interface index will be used. The If Index of a router interface can be displayed using the command **show router interface detail**.

if-name

Specifies the interface name.

port-id

Specifies the port ID.

vlan-ascii-tuple

Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Therefore, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

none

Specifies that no circuit should be used.

Platforms

7705 SAR Gen 2

7.41 ckn

ckn

Syntax

ckn *hex-string*

no ckn

Context

[Tree] (config>macsec>conn-assoc>static-cak>pre-shared-key ckn)

Full Context

configure macsec connectivity-association static-cak pre-shared-key ckn

Description

Specifies the connectivity association key name (CKN) for a pre-shared key.

CKN is appended to the MKA for identification of the appropriate CAK by the peer.

The **no** form of this command reverts to the default value.

Parameters

hex-string

Specifies the value of the CKN.

Values 32 octets char (64 hex)

Platforms

7705 SAR Gen 2

7.42 class-type

class-type

Syntax

class-type *ct-number*

no class-type

Context

[Tree] (config>router>mpls>lsp>primary class-type)

[Tree] (config>router>mpls>lsp-template class-type)

[Tree] (config>router>mpls>lsp class-type)

[Tree] (config>router>mpls>lsp>secondary class-type)

Full Context

configure router mpls lsp primary class-type

configure router mpls lsp-template class-type

configure router mpls lsp class-type

configure router mpls lsp secondary class-type

Description

This command configures the Diff-Serv Class Type (CT) for an LSP, the LSP primary path, or the LSP secondary path. The path level configuration overrides the LSP level configuration. However, only one CT per LSP path will be allowed as per RFC 4124.

The signaled CT of a dynamic bypass is always be CT0 regardless of the CT of the primary LSP path. The setup and hold priorities must be set to default values, that is, 7 and 0 respectively. This assumes that

the operator configured a couple of TE classes, one which combines CT0 and a priority of 7 and the other which combines CT0 and a priority of 0. If not, the bypass LSP will not be signaled and will go into the down state.

The operator cannot configure the CT, setup priority, and hold priority of a manual bypass. They are always signaled with CT0 and the default setup and holding priorities.

The signaled CT and setup priority of a detour LSP must match those of the primary LSP path it is associated with.

If the operator changes the CT of an LSP or of an LSP path, or changes the setup and holding priorities of an LSP path, the path will be torn down and retried.

An LSP which does not have the CT explicitly configured will behave like a CT0 LSP when Diff-Serv is enabled.

If the operator configured a combination of a CT and a setup priority and/or a combination of a CT and a holding priority for an LSP path that are not supported by the user-defined TE classes, the LSP path will be kept in a down state and an error code will be displayed in the show command output for the LSP path.

The **no** form of this command reverts to the default value.

Default

class-type 0

Parameters

ct-number

Specifies the Diff-Serv Class Type number.

Values 0 to 7

Platforms

7705 SAR Gen 2

7.43 class-type-bw

class-type-bw

Syntax

class-type-bw ct0 *%-link-bandwidth* **ct1***%-link-bandwidth* **ct2***%-link-bandwidth* **ct3***%-link-bandwidth* **ct4***%-link-bandwidth* **ct5***%-link-bandwidth* **ct6***%-link-bandwidth* **ct7***%-link-bandwidth*

no class-type-bw

Context

[Tree] (config>router>rsvp>interface class-type-bw)

[Tree] (config>router>rsvp>diffserv-te class-type-bw)

Full Context

configure router rsvp interface class-type-bw

configure router rsvp diffserv-te class-type-bw

Description

This command configures the percentage of RSVP interface bandwidth each CT shares, for example, the Bandwidth Constraint (BC).

The absolute value of the CT share of the interface bandwidth is derived as the percentage of the bandwidth advertised by IGP in the Maximum Reservable Link Bandwidth TE parameter, for example, the link bandwidth multiplied by the RSVP interface **subscription** *percentage* parameter.

Note:
This configuration also exists at RSVP interface level and the interface specific configured value overrides the global configured value. The BC value can be changed at any time.

The RSVP interface **subscription** *percentage* parameter is configured in the **config>router>rsvp>interface** context.

The operator can specify the Bandwidth Constraint (BC) for a CT which is not used in any of the TE class definition but that does not get used by any LSP originating or transiting this node.

When Diff-Serv is disabled on the node, this model degenerates into a single default CT internally with eight preemption priorities and a non-configurable BC equal to the Maximum Reservable Link Bandwidth. This would behave exactly like CT0 with eight preemption priorities and BC= Maximum Reservable Link Bandwidth if Diff-Serv was enabled.

The **no** form of this command reverts to the default value.

Parameters

ct0 (ct1/ct2/ —ct7) %*link-bandwidth*

The Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.

Values 0 to 100 %

Default 0

Platforms

7705 SAR Gen 2

7.44 classic-cli

classic-cli

Syntax

classic-cli

Context

[Tree] (config>system>management-interface>cli classic-cli)

Full Context

configure system management-interface cli classic-cli

Description

Commands in this context configure the classic CLI management interface.

Platforms

7705 SAR Gen 2

classic-cli

Syntax

classic-cli

Context

[Tree] (config>system>security>management-interface classic-cli)

Full Context

configure system security management-interface classic-cli

Description

Commands in this context configure hash-control for the classic CLI interface.

Platforms

7705 SAR Gen 2

7.45 classic-lsn-max-subscriber-limit

classic-lsn-max-subscriber-limit

Syntax

classic-lsn-max-subscriber-limit *max*

no classic-lsn-max-subscriber-limit

Context

[Tree] (config>router>nat>inside classic-lsn-max-subscriber-limit)

[Tree] (config>service>vprn>nat>inside classic-lsn-max-subscriber-limit)

Full Context

configure router nat inside classic-lsn-max-subscriber-limit

configure service vprn nat inside classic-lsn-max-subscriber-limit

Description

This command sets the granularity of traffic distribution in the upstream direction across the MS-ISA within the scope of an inside routing context. Traffic distribution mechanism is based on the source IPv4 addresses/prefixes. More granular distribution is based on the IPv4 address, while distribution based on the IPv4 prefix (determined by prefix length) will be less granular. The granularity will further decrease with shorter prefix length.

For example, a prefix length of 32 will distribute individual /32 IPv4 addresses over multiple MS-ISAs in an ISA group. This will ensure better traffic load balancing at the expense of forwarding table utilization on the outside (public side) where each /32 is installed in the forwarding table. On the contrary, shorter prefixes will ensure better utilization of the forwarding table on the outside, at the expense of coarser spread of IP addresses over multiple MS-ISAs.

This command affects all flavors of LSN44 within the inside routing contexts, although its primary use is intended for deterministic NAT and dnat-only.

The length of the prefix that is used for distribution purposes is (32-n), where 2^n= classic-lsn-max-subscriber-limit. For example, if traffic distribution is based on the IPv4 address (prefix length = 32), then n must be 0. From here, it follows that classic-lsn-max-subscriber-limit must be set to 1:

Prefix length = 32 -> 32-n = 32 -> n=0 -> 2^0= 1 = *classic-lsn-max-subscriber-limit classic-lsn-max-subscriber-limit* = 1

The implicit method given by this command uses power of 2 calculations to provide prefix length for traffic distribution purposes. This roundabout approach to determine the prefix-length has roots in deterministic NAT where this command was originally introduced.

Even though deterministic NAT and dnat-only have very little in common, the method (and CLI syntax) for calculating the prefix length using the classic-lsn-max-subscriber-limit parameter for traffic distribution purposes is shared between the two. In dnat-only, this parameter is important from an operational perspective since it affects traffic load balancing over MS-ISA and the size of the routing table.

This command must be configured before any prefix is configured and can be modified only if there are no prefixes configured under the deterministic NAT.

Parameters

max

The power of 2 (2^n) value which in deterministic NAT must match the largest subscriber-limit value in any deterministic pool referenced from this inside routing instance.

In **dnat-only**, this value can be set to any value from the allowed range.

In both cases, this value will determine the prefix-length (17-32) that will directly influence load distribution between the MS-ISAs and the size of the routing table.

Values 1,2,4,8 to 32768

Platforms

7705 SAR Gen 2

7.46 clear

clear

Syntax

clear

Context

[Tree] (admin clear)

Full Context

admin clear

Description

Commands in this context clear statistics.

Platforms

7705 SAR Gen 2

clear

Syntax

clear [**now**]

Context

[Tree] (admin>system>license clear)

Full Context

admin system license clear

Description

This command removes the entitlements that were installed using a license file.

All the entitlements must be unallocated; otherwise, the command fails.

Parameters

now

Keyword used to specify the immediate removal of the license file entitlements. If the **now** keyword is not present, the user is prompted to confirm the removal.

Platforms

7705 SAR Gen 2

7.47 clear-df-bit

clear-df-bit

Syntax

[**no**] **clear-df-bit**

Context

[Tree] (config>service>ies>if>sap>ip-tunnel clear-df-bit)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel clear-df-bit)

[Tree] (config>router>if>ipsec>ipsec-tunnel clear-df-bit)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel clear-df-bit)

Full Context

configure service ies interface sap ip-tunnel clear-df-bit

configure service vprn interface sap ipsec-tunnel clear-df-bit

configure router interface ipsec ipsec-tunnel clear-df-bit

configure service ies interface ipsec ipsec-tunnel clear-df-bit

Description

This command instructs the MS-ISA to reset the DF bit to 0 in all payload IP packets associated with the GRE or IPsec tunnel, before any potential fragmentation resulting from the **ip-mtu** command (this requires a modification of the header checksum).

The **no** form of this command disables the DF bit reset.

Default

no clear-df-bit

Platforms

7705 SAR Gen 2

clear-df-bit

Syntax

[**no**] **clear-df-bit**

Context

[Tree] (config>ipsec>tnl-temp clear-df-bit)

Full Context

configure ipsec tunnel-template clear-df-bit

Description

This command enables clearing of the Do-not-Fragment bit.

Default

no clear-df-bit

Platforms

7705 SAR Gen 2

7.48 clear-ocsp-cache

clear-ocsp-cache

Syntax

clear-ocsp-cache [*entry-id*]

Context

[Tree] (admin>certificate clear-ocsp-cache)

Full Context

admin certificate clear-ocsp-cache

Description

This command clears the current OCSP response cache. If optional issuer and serial-number are not specified, then all current cached results are cleared.

Parameters

entry-id

Specifies the local cache entry identifier of the certificate to clear.

Values 1 to 2000

Platforms

7705 SAR Gen 2

7.49 clear-request

clear-request

Syntax

clear-request ca *ca-profile-name*

Context

[Tree] (admin>certificate>cmpv2 clear-request)

Full Context

admin certificate cmpv2 clear-request

Description

This command clears current pending CMPv2 requests toward the specified CA. If there are no pending requests, it will clear the saved result of prior request.

Parameters

ca *ca-profile-name*

Specifies a ca-profile name up to 32 characters.

Platforms

7705 SAR Gen 2

7.50 clear-tag-mode

clear-tag-mode

Syntax

clear-tag-mode *clear-tag-mode*

no clear-tag-mode

Context

[Tree] (config>macsec>connectivity-association clear-tag-mode)

Full Context

configure macsec connectivity-association clear-tag-mode

Description

This command puts 802.1Q tags in cleartext before the SecTAG. There are two modes: **single-tag** and **dual-tag**.

Table 22: Encrypted Dot1q and QinQ Packet Format explains the encrypted dot1q and QinQ packet format when clear-tag-mode single-tag or dual-tag is configured.

The **no** form of this command puts all dot1q tags encrypted after the SecTAG.

Table 22: Encrypted Dot1q and QinQ Packet Format

Unencrypted format	Clear-tag-mode	Pre-encryption (Tx)	Pre-decryption (Rx)
Single tag (dot1q)	single-tag	DA, SA, TPID, VID, Etype	DA, SA, TPID, VID, SecTag
Single tag (dot1q)	dual-tag	DA, SA, TPID, VID, Etype	DA, SA, TPID, VID, SecTag
Double tag (q-in-q)	single-tag	DA, SA, TPID1, VID1, IPID2, VID2, Etype	DA, SA, TPID1, VID1, SecTag
Double tag (QinQ)	dual-tag	DA, SA, TPID1, VID1, IPID2, VID2, Etype	DA, SA, TPID1, VID1, IPID2, VID2, SecTag

Default

no clear-tag-mode

Parameters

clear-tag-mode

Specifies the clear tag mode.

Values single-tag, dual-tag

Platforms

7705 SAR Gen 2

7.51 cli

cli

Syntax

cli

Context

[Tree] (config>system>management-interface cli)

Full Context

configure system management-interface cli

Description

Commands in this context configure the CLI management interfaces.

Platforms

7705 SAR Gen 2

cli

Syntax

cli {**warning** | **info**}

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>message-severity-level cli)

Full Context

configure system management-interface cli md-cli environment message-severity-level cli

Description

This command specifies the threshold for CLI messages.

Default

cli info

Parameters

warning

Specifies that WARNING messages are displayed but INFO messages are suppressed.

info

Specifies that INFO messages and WARNING messages are displayed.

Platforms

7705 SAR Gen 2

7.52 cli-engine

cli-engine

Syntax

cli-engine {**classic-cli** | **md-cli**} [{**classic-cli** | **md-cli**}]

no cli-engine

Context

[Tree] (config>system>management-interface>cli cli-engine)

Full Context

configure system management-interface cli cli-engine

Description

This command configures the system-wide CLI engine. The operator can configure one or both engines. For the configuration to take effect, exit the running CLI session and start a new session after committing the new value.

Parameters

classic-cli

Specifies the classic CLI.

md-cli

Specifies the MD-CLI.

Platforms

7705 SAR Gen 2

7.53 cli-script

cli-script

Syntax

cli-script

Context

[Tree] (config>system>security cli-script)

Full Context

configure system security cli-script

Description

Commands in this context configure the security parameters in the system.

Platforms

7705 SAR Gen 2

7.54 cli-session-group

cli-session-group

Syntax

cli-session-group *session-group-name* [**create**]

no cli-session-group *session-group-name*

Context

[Tree] (config>system>security cli-session-group)

Full Context

configure system security cli-session-group

Description

This command is used to configure a session group that can be used to limit the number of CLI sessions available to members of the group.

Parameters

session-group-name

Specifies a particular session group.

Platforms

7705 SAR Gen 2

7.55 cli-user

cli-user

Syntax

cli-user *user-name*

no cli-user

Context

[Tree] (config>system>security>cli-script>authorization>event-handler cli-user)

[Tree] (config>system>security>cli-script>authorization>cron cli-user)

Full Context

configure system security cli-script authorization event-handler cli-user

configure system security cli-script authorization cron cli-user

Description

This command configures the user context under which various types of CLI scripts should execute in order to authorize the script commands. TACACS+ and RADIUS users and authorization are not permitted for **cli-script** authorization.

The **no** form of this command configures scripts to execute with no restrictions and without performing authorization.

Default

no cli-user

Parameters

user-name

The name of a user in the local node database. TACACS+ or RADIUS users cannot be used. The user configuration should reference a valid local profile for authorization.

Platforms

7705 SAR Gen 2

7.56 client

client

Syntax

client *client-index* [**create**]

no client *client-index*

Context

[Tree] (config>ipsec>client-db client)

Full Context

configure ipsec client-db client

Description

This command creates a new IPsec client entry in the client-db or enters the configuration context of an existing client entry.

There may be multiple client entries defined in the same client-db. If there are multiple entries that match the new tunnel request, then the system will select the entry that has smallest client-index.

The **no** form of this command reverts to the default.

Parameters

client-index

Specifies the ID of the client entry.

Values 1 to 8000

create

Keyword used to create the security policy instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

client

Syntax

client all

client *ip-address*

no client

Context

[Tree] (debug>system>grpc client)

Full Context

debug system grpc client

Description

This command enables debug output for all clients for a particular client.

The **no** form of this command deactivates debugging for all clients.

Parameters

all

Specifies that debugging will occur for all clients.

ip-address

Specifies the IPv4 or IPv6 address of the client.

Platforms

7705 SAR Gen 2

client

Syntax

client

Context

[Tree] (config>system>security>ssh>key-re-exchange client)

Full Context

configure system security ssh key-re-exchange client

Description

Commands in this context enable the key re-exchange for SR OS as an SSH client.

Platforms

7705 SAR Gen 2

7.57 client-cert-subject-key-id

client-cert-subject-key-id

Syntax

[**no**] **client-cert-subject-key-id**

Context

[Tree] (config>ipsec>rad-auth-plcy>include client-cert-subject-key-id)

Full Context

configure ipsec radius-authentication-policy include-radius-attribute client-cert-subject-key-id

Description

This command enables the inclusion of the Subject Key Identifier of the peer's certificate in the RADIUS Access-Request packet as VSA: Alc-Subject-Key-Identifier.

Default

no client-cert-subject-key-id

Platforms

7705 SAR Gen 2

7.58 client-cipher-list

client-cipher-list

Syntax

client-cipher-list

Context

[Tree] (config>system>security>ssh client-cipher-list)

Full Context

configure system security ssh client-cipher-list

Description

Commands in this context configure a list of allowed ciphers by the SSH client.

Platforms

7705 SAR Gen 2

client-cipher-list

Syntax

client-cipher-list *name* [**create**]

no client-cipher-list *name*

Context

[Tree] (config>system>security>tls client-cipher-list)

Full Context

configure system security tls client-cipher-list

Description

This command creates a cipher list that the client sends to the server in the client Hello message. It is a list of ciphers that are supported and preferred by the SR OS to be used in the TLS session. The server matches this list against the server cipher list. The most preferred cipher found in both lists is chosen.

Parameters

name

Specifies the name of the client cipher list, up to 32 characters in length.

create

Keyword used to create the client cipher list.

Platforms

7705 SAR Gen 2

7.59 client-db

client-db

Syntax

client-db *db-name* [**create**]

no client-db *db-name*

Context

[Tree] (config>ipsec client-db)

Full Context

configure ipsec client-db

Description

This command creates a new IPsec client-db or enters the configuration context of an existing client-db.

An IPsec client-db can be used for IKEv2 dynamic LAN-to-LAN tunnel authentication and authorization. When a new tunnel request is received, the system will match the request to the client entries configured in client-db and use credentials returned by the matched client entry for authentication. If authentication succeeds, the system could also use the IPsec configuration parameters (such as **private-service-id**) returned by the matched entry to set up the tunnel.

The configured client-db is referenced under the ipsec-gw configuration context using the **client-db** command.

The **no** form of this command removes the *db-name* from the configuration.

Parameters

db-name

Specifies the name of this IPsec client up to 32 characters.

create

Keyword used to create the security policy instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

client-db

Syntax

client-db *name*

client-db *name* **fallback**

client-db *name* **no-fallback**

no client-db

Context

[Tree] (config>service>ies>if>sap>ipsec-gw client-db)

[Tree] (config>service>vprn>if>sap>ipsec-gw client-db)

Full Context

configure service ies interface sap ipsec-gw client-db

configure service vprn interface sap ipsec-gw client-db

Description

This command enables the use of an IPsec client database. The system uses the specified client database to authenticate IKEv2 dynamic LAN-to-LAN tunnel.

Default

no client-db

Parameters

name

Specifies the name of the client database.

fallback

Specifies whether or not this IPsec gateway falls back to the default authentication policy when the IPsec tunnel authentication request fails to match any clients in the IPsec database.

no-fallback

Specifies that if the client database lookup fails to return a matched result, the system will fail the tunnel setup.

Platforms

7705 SAR Gen 2

client-db

Syntax

[**no**] **no client-db** *db-name*

Context

[Tree] (debug>ipsec client-db)

Full Context

debug ipsec client-db

Description

This command enables debugging for the specified IPsec client-db.

Parameters

db-name

Specifies the IPsec client database name, up to 32 characters.

Platforms

7705 SAR Gen 2

7.60 client-group-list

client-group-list

Syntax

client-group-list *name* [**create**]

no client-group-list *name*

Context

[Tree] (config>system>security>tls client-group-list)

Full Context

configure system security tls client-group-list

Description

This command configures a list of group suite codes that the client sends in a client Hello message.

The **no** form of this command removes the client group list.

Parameters

name

Specifies the name of the client group list, up to 32 characters.

create

Keyword used to create the client group list.

Platforms

7705 SAR Gen 2

7.61 client-host-key-list

client-host-key-list

Syntax

client-host-key-list

Context

[Tree] (config>system>security>ssh client-host-key-list)

Full Context

configure system security ssh client-host-key-list

Description

Commands in this context configure the list of host key algorithms negotiated by the SR OS acting as the SSH client.

Platforms

7705 SAR Gen 2

7.62 client-identification

client-identification

Syntax

client-identification

Context

[Tree] (config>ipsec>client-db>client client-identification)

Full Context

configure ipsec client-db client client-identification

Description

Commands in this context configure client ID information of this IPsec client.

If there are multiple match input are configured in the match-list of the client-db, then all corresponding match criteria must be configured for the client-entry.

Platforms

7705 SAR Gen 2

7.63 client-kex-list

client-kex-list

Syntax

client-kex-list

Context

[Tree] (config>system>security>ssh client-kex-list)

Full Context

configure system security ssh client-kex-list

Description

Commands in this context configure SSH KEX algorithms for SR OS as a client.

An empty list is the default list that the SSH KEX advertises. The default list contains the following:

ecdh-sha2-nistp512

ecdh-sha2-nistp384

ecdh-sha2-nistp256

diffie-hellman-group16-sha512

diffie-hellman-group14-sha256

diffie-hellman-group14-sha1

diffie-hellman-group-exchange-sha1

diffie-hellman-group1-sha1

Platforms

7705 SAR Gen 2

7.64 client-mac-address

client-mac-address

Syntax

[**no**] **client-mac-address**

Context

[Tree] (config>service>vpls>sap>dhcp>option>vendor client-mac-address)

[Tree] (config>service>vprn>if>dhcp>option>vendor client-mac-address)

[Tree] (config>service>ies>if>dhcp>option>vendor client-mac-address)

Full Context

configure service vpls sap dhcp option vendor-specific-option client-mac-address

configure service vprn interface dhcp option vendor-specific-option client-mac-address

configure service ies interface dhcp option vendor-specific-option client-mac-address

Description

This command enables the sending of the MAC address in the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the MAC address in the Nokia vendor-specific sub-option of the DHCP relay packet.

Platforms

7705 SAR Gen 2

7.65 client-mac-list

client-mac-list

Syntax

client-mac-list

Context

[Tree] (config>system>security>ssh client-mac-list)

Full Context

configure system security ssh client-mac-list

Description

Commands in this context configure SSH MAC algorithms for SR OS as a client.

Platforms

7705 SAR Gen 2

7.66 client-name

client-name

Syntax

client-name *name*

no client-name

Context

[Tree] (config>ipsec>client-db>client client-name)

Full Context

configure ipsec client-db client client-name

Description

This command specifies the name of the client entry. The client name can be used in CLI navigation or in show commands.

Default

no client-name

Parameters

name

Specifies the name of the client.

Platforms

7705 SAR Gen 2

7.67 client-signature-list

client-signature-list

Syntax

client-signature-list *name* [**create**]

no client-signature-list *name*

Context

[Tree] (config>system>security>tls client-signature-list)

Full Context

configure system security tls client-signature-list

Description

This command configures a list of TLS 1.3-supported signature suite codes that the client sends in a client Hello message.

The **no** form of this command removes the client signature list.

Parameters

name

Specifies the name of the client signature list, up to 32 characters.

create

Keyword used to create the client signature list.

Platforms

7705 SAR Gen 2

7.68 client-tls-profile

client-tls-profile

Syntax

client-tls-profile *name*

no client-tls-profile

Context

[Tree] (config>system>security>pki>est-profile client-tls-profile)

Full Context

configure system security pki est-profile client-tls-profile

Description

This command configures the TLS client profile to be assigned to applications for encryption. The profile creates the TLS connection to the EST server.

The **no** form of this command removes the name from the configuration.

Default

no client-tls-profile

Parameters

name

Specifies the name of the client TLS profile, up to 32 characters

Platforms

7705 SAR Gen 2

client-tls-profile

Syntax

client-tls-profile *name* [**create**]

no client-tls-profile *name*

Context

[Tree] (config>system>security>tls client-tls-profile)

Full Context

configure system security tls client-tls-profile

Description

This command configures the TLS client profile to be assigned to applications for encryption.

Parameters

name

Specifies the name of the client TLS profile, up to 32 characters in length.

create

Keyword used to create the client TLS profile.

Platforms

7705 SAR Gen 2

client-tls-profile

Syntax

client-tls-profile *name*

no client-tls-profile

Context

[Tree] (config>system>management-interface>remote-management client-tls-profile)

Full Context

configure system management-interface remote-management client-tls-profile

Description

This command configures the TLS client profile used for encryption by all remote managers. This command and **allow-unsecure-connection** are mutually exclusive.

If this command is also configured for a specific manager in the **config**>**system**> **management-interface**>**remote-management**>**manager** context, that configuration takes precedence.

The **no** form of this command causes the profile configuration not to be used.

Parameters

name

Specifies the name of the client TLS profile, up to 32 characters.

Platforms

7705 SAR Gen 2

client-tls-profile

Syntax

client-tls-profile *name*

no client-tls-profile

Context

[Tree] (config>system>management-interface>remote-management>manager client-tls-profile)

Full Context

configure system management-interface remote-management manager client-tls-profile

Description

This command configures the TLS client profile used for encryption by this remote manager. This command and **allow-unsecure-connection** are mutually exclusive.

This command takes precedence over the same command configured in the global context (**config**>**system**>**management-interface**>**remote-management**).

The **no** form of this command causes the profile configuration to be inherited from the global context (**config**>**system**>**management-interface**>**remote-management**).

Parameters

name

Specifies the name of the client TLS profile, up to 32 characters.

Platforms

7705 SAR Gen 2

7.69 clli-code

clli-code

Syntax

clli-code *clli-code*

no clli-code

Context

[Tree] (config>system clli-code)

Full Context

configure system clli-code

Description

This command creates a Common Language Location Identifier (CLLI) code string for the router. A CLLI code is an 11-character standardized geographic identifier that uniquely identifies geographic locations and certain functional categories of equipment unique to the telecommunications industry.

No CLLI validity checks other than truncating or padding the string to eleven characters are performed.

Only one CLLI code can be configured, if multiple CLLI codes are configured the last one entered overwrites the previous entry.

The **no** form of the command removes the CLLI code.

Default

no clli-code

Parameters

clli-code

Specifies the 11 character string CLLI code. Any printable, seven bit ASCII characters can be used within the string. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. If more than 11 characters are entered, the string is truncated. If less than 11 characters are entered the string is padded with spaces.

Platforms

7705 SAR Gen 2

7.70 clock-offset

clock-offset

Syntax

clock-offset *seconds*

no clock-offset

Context

[Tree] (config>oam-pm>session>meas-interval clock-offset)

Full Context

configure oam-pm session meas-interval clock-offset

Description

This command allows measurement intervals with a boundary-type of clock aligned to be offset from the default time of day clock. The configured offset must be smaller than the size of the measurement interval. As an example, an offset of 120 (seconds) shifts the start times of the measurement intervals by two minutes from their default alignments with respect to the time of day clock.

The **no** form of this command sets the offset to 0.

Default

clock-offset 0

Parameters

seconds

Specifies the number of seconds to offset a clock-alignment measurement interval from its default.

Values 0 to 86399

Default 0

Platforms

7705 SAR Gen 2

7.71 close-session

close-session

Syntax

[**no**] **close-session**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization close-session)

Full Context

configure system security profile netconf base-op-authorization close-session

Description

This command enables the NETCONF <close-session> RPC.

The **no** form of this command disables the RPC.

Default

no close-session

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

7.72 cluster

cluster

Syntax

cluster *cluster-id*

no cluster

Context

[Tree] (config>service>vprn>bgp cluster)

[Tree] (config>service>vprn>bgp>group>neighbor cluster)

[Tree] (config>service>vprn>bgp>group cluster)

Full Context

configure service vprn bgp cluster

configure service vprn bgp group neighbor cluster

configure service vprn bgp group cluster

Description

This command configures the cluster ID for a route reflector server.

Route reflectors are used to reduce the number of IBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in an AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients.

When a route reflector receives a route, first it must select the best path from all the paths received. If the route was received from a non-client peer, then the route reflector sends the route to all clients in the cluster. If the route came from a client peer, the route reflector sends the route to all non-client peers and to all client peers except the originator.

For redundancy, a cluster can have multiple route reflectors.

Confederations can also be used to remove the full IBGP mesh requirement within an AS.

The **no** form of this command deletes the cluster ID and effectively disables the Route Reflection for the given group.

Default

no cluster — No cluster ID is defined.

Parameters

cluster-id

The route reflector cluster ID is expressed in dot decimal notation.

Values Any 32 bit number in dot decimal notation. (0.0.0.1 to 255.255.255.255)

Platforms

7705 SAR Gen 2

cluster

Syntax

cluster *cluster-id* **orr-location** *location-id* [**allow-local-fallback**]]

Context

[Tree] (config>router>bgp cluster)

Full Context

configure router bgp cluster

Description

This command configures the cluster ID for a route reflector server ID and implicitly configures the associated BGP sessions as route reflector clients of the BGP instance. If an ORR location ID is specified with the cluster ID, the clients in that cluster receive routes optimal for that specific location; refer to *draft-ietf-idr-bgp-optimal-route-reflection* for more information.

Route reflectors are used to reduce the number of IBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in an AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients.

When a route reflector receives best path from a non-client peer, it sends the route to all clients. When the route reflector receives a best path from a client peer it sends the route to all non-client and all client peers except the originator.

With optimal route reflection, the best path advertised to a client takes location ID into account, which means that if the tie-break for best path (or Add-Paths) comes down to next-hop IGP cost, the IGP costs will be calculated relative to the specified location. In the SR OS implementation, the IGP costs from arbitrary ORR locations are calculated using OSPF/OSPFv3, IS-IS, or BGP-LS information in the TE DB.

Default

no cluster

Parameters

ip-address

Specifies the route reflector cluster ID is expressed in dot decimal notation.

Values Any 32 bit number in dot decimal notation. (0.0.0.1 to 255.255.255.255)

orr-location location-id

Specifies the optimal route reflection location index for this set of route reflector clients.

Values 1 to 255

allow-local-fallback

Controls the behavior when there are no BGP routes to advertise to the RR clients that are reachable from the perspective of their ORR location. If this option is configured, the RR is allowed (in this circumstance only), to advertise the best reachable BGP path from its own topology location. If this option is not configured and this situation applies, then no route is advertised to the clients.

Platforms

7705 SAR Gen 2

cluster

Syntax

cluster *cluster-id* **orr-location** *location-id* [**allow-local-fallback**]]

cluster *cluster-id*

no cluster

Context

[Tree] (config>router>bgp>group cluster)

[Tree] (config>router>bgp>group>neighbor cluster)

Full Context

configure router bgp group cluster

configure router bgp group neighbor cluster

Description

This command configures the cluster ID for a route reflector server ID and implicitly configures the associated BGP sessions as route reflector clients of the BGP instance. If an ORR location ID is specified with the cluster ID, the clients in that cluster receive routes optimal for that specific location; see *draft-ietf-idr-bgp-optimal-route-reflection* for more information.

Route reflectors are used to reduce the number of IBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in an AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients.

When a route reflector receives best path from a non-client peer, it sends the route to all clients. When the route reflector receives a best path from a client peer it sends the route to all non-client and all client peers except the originator.

With optimal route reflection, the best path advertised to a client takes location ID into account, which means that if the tie-break for best path (or Add-Paths) comes down to next-hop IGP cost, the IGP costs

will be calculated relative to the specified location. In the SR OS implementation, the IGP costs from arbitrary ORR locations are calculated using OSPF/OSPFv3, IS-IS, or BGP-LS information in the TE DB.

The **no** form of this command deletes the cluster ID and effectively disables route reflection for the group.

Default

no cluster

Parameters

ip-address

Specifies the route reflector cluster ID is expressed in dot decimal notation.

Values Any 32 bit number in dot decimal notation. (0.0.0.1 to 255.255.255.255)

orr-location location-id

Specifies the optimal route reflection location index for this set of route reflector clients.

Values 1 to 255

allow-local-fallback

Controls the behavior when there are no BGP routes to advertise to the RR clients that are reachable from the perspective of their ORR location. If this option is configured, the RR is allowed (in this circumstance only), to advertise the best reachable BGP path from its own topology location. If this option is not configured and this situation applies, then no route is advertised to the clients.

Platforms

7705 SAR Gen 2

7.73 cluster-id

cluster-id

Syntax

cluster-id *ip-address/mask* [*ip-address/mask*]

cluster-id none

no cluster-id

Context

[Tree] (config>router>policy-options>policy-statement>entry>from cluster-id)

Full Context

configure router policy-options policy-statement entry from cluster-id

Description

This command enables BGP routes to be matched based on the IP addresses encoded in the CLUSTER_LIST attribute.

The first i*p-address/mask pair* is matched against the most recently added cluster ID. Each subsequent *ip-address/mask* pair is tested against the next most recent cluster ID.

For example, to match all routes reflected by the RR with cluster ID 1.1.1.1 and then any other RR before reaching the router where the policy is applied, use the command **cluster-id 0.0.0.0/0 1.1.1.1/32**.

Note:
The command matches routes with two or more cluster IDs; the third and older cluster IDs are not evaluated and are automatically considered matching.

The **cluster-id none** form of this command only matches BGP routes without any CLUSTER_LIST attribute.

A non-BGP route does not match a policy entry if it contains the **cluster-id** command.

Default

no cluster-id

Parameters

ip-address

Specifies the 32-bit cluster ID in dotted decimal notation.

Values a.b.c.d

mask

Specifies a bit mask to apply to the *ip-address* parameter.

Values 0 to 32 (0 is only allowed if the *ip-address* is 0.0.0.0)

none

Specifies that only BGP routes without a CLUSTER_LIST attribute should be matched.

Platforms

7705 SAR Gen 2

7.74 cmpv2

cmpv2

Syntax

cmpv2

Context
[Tree] (admin>certificate cmpv2)

Full Context
admin certificate cmpv2

Description
Commands in this context configure CMPv2 operations.

Platforms
7705 SAR Gen 2

cmpv2

Syntax
cmpv2

Context
[Tree] (config>system>security>pki>ca-profile cmpv2)

Full Context
configure system security pki ca-profile cmpv2

Description
Commands in this context configure CMPv2 parameters.

Platforms
7705 SAR Gen 2

cmpv2

Syntax
[**no**] **cmpv2**

Context
[Tree] (debug>certificate cmpv2)

Full Context
debug certificate cmpv2

Description
This command enables debugging of CMPv2 operations.

Platforms

7705 SAR Gen 2

7.75 coa-script-policy

coa-script-policy

Syntax

coa-script-policy *policy-name*

no coa-script-policy

Context

[Tree] (config>service>vprn>radius-server>server coa-script-policy)

[Tree] (config>router>radius-server>server coa-script-policy)

Full Context

configure service vprn radius-server server coa-script-policy

configure router radius-server server coa-script-policy

Description

This command specifies the RADIUS script policy to modify the Change-of-Authorization messages sent from this RADIUS server.

The **no** form of this command removes the policy name from the configuration.

Parameters

policy-name

Specifies the name of radius-script-policy up to 80 characters.

Platforms

7705 SAR Gen 2

7.76 code-type

code-type

Syntax

code-type [**sonet** | **sdh**]

[**no**] **code-type**

Context

[Tree] (config>port>ethernet>ssm code-type)

Full Context

configure port ethernet ssm code-type

Description

This command configures the encoding of synchronization status messages. For example, whether to use an SDH or SONET set of values. Configuring the network-type is only applicable to SyncE ports. It is not configurable on SONET/SDH ports. For the network-type, sdh refers to ITU-T G.781 Option I, while sonet refers to G.781 Option II (equivalent to Telcordia GR-253-CORE).

Default

code-type sdh

Parameters

sdh

Specifies the values used on a G.781 Option 1 compliant network.

sonet

Specifies the values used on a G.781 Option 2 compliant network.

Platforms

7705 SAR Gen 2

7.77 coherent

coherent

Syntax

coherent

Context

[Tree] (config>port>dwdm coherent)

Full Context

configure port dwdm coherent

Description

This command configures the coherent optical module parameters.

Platforms

7705 SAR Gen 2

7.78 cold-start-wait

cold-start-wait

Syntax

cold-start-wait *seconds*

no cold-start-wait

Context

[Tree] (config>log>app-route-notifications cold-start-wait)

Full Context

configure log app-route-notifications cold-start-wait

Description

The time delay that must pass before notifying specific CPM applications that a route is available after a cold reboot.

Default

no cold-start-wait

Parameters

seconds

Time delay in seconds.

Values 1 to 300

Platforms

7705 SAR Gen 2

7.79 collect-stats

collect-stats

Syntax

[**no**] **collect-stats**

Context

[Tree] (config>service>vpls>mesh-sdp collect-stats)

[Tree] (config>service>vpls>spoke-sdp collect-stats)

[Tree] (config>service>ies>if>sap collect-stats)

[Tree] (config>service>vpls>sap collect-stats)

Full Context

configure service vpls mesh-sdp collect-stats

configure service vpls spoke-sdp collect-stats

configure service ies interface sap collect-stats

configure service vpls sap collect-stats

Description

This command enables accounting and statistical data collection for either the SAP or SDP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU does not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default

collect-stats

Platforms

7705 SAR Gen 2

collect-stats

Syntax

[**no**] **collect-stats**

Context

[Tree] (config>card>fp>ingress>access>queue-group collect-stats)

[Tree] (config>card>fp>ingress>network>queue-group collect-stats)

Full Context

configure card fp ingress access queue-group collect-stats

configure card fp ingress network queue-group collect-stats

Description

This command enables the collection of accounting and statistical data for the queue group on the forwarding plane. When applying accounting policies, the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued, the statistics are still accumulated, however, the CPU does not obtain the results and write them to the billing file. If the **collect-stats** command is issued again (enabled), then the counters written to the billing file will include the traffic collected while the **no collect-stats** command was in effect.

Default

no collect-stats

Platforms

7705 SAR Gen 2

collect-stats

Syntax

[**no**] **collect-stats**

Context

[Tree] (config>port>ethernet collect-stats)

[Tree] (config>port>ethernet>access>egr>qgrp collect-stats)

[Tree] (config>port>ethernet>network>egr>qgrp collect-stats)

[Tree] (config>port>ethernet>access>ing>qgrp collect-stats)

[Tree] (config>port>ethernet>network collect-stats)

Full Context

configure port ethernet collect-stats

configure port ethernet access egress queue-group collect-stats

configure port ethernet network egress queue-group collect-stats

configure port ethernet access ingress queue-group collect-stats

configure port ethernet network collect-stats

Description

This command enables the collection of accounting and statistical data for the network interface. When applying accounting policies, the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued, the statistics are still accumulated by the XCM/IOM cards, however, the CPU does not obtain the results and write them to the billing file. If the **collect-stats** command is issued again (enabled), then the counters written to the billing file will include the traffic collected while the **no collect-stats** command was in effect.

Default

no collect-stats

Platforms

7705 SAR Gen 2

collect-stats

Syntax

[**no**] **collect-stats**

Context

[Tree] (config>service>epipe>sap collect-stats)

[Tree] (config>service>epipe>spoke-sdp collect-stats)

Full Context

configure service epipe sap collect-stats

configure service epipe spoke-sdp collect-stats

Description

This command enables accounting and statistical data collection for either the SAP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued, then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default

no collect-stats

Platforms

7705 SAR Gen 2

collect-stats

Syntax

[**no**] **collect-stats**

Context

[Tree] (config>service>ies>if>spoke-sdp collect-stats)

Full Context

configure service ies interface spoke-sdp collect-stats

Description

This command enables statistics collection.

Platforms

7705 SAR Gen 2

collect-stats

Syntax

[**no**] **collect-stats**

Context

[Tree] (config>service>vprn>if>spoke-sdp collect-stats)

[Tree] (config>service>vprn>if>sap collect-stats)

Full Context

configure service vprn interface spoke-sdp collect-stats

configure service vprn interface sap collect-stats

Description

This command enables accounting and statistical data collection for either an interface SAP or interface SAP spoke SDP, or network port. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default

no collect-stats

Platforms

7705 SAR Gen 2

collect-stats

Syntax

[**no**] **collect-stats**

Context

[Tree] (config>service>sdp collect-stats)

[Tree] (config>service>pw-template collect-stats)

Full Context

configure service sdp collect-stats

configure service pw-template collect-stats

Description

This command enables accounting and statistical data collection for either the SDP. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM or XCM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default

no collect-stats

Platforms

7705 SAR Gen 2

7.80 collection-interval

collection-interval

Syntax

collection-interval *minutes*

no collection-interval

Context

[Tree] (config>log>acct-policy collection-interval)

Full Context

configure log accounting-policy collection-interval

Description

This command configures the accounting collection interval.

Parameters

minutes

Specifies the interval between collections, in minutes.

Values 1 to 120 A range of 1 to 4 is only allowed when the record type is set to SAA.

Platforms

7705 SAR Gen 2

7.81 color

color

Syntax

color *color*

no color

Context

[Tree] (conf>router>segment-routing>sr-policies>policy color)

Full Context

configure router segment-routing sr-policies static-policy color

Description

This command associates a color value with a statically defined segment routing policy. This is a mandatory parameter and configuration command to enable the segment routing policy; if the color parameter value is not configured, the execution of the **no shutdown** command on the static segment routing policy fails.

The **no** form of this command removes the color association.

Default

no color

Parameters

color

Specifies the color ID.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

color

Syntax

color *color-id*

no color

Context

[Tree] (config>router>policy-options>policy-statement>entry>from color)

Full Context

configure router policy-options policy-statement entry from color

Description

This command configures an SR Policy color ID as a route policy match criterion.

This match criterion is only used in import policies.

The **no** form of this command removes the configuration.

Parameters

color-id

Specifies the SR policy color ID.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

color

Syntax

color *color-id*

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls>sr-policy color)

Full Context

configure oam-pm session ip tunnel mpls sr-policy color

Description

This command configures the color for associating the SR policy with an objective.

Default

color 0

Parameters

color-id

Specifies the color ID.

Values 0 to 4294967295

Default 0

Platforms

7705 SAR Gen 2

7.82 combined-max-sessions

combined-max-sessions

Syntax

combined-max-sessions *number-of-sessions*

no combined-max-sessions

Context

[Tree] (config>system>security>profile combined-max-sessions)

[Tree] (config>system>security>cli-session-group combined-max-sessions)

Full Context

configure system security profile combined-max-sessions

configure system security cli-session-group combined-max-sessions

Description

This command is used to limit the number of combined SSH/TELNET based sessions available to all users that are part of a specific profile, or to all users of all profiles that are part of the same **cli-session-group**.

The **no** form of this command disables the command and the profile or group limit is not applied to the number of combined sessions.

Default

no combined-max-sessions

Parameters

number-of-sessions

Specifies the maximum number of allowed combined SSH/TELNET based sessions.

Values 0 to 50

Platforms

7705 SAR Gen 2

7.83 command-accounting-during-load

command-accounting-during-load

Syntax

[**no**] **command-accounting-during-load**

Context

[Tree] (config>system>security>management-interface>md-cli command-accounting-during-load)

Full Context

configure system security management-interface md-cli command-accounting-during-load

Description

This command controls command accounting performed on the contents of a file loaded using the MD-CLI **load** or **rollback** command.

When enabled, all commands in the loaded file are logged, which may decrease the system response time with large files.

When disabled, command accounting is not performed during a load or rollback operation, which may increase the system response time by reducing the number of command accounting messages, especially when remote AAA servers are used.

The **load** or **rollback** command itself is always logged.

The **no** form of this command disables command accounting during a load or rollback operation.

Default

command-accounting-during-load

Platforms

7705 SAR Gen 2

7.84 command-completion

command-completion

Syntax

command-completion

Context

[Tree] (config>system>management-interface>cli>md-cli>environment command-completion)

Full Context

configure system management-interface cli md-cli environment command-completion

Description

This command configures keystrokes to trigger command completion.

Platforms

7705 SAR Gen 2

7.85 comment

comment

Syntax

[**no**] **comment**

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>commit-options comment)

Full Context

configure system management-interface cli md-cli environment commit-options comment

Description

This command configures the requirement for a commit comment when committing configuration.

The **no** form of this command does not require a commit comment when committing configuration

Default

no comment

Platforms

7705 SAR Gen 2

7.86 commit

commit

Syntax

commit

Context

[Tree] (config>router>bfd commit)

Full Context

configure router bfd commit

Description

This command saves the changes made to a BFD template during an active session and makes the changes active.

Platforms

7705 SAR Gen 2

commit

Syntax

commit

Context

[Tree] (config>router>route-next-hop-policy commit)

Full Context

configure router route-next-hop-policy commit

Description

This command saves the changes made to route next-hop templates during an active session.

Default

commit

Platforms

7705 SAR Gen 2

commit

Syntax

commit [**confirmed** *timeout*] [**comment** *comment*]

commit no-checkpoint [**confirmed** *timeout*]

Context

[Tree] (candidate commit)

Full Context

candidate commit

Description

This command applies the changes in the candidate configuration to the active running configuration. The candidate changes will take operational effect.

If a commit operation is successful then all of the candidate changes will take operational effect and the candidate is cleared. If there is an error in the processing of the commit, or a 'commit confirmed' is not confirmed and an auto-revert occurs, then the router will return to a configuration state with none of the candidate changes applied. The operator can then continue editing the candidate and try a commit later.

By default, the SR OS will automatically create a new rollback checkpoint after a commit operation. The rollback checkpoint will contain the new configuration changes made by the commit. An optional **no-checkpoint** keyword can be used to avoid the auto-creation of a rollback checkpoint after a commit.

A commit operation is blocked if a rollback revert is currently being processed.

Parameters

confirmed

specifies that the commit operation (if successful) should be automatically reverted (undone) at the end of the timeout period unless the operator issues the confirm command before the timeout period expires. A rollback checkpoint is created after the commit operation (if successful) and will remain available whether the commit is auto-reverted or not. The contents of the candidate will remain visible (candidate view) and changes to the candidate are blocked until the timeout is completed or the **candidate confirm** command is executed. If the timeout expires and an auto-revert occurs, then the original candidate config will be available in edit-cfg mode.

Standard line-by-line non-transactional configuration commands (including via SNMP) are not blocked during the countdown period and any changes made to the configuration during the countdown period will be rolled back if the timeout expires. The confirmed option is useful when changes are being made that could impact management reachability to the router.

A rollback revert is blocked during the countdown period until the commit has been confirmed.

timeout

Specifies the auto-revert timeout period, in minutes.

Values 1 to 168

no-checkpoint

Specifies to avoid the automatic creation of a rollback checkpoint for a successful commit.

comment *comment*

Adds a comment up to 255 characters to the automatic rollback checkpoint.

Platforms

7705 SAR Gen 2

commit

Syntax

commit

Context

[Tree] (config>router>policy-options commit)

Full Context

configure router policy-options commit

Description

This command is required to save changes made to a route policy.

Platforms

7705 SAR Gen 2

commit

Syntax

[**no**] **commit**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization commit)

Full Context

configure system security profile netconf base-op-authorization commit

Description

This command enables the NETCONF <commit> RPC.

The **no** form of this command disables the RPC.

Default

no commit

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

7.87 commit-options

commit-options

Syntax

commit-options

Context

[Tree] (config>system>management-interface>cli>md-cli>environment commit-options)

Full Context

configure system management-interface cli md-cli environment commit-options

Description

Commands in this context configure commit options.

Platforms

7705 SAR Gen 2

7.88 common-name-list

common-name-list

Syntax

common-name-list *name* [**create**]

Context

[Tree] (config>system>security>pki common-name-list)

Full Context

configure system security pki common-name-list

Description

This command configures a list of common names (CNs) that will be used to authenticate X.509.3 certificates. If the CN field of the X.509.3 certificate matches any of the CNs in the list, then the certificate can be used.

Parameters

name

Specifies the name of the CN list, up to 32 characters maximum.

Platforms

7705 SAR Gen 2

7.89 community

community

Syntax

community *community-name* [**hash** | **hash2** | **custom**] [**access-permissions**] [**version** *SNMP-version*] [**src-access-list** *list-name*]

no community *community-name* [**hash** | **hash2** | **custom**]

Context

[Tree] (config>service>vprn>snmp community)

Full Context

configure service vprn snmp community

Description

This command sets the SNMP community name(s) to be used with the associated VPRN instance. These VPRN community names are used to associate SNMP v1/v2c requests with a particular vprn context and to return a reply that contains VPRN-specific data or limit SNMP access to data in a specific VPRN instance.

VPRN snmp communities configured with an access permission of 'r' are automatically associated with the default access group "snmp-vprn-ro" and the "vprn-view" view (read only). VPRN snmp communities configured with an access permission of 'rw' are automatically associated with the default access group "snmp-vprn" and the "vprn-view" view (read/write).

The community in an SNMP v1/v2 request determines the SNMP context (i.e., the vprn# for accessing SNMP tables) and not the VPRN of the incoming interface on which the request was received. When an SNMP request arrives on VPRN 5 interface "ringo" with a destination IP address equal to the "ringo"

interface, but the community in the SNMP request is the community configured against VPRN 101, then the SNMP request will be processed using the VPRN 101 context. (the response will contain information about VPRN 101). It is recommended to avoid using a simple series of vprn snmp-community values that are similar to each other (for example, avoid my-vprncomm-1, my-vprn-comm-2, etc).

The **no** form of this command removes the SNMP community name from the given VPRN context.

Parameters

community-name

Specifies the SNMP v1/v2c community name. This is a secret/confidential key used to access SNMP and specify a context (base vs vprn1 vs vprn2).

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

version *SNMP-version*

Specifies the SNMP version.

Values v1, v2c, both

access-permissions

Specifies the access rights to MIB objects.

Values **r** — Grants only read access to MIB objects. Creates an association of the community-name with the **snmp-vprn-ro** access group.**rw** — Grants read and write access to MIB objects. Creates an association of the community-name with the **snmp-vprn** access group.

list-name

Configures the **community** to reference a specific **src-access-list** (created under **configure system security snmp**), which will be used to validate the source IP address of all received SNMP requests that use this **community**. Multiple **community** (vprn or base router) and **usm-community** instances can reference the same **src-access-list**.

Platforms

7705 SAR Gen 2

community

Syntax

community *comm-id* [*comm-id*]

no community [*comm-id* [*comm-id*]]

Context

[Tree] (config>service>vprn>static-route-entry community)

Full Context

configure service vprn static-route-entry community

Description

This command associates a list of up to 12 BGP communities (any mix of standard, extended, and large communities) with the static route. These communities can be matched in route policies and are automatically added to BGP routes that are created from the static route.

The communities specified at this level of the static route causes communities configured under the next-hop, black-hole, and indirect contexts of the static route to be ignored.

The **no** form of this command removes the association.

Default

no community

Parameters

comm-id

Specifies a BGP community value, up to 72 characters.

Values [*as-num:comm-val* | *well-known-comm* | *ext-comm* | *large-comm*]

where:

- *as-num* — 0 to 65535
- *comm-val* — 0 to 65535
- *well-known-comm* — **null** | **no-export** | **no-export-subconfed** | **no-advertise** | **llgr-stale** | **no-llgr** | **blackhole**
- *ext-comm* — the extended community, defined as one of the following:
  - {*target* | *origin*}:*ip-address*:*comm-val*
  - {*target* | *origin*}:*asnum*:*ext-comm-val*
  - {*target* | *origin*}:*ext-asnum*:*comm-val*
  - **bandwidth**:*asnum*:*val-in-mbps*
  - **ext:4300**:*ovstate*
  - **ext**:*value1*:*value2*

– color:co-bits:*color-value*

where:

– *target* — route target
– *origin* — route origin
– *ip-address* — a.b.c.d
– *ext-comm-val* — 0 to 4294967295
– *ext-asnum* — 0 to 4294967295
– *val-in-mbps* — 0 to 16777215
– *ovstate* — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid)
– *value1* — 0000 to FFFF
– *value2* — 0 to FFFFFFFFFFFF
– *co-bits* — 00, 01, 10 or 11
– *color-value* — 0 to 4294967295

- *large-comm — asn-or-ex:val-or-ex:val-or-ex*

Platforms

7705 SAR Gen 2

community

Syntax

community *comm-id*

no community [*comm-id*]

Context

[Tree] (config>service>vprn>static-route-entry>next-hop community)

[Tree] (config>service>vprn>static-route-entry>indirect community)

[Tree] (config>service>vprn>static-route-entry>black-hole community)

Full Context

configure service vprn static-route-entry next-hop community

configure service vprn static-route-entry indirect community

configure service vprn static-route-entry black-hole community

Description

This command associates one BGP community (standard, extended or large) with a next-hop of the static route. This community can be matched in route policies and automatically added to BGP routes that are created from the static route.

Any community specified in one of these contexts is overridden by any communities specified at the prefix level of the static route entry.

The **no** form of this command removes the association.

Default

no community

Parameters

comm-id

Specifies a BGP community value, up to 72 characters.

Values [*as-num:comm-val* | *well-known-comm* | *ext-comm* | *large-comm*]

where:

- *as-num* — 0 to 65535
- *comm-val* — 0 to 65535
- *well-known-comm* — **null** | **no-export** | **no-export-subconfed** | **no-advertise** | **llgr-stale** | **no-llgr** | **blackhole**
- *ext-comm* — the extended community, defined as one of the following:
  - {*target* | *origin*}:*ip-address*:*comm-val*
  - {*target* | *origin*}:*asnum*:*ext-comm-val*
  - {*target* | *origin*}:*ext-asnum*:*comm-val*
  - **bandwidth**:*asnum*:*val-in-mbps*
  - **ext:4300**:*ovstate*
  - **ext**:*value1*:*value2*
  - color:co-bits:*color-value*

  where:
  - *target* — route target
  - *origin* — route origin
  - *ip-address* — a.b.c.d
  - *ext-comm-val* — 0 to 4294967295
  - *ext-asnum* — 0 to 4294967295
  - *val-in-mbps* — 0 to 16777215
  - *ovstate* — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid)
  - *value1* — 0000 to FFFF
  - *value2* — 0 to FFFFFFFFFFFF
  - *co-bits* — 00, 01, 10 or 11
  - *color-value* — 0 to 4294967295
- *large-comm* — *asn-or-ex:val-or-ex:val-or-ex*

Platforms

7705 SAR Gen 2

community

Syntax

community *comm-id*

no community [*comm-id*]

Context

[Tree] (config>service>vprn>static-route-entry>ipsec-tunnel community)

Full Context

configure service vprn static-route-entry ipsec-tunnel community

Description

This configuration option associates a BGP community with the static route. The community can be matched in route policies and is automatically added to BGP routes exported from the static route.

The **no** form of this command removes the community association.

Default

no community

Parameters

comm-id

Specifies community IDs, up to 72 characters.

Values [*2 byte asnumber:comm-val* | *well-known-comm*]

where:

- *2 byte as-number* — 0 to 65535
- *comm-val* — 0 to 65535
- *well-known-comm* — **no-export** | **no-export-subconfed** | **no-advertise**

Platforms

7705 SAR Gen 2

community

Syntax

community *community-name*

no community

Context

[Tree] (config>router>ldp>session-params>peer community)

[Tree] (config>router>ldp>targeted-session>peer-template community)

Full Context

configure router ldp session-parameters peer community

configure router ldp targeted-session peer-template community

Description

This command configures a community name associated with a targeted session to a specified peer. The community is a local configuration for a targeted session. FECs received over a session of a given community are taken to belong to that community, and are redistributed over sessions of the same community.

The SR OS router uses the following rules for community:

- If both the session parameters for a specified peer and targeted peer template that is applied to session have the default configuration then no community applies.
- If the session parameters for a peer have the default configuration, but targeted session peer template has an explicit configuration for community, then the targeted peer template configuration will be used.
- If the session parameters have an explicit configuration for community, and the targeted session peer template has the default configuration, then the session parameter configuration applies.
- If both session parameters and targeted peer template have an explicit configuration for community, then the session parameter configuration is used.

The **no** form of this command removes the community from the session to the peer. FEC subsequently received over the session are treated as having no community.

Default

no community

Parameters

community-name

Specifies the string defining the LDP community assigned to the session. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string.

Platforms

7705 SAR Gen 2

community

Syntax

community *comm-id*

no community [*comm-id*]

Context

[Tree] (config>router>static-route-entry>black-hole community)

[Tree] (config>router>static-route-entry>indirect community)

[Tree] (config>router>static-route-entry>next-hop community)

Full Context

configure router static-route-entry black-hole community

configure router static-route-entry indirect community

configure router static-route-entry next-hop community

Description

This command associates one BGP community (standard, extended or large) with a next-hop of the static route. This community can be matched in route policies and automatically added to BGP routes that are created from the static route.

Any community specified in one of these contexts is overridden by any communities specified at the prefix level of the static route entry.

The **no** form of this command removes the association.

Default

no community

Parameters

comm-id

Specifies a BGP community value, up to 72 characters.

Values [*as-num:comm-val* | *well-known-comm* | *ext-comm* | *large-comm*]

where:

- *as-num* — 0 to 65535
- *comm-val* — 0 to 65535
- *well-known-comm* — **null** | **no-export** | **no-export-subconfed** | **no-advertise** | **llgr-stale** | **no-llgr** | **blackhole**
- *ext-comm* — the extended community, defined as one of the following:
  - {*target* | *origin*}:*ip-address*:*comm-val*
  - {*target* | *origin*}:*asnum*:*ext-comm-val*

– {*target* | *origin*}:*ext-asnum*:*comm-val*
– **bandwidth**:*asnum*:*val-in-mbps*
– **ext:4300**:*ovstate*
– **ext**:*value1*:*value2*
– color:co-bits:*color-value*

where:

– *target* — route target
– *origin* — route origin
– *ip-address* — a.b.c.d
– *ext-comm-val* — 0 to 4294967295
– *ext-asnum* — 0 to 4294967295
– *val-in-mbps* — 0 to 16777215
– *ovstate* — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid)
– *value1* — 0000 to FFFF
– *value2* — 0 to FFFFFFFFFFFF
– *co-bits* — 00, 01, 10 or 11
– *color-value* — 0 to 4294967295

- *large-comm — asn-or-ex:val-or-ex:val-or-ex*

Platforms

7705 SAR Gen 2

community

Syntax

community *comm-id* [*comm-id*]

no community [*comm-id* [*comm-id*]]

Context

[Tree] (config>router>static-route-entry community)

Full Context

configure router static-route-entry community

Description

This command associates a list of up to 12 BGP communities (any mix of standard, extended, and large communities) with the static route. These communities can be matched in route policies and are automatically added to BGP routes that are created from the static route.

The communities specified at this level of the static route causes communities configured under the next-hop, black-hole and indirect contexts of the static route to be ignored.

The **no** form of this command removes the association.

Default

no community

Parameters

comm-id

Specifies a BGP community value, up to 72 characters.

Values [*as-num:comm-val* | *well-known-comm* | *ext-comm* | *large-comm*]

where:

- *as-num* — 0 to 65535
- *comm-val* — 0 to 65535
- *well-known-comm* — **null** | **no-export** | **no-export-subconfed** | **no-advertise** | **llgr-stale** | **no-llgr** | **blackhole**
- *ext-comm* — the extended community, defined as one of the following:
  - {*target* | *origin*}:*ip-address*:*comm-val*
  - {*target* | *origin*}:*asnum*:*ext-comm-val*
  - {*target* | *origin*}:*ext-asnum*:*comm-val*
  - **bandwidth**:*asnum*:*val-in-mbps*
  - **ext:4300**:*ovstate*
  - **ext**:*value1*:*value2*
  - color:co-bits:*color-value*

  where:
  - *target* — route target
  - *origin* — route origin
  - *ip-address* — a.b.c.d
  - *ext-comm-val* — 0 to 4294967295
  - *ext-asnum* — 0 to 4294967295
  - *val-in-mbps* — 0 to 16777215
  - *ovstate* — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid)
  - *value1* — 0000 to FFFF
  - *value2* — 0 to FFFFFFFFFFFF
  - *co-bits* — 00, 01, 10 or 11
  - *color-value* — 0 to 4294967295
- *large-comm* — *asn-or-ex:val-or-ex:val-or-ex*

Platforms

7705 SAR Gen 2

community

Syntax

community *community-string* [**hash** | **hash2** | **custom**] *access-permissions* [**version** *SNMP-version*] [**src-access-list** *list-name*]

no community *community-string* [**hash** | **hash2** | **custom**]

Context

[Tree] (config>system>security>snmp community)

Full Context

configure system security snmp community

Description

This command creates SNMP community strings for SNMPv1 and SNMPv2c access. This command is used in combination with the predefined access groups and views. To create custom access groups and views and associate them with SNMPv1 or SNMPv2c access use the **usm-community** command.

When configured, community implies a security model for SNMPv1 and SNMPv2c only.

For SNMPv3 security, the **access group** command must be configured.

The **no** form of the command removes the specified community string.

Parameters

community-string

Configures the SNMPv1 and/or SNMPv2c community string.

Values community-string — Specifies the community string. Allowed values are any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (for example, #, $, spaces), the entire string must be enclosed within double quotes.

hash-key — Up to 33 characters

hash2-key — Up to 96 characters

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be

in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

access-permissions

Configures the access permissions for objects in the MIB.

r — Grants only read access to objects in the MIB, except security objects, using the internal "snmp-ro" access group and the "no-security" snmp view.

rw — Grants read and write access to all objects in the MIB, using the internal "snmp-rw" access group and the "no-security" snmp view.

rwa — Grants read and write access to all objects in the MIB, including security, using the internal snmp-rwa access group and the iso snmp view.

mgmt — Assigns a unique SNMP community string for SNMP access via the management router instance. This community uses the internal snmp-mgmt access group and the mgmt snmp view.

vpls-mgmt — Assigns a unique SNMP community string for SNMP access via the vpls-management router instance. This community uses the internal snmp-vpls-mgmt access group and mgmt-view snmp view.

version {v1 | v2c | both}

Configures the scope of the community string to be for SNMPv1, SNMPv2c, or both SNMPv1 and SNMPv2c access.

Default both

list-name

Configures the **community** to reference a specific **src-access-list**, which will be used to validate the source IP address of all received SNMP requests that use this community. Multiple community, usm-community, or VPRN SNMP community instances can reference the same src-access-list.

Platforms

7705 SAR Gen 2

community

Syntax

[**no**] **community** *name*

Context

[Tree] (config>router>policy-options community)

Full Context

configure router policy-options community

Description

This command creates a route policy community list or expression to use in route policy entries. A community list is an unordered set of community values (members). In general a route matches a community list if it has any of the member values. A community expression is a set of community values that are arranged in a logical expression using operators such as AND, OR, and NOT. A route matches a community expression if it satisfies the logic of the expression.

For additional information, see the **expression** and **members** commands in the **config**> **router**>**policy-options**>**community** context.

The **no** form of this command deletes the community list or the provided community ID.

Default

no community

Parameters

name

Specifies the community list name. Allowed values are any string up to 64 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (for example, #, $, spaces), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

community

Syntax

community add *name* [*name*]

community remove *name* [*name*]

community replace *name* [*name*]

no community

Context

[Tree] (config>router>policy-options>policy-statement>default-action community)

[Tree] (config>router>policy-options>policy-statement>entry>action community)

Full Context

configure router policy-options policy-statement default-action community

configure router policy-options policy-statement entry action community

Description

This command adds or removes a BGP community list to or from routes matching the route policy statement entry.

If no community list is specified, the community path attribute is not changed.

The community list changes the community path attribute according to the **add** and **remove** keywords.

The **no** form of this command disables the action to edit the community path attribute for the route policy entry.

Default

no community

Parameters

name

Specifies up to 28 names.

add

The specified community list is added to any existing list of communities.

remove

The specified community list is removed from the existing list of communities.

replace

The specified community list replaces any existing community attribute. name — The community list name. Allowed values are any string up to 64 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@," "start@variable@end"," @variable@end", or "start@variable@".

Platforms

7705 SAR Gen 2

community

Syntax

community *comm-name*

community expression *expression*

no community

Context

[Tree] (config>router>policy-options>policy-statement>entry>from community)

Full Context

configure router policy-options policy-statement entry from community

Description

This command configures a community list as a match criterion for the route policy entry.

If no community list is specified, any community is considered a match.

The **no** form of this command removes the community list match criterion.

Default

no community

Parameters

comm-name

Specifies the community list name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

expression

Specifies that the parameters are applied to routes matching the entry.

Values *expression* is one of the following, up to 900 characters:

<expression> {AND| OR} <expression>

[NOT] (<expression>)

[NOT] "["<comm-name> "]

The following are examples of valid logical expressions:

- "[community_list_A] OR ([community_list_B] AND [community_list_C])"
- "NOT [community_list_A]”
- "[community_list_A] AND [community_list_B] OR [community_list_C]”
- "NOT ([community_list_A] OR [community_list_B] OR [community_list_C])”

Platforms

7705 SAR Gen 2

7.90 community-count

community-count

Syntax

community-count *count* [**equal** | **or-higher** | **or-lower**] [**standard** | **extended** | **large**]

no community-count

Context

[Tree] (config>router>policy-options>policy-statement>entry>from community-count)

Full Context

configure router policy-options policy-statement entry from community-count

Description

This command matches BGP routes based on community length (that is, the number of community members in the COMMUNITY, EXTENDED_COMMUNITY, or LARGE_COMMUNITY the attributes).

If no comparison qualifiers are present (**equal**, **or-higher**, **or-lower**), then **equal** is the implied default.

Without the optional **standard**, **extended**, or **large** keyword, the community length applies to the total number of communities, of all types. If some keywords are present, then only the types specified are counted against the limit.

A non-BGP route does not match a policy entry if it contains the **community-count** command.

Default

no community-count

Parameters

count

Specifies the number of community members.

Values 0 to 1024, or a parameter, up to 32 characters, name delimited by a starting and ending at-sign (@) character

equal

Specifies that matched routes should have the same number of AS path elements as the value specified.

or-higher

Specifies that matched routes should have the same or a greater number of community members as the value specified.

or-lower

Specifies that matched routes should have the same or a lower number of community members as the value specified.

standard

Specifies that only communities in the COMMUNITY attribute should be counted.

extended

Specifies that only communities in the EXTENDED_COMMUNITY attribute should be counted.

large

Specifies that only communities in the LARGE_COMMUNITY attribute should be counted.

Platforms

7705 SAR Gen 2

7.91 compare

compare

Syntax

compare *source1* **to** *source2*

Context

[Tree] (admin compare)

Full Context

admin compare

Description

This command displays the differences between rollback checkpoints and the active operational configuration, with source1 as the base/first file to which source2 is compared.

A compare operation does not check authorization of each line of output. Permission to execute the compare operation from the admin branch of CLI (authorization for the **admin rollback compare** or **admin compare** command itself) should only be given to users who are allowed to view the entire configuration, similar to permissions for **admin display-config**.

Default

The defaults for source1 and source2 are context aware and differ based on the branch in which the command is executed. In general, the default for source1 matches the context from which the command is issued.

- In the admin node: No defaults. source1 and source2 must be specified.
- In the admin>rollback node:

  source1 default = active-cfg, source2 default = latest-rb

  compare: equivalent to "compare active-cfg to latest-rb”

  compare to source2: equivalent to "compare active-cfg to source2”
- In a config>*xx* node:

  compare to source2: equivalent to "compare active-cfg to source2”

Parameters

source1, source2

Specifies comparison information.

Values **active-cfg** — The current operational configuration that is active in the node.

latest-rb — The most recent rollback checkpoint (the checkpoint file at the configured rollback-location with "*.rb” as the suffix).

rescue — The rescue configuration (at the configured rescue-location).

checkpoint-id — An ID indicating a specific rollback checkpoint. A checkpoint-id of 1 indicates the rollback checkpoint file (at the configured rollback-location) with "*.rb.1" as the suffix, 2 for file "*.rb.2", and so on.

Platforms

7705 SAR Gen 2

compare

Syntax

compare [**to** *checkpoint2*]

compare *checkpoint1* **to** *checkpoint2*

Context

[Tree] (admin>rollback compare)

Full Context

admin rollback compare

Description

This command can be used in any branch under configure, but not with configure itself. The command syntax, parameter names, and default values are context aware and will differ based on the branch in which the command is executed.

This command displays the differences between rollback checkpoints and the active operational configuration, with checkpoint1 as the base/first file to which checkpoint2 is compared. This command displays the comparison for the configuration context where it is entered and all branches below that context level.

A compare operation does not check authorization of each line of output. Permission to execute the compare operation from the admin branch of CLI (authorization for the **admin rollback compare** or **admin compare** command itself) should only be given to users who are allowed to view the entire configuration, similar to permissions for **admin display-config**.

Default

The defaults for checkpoint1 and checkpoint2 are context-aware and differ based on the branch in which the command is executed. In general, the default for checkpoint1 matches the context from which the command is issued.

- In the admin node: No defaults. checkpoint1 and checkpoint2 must be specified.
- In the admin>rollback node:

  checkpoint1 default = active-cfg, checkpoint2 default = latest-rb

  compare: equivalent to "compare active-cfg to latest-rb"

  compare to checkpoint2: equivalent to "compare active-cfg to checkpoint2"

- In a config>*xx* node:

  compare to checkpoint2: equivalent to "compare active-cfg to checkpoint2”

Parameters

checkpoint1, checkpoint2

Specifies comparison information.

Values **active-cfg** — The current operational configuration that is active in the node.

latest-rb — The most recent rollback checkpoint (the checkpoint file at the configured rollback-location with "*.rb” as the suffix).

rescue — The rescue configuration (at the configured rescue-location).

checkpoint-id — An ID indicating a specific rollback checkpoint. A checkpoint-id of 1 indicates the rollback checkpoint file (at the configured rollback-location) with "*.rb.1” as the suffix, 2 for file "*.rb.2”, and so on.

Platforms

7705 SAR Gen 2

7.92 compare-chain-include

compare-chain-include

Syntax

compare-chain-include *ca-profile-name*

no compare-chain-include

Context

[Tree] (config>ipsec>cert-profile>entry compare-chain-include)

Full Context

configure ipsec cert-profile entry compare-chain-include

Description

This command configures the Certificate Authority (CA) profile that needs to be included in the compare-chain for the entry. This configuration is required in instances where there are multiple overlapping compare-chains, for example, the configured root CA is cross-signed by another CA.

Default

no compare-chain-include

Parameters

ca-profile-name

Specifies the name of the CA profile.

Platforms

7705 SAR Gen 2

7.93 compare-origin-validation-state

compare-origin-validation-state

Syntax

[**no**] **compare-origin-validation-state**

Context

[Tree] (config>service>vprn>bgp>best-path-selection compare-origin-validation-state)

Full Context

configure service vprn bgp best-path-selection compare-origin-validation-state

Description

This command enables the comparison of origin validation states during the BGP decision process. When this command is configured, a new step is inserted in the BGP decision process after the removal of invalid routes and before the comparison of Local Preference. This step compares the origin validation state so a BGP route with a "Valid" state is preferred over a BGP route with a "Not-Found" state. A BGP route with a "Not-Found" state is preferred over a BGP route with an "Invalid" state assuming that these routes are considered "usable".

This comparison only applies to BGP routes learned from VPRN BGP peers. It does not apply to any comparison involving BGP-VPN routes that have been imported into the VPRN.

The **no** form of this command causes the new step to be skipped during the BGP decision process.

Default

no compare-origin-validation-state

Platforms

7705 SAR Gen 2

compare-origin-validation-state

Syntax

[**no**] **compare-origin-validation-state**

Context

[Tree] (config>router>bgp>best-path-selection compare-origin-validation-state)

Full Context

configure router bgp best-path-selection compare-origin-validation-state

Description

When this command is configured, a new step is inserted in the BGP decision process after removal of invalid routes and before the comparison of Local Preference. The new step compares the RPKI origin validation state so that a BGP route with a 'Valid' state is preferred over a BGP route with a 'Not-Found' state, and a BGP route with a 'Not-Found' state is preferred over a BGP route with an 'Invalid' state assuming that these routes are considered 'usable'.

The new step is skipped when **no compare-origin-validation-state** is configured.

Default

no compare-origin-validation-state

Platforms

7705 SAR Gen 2

7.94 compatibility

compatibility

Syntax

compatibility *mode*

Context

[Tree] (config>port>dwdm>coherent compatibility)

Full Context

configure port dwdm coherent compatibility

Description

This command configures the optical mode and rate of operation.

Parameters

mode

Specifies the optical mode.

Values **long-haul** - The port operates in the native long-haul mode.

long-haul-non-diff - The port operates in the native long-haul mode using non-differential encoding.

metro - The port operates in the native metro regional mode.

access - The port operates in the native access mode (80km reach).

interop - The port operates in the third party interop mode.

interop2 - The port operates in the third party interop mode with alternate differential encoding.

interop3 - The port operates in the CFP2-DCO Rev A0 Staircase FEC interop mode.

oif-400g-zr - The port operates in compliance with the OIF 400G ZR implementation agreement (IA). This parameter is only supported for use with 400G ZR and 400G ZR+ pluggable transceiver modules.

open-zrp-ofec1 - The port operates in compliance with the OpenZR + multi-source agreement (MSA) (100GHz spacing). This parameter is only supported for use with 400G ZR and 400G ZR+ pluggable transceiver modules.

open-zrp-ofec2 - The port operates in compliance with the OpenZR+ MSA (75 GHz spacing). This parameter is only supported for use with 400G ZR and 400G ZR+ pluggable transceiver modules.

Default long-haul

Platforms

7705 SAR Gen 2

7.95 compatible-rfc1583

compatible-rfc1583

Syntax

[**no**] **compatible-rfc1583**

Context

[Tree] (config>service>vprn>ospf compatible-rfc1583)

Full Context

configure service vprn ospf compatible-rfc1583

Description

This command enables OSPF summary and external route calculations in compliance with RFC 1583 and earlier RFCs.

RFC 1583 and earlier RFCs use a different method to calculate summary and external route costs. To avoid routing loops, all routers in an OSPF domain should perform the same calculation method.

Although it would be favorable to require all routers to run a more current compliance level, this command allows the router to use obsolete methods of calculation.

This command is not supported in OSPF3.

The **no** form of this command enables the post-RFC 1583 method of summary and external route calculation.

Default

compatible-rfc1583 — RFC 1583 compliance is enabled.

Platforms

7705 SAR Gen 2

compatible-rfc1583

Syntax

[**no**] **compatible-rfc1583**

Context

[Tree] (config>router>ospf compatible-rfc1583)

Full Context

configure router ospf compatible-rfc1583

Description

This command enables OSPF summary and external route calculations in compliance with RFC 1583 and earlier RFCs.

RFC 1583 and earlier RFCs use a different method to calculate summary and external route costs. To avoid routing loops, all routers in an OSPF domain should perform the same calculation method.

Although it would be favorable to require all routers to run a more current compliance level, this command allows the router to use obsolete methods of calculation.

The **no** form of this command enables the post-RFC 1583 method of summary and external route calculation.

Default

compatible-rfc1583

Platforms

7705 SAR Gen 2

7.96 complexity-rules

complexity-rules

Syntax

complexity-rules

Context

[Tree] (config>system>security>password complexity-rules)

Full Context

configure system security password complexity-rules

Description

This command defines a list of rules for configurable password options.

Note:
This command applies to local users.

Platforms

7705 SAR Gen 2

7.97 conditional-expression

conditional-expression

Syntax

conditional-expression

Context

[Tree] (config>router>policy-options>policy-statement>entry conditional-expression)

Full Context

configure router policy-options policy-statement entry conditional-expression

Description

This command creates the context to configure a route existence expression.

Platforms

7705 SAR Gen 2

7.98 confederation

confederation

Syntax

confederation *confed-as-num* [**members** *as-number* [*as-number*]]

no confederation *confed-as-num* **members** *as-number* [*as-number*]

no confederation

Context

[Tree] (config>service>vprn confederation)

Full Context

configure service vprn confederation

Description

This command configures the VPRN BGP instance to participate in a BGP confederation. BGP confederations can be used to reduce the number of IBGP sessions required within an AS.

When a VPRN BGP instance is part of a confederation, it can form confederation-EBGP sessions with CE router peers in a different sub-autonomous systems of the same confederation as well as regular EBGP sessions with CE router peers outside the confederation. A VPRN BGP instance that is part of a confederation cannot import or export its routes to the base router instance (as VPN-IP routes).

The **no** form of this command deletes the specified member AS from the confederation. When members are not specified in the no statement, the entire list is removed and confederations is disabled. When the last member of the list is removed, confederations is disabled.

Default

no confederation

Parameters

confed-as-num

The confederation AS number defined as a decimal value.

Values 1 to 4294967295

members *as-number*

The AS number(s) that are members of the confederation, each expressed as a decimal integer. Configure up to 15 members per confed-as-num.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

confederation

Syntax

confederation *confed-as-num* [**members** *as-number* [*as-number*]]

no confederation *confed-as-num* **members** *as-number* [*as-number*]

no confederation

Context

[Tree] (config>router confederation)

Full Context

configure router confederation

Description

This command creates confederation autonomous systems within an AS.

This technique is used to reduce the number of IBGP sessions required within an AS. Route reflection is another technique that is commonly deployed to reduce the number of IBGP sessions.

The **no** form of this command deletes the specified member AS from the confederation.

When no members are specified in the **no** statement, the entire list is removed and **confederation** is disabled.

When the last member of the list is removed, **confederation** is disabled.

Default

no confederation - no confederations are defined.

Parameters

confed-as-num

Specifies the confederation AS number expressed as a decimal integer.

Values 1 to 65535

as-number

Specifies the AS number of members that are part of the confederation, expressed as a decimal integer. Up to 15 members per *confed-as-num* can be configured.

Values 1 to 65535

Platforms

7705 SAR Gen 2

7.99 config-backup

config-backup

Syntax

config-backup *count*

no config-backup

Context

[Tree] (config>system config-backup)

Full Context

configure system config-backup

Description

This command configures the maximum number of backup versions maintained for configuration files and BOF.

For example, assume the **config-backup** *count* is set to 5 and the configuration file is called *xyz.cfg.* When the configuration is saved, the file *xyz.cfg* is saved with a 1 extension. Each configuration save increments the numeric extension until the maximum count is reached.

xyz.cfg xyz.cfg.1 xyz.cfg.2 xyz.cfg.3 xyz.cfg.4 xyz.cfg.5

Each classic CLI persistent index file is updated at the same time as the associated configuration file. When the index file is updated, then the save is performed to *xyz.cfg* and the index file is created as *xyz.ndx*. Synchronization between the active and standby CPM is performed for all configurations and their associated persistent index files.

The **no** form of the command returns the configuration to the default value.

Default

config-backup 50

Parameters

count

Specifies the maximum number of backup revisions.

Values 1 to 200

Platforms

7705 SAR Gen 2

7.100 configuration-mode

configuration-mode

Syntax

configuration-mode {**classic** | **mixed** | **model-driven**}

Context

[Tree] (config>system>management-interface configuration-mode)

Full Context

configure system management-interface configuration-mode

Description

This command controls which management interfaces are used for editing and changing the configuration of the router.

Any management interface can be used in any configuration mode (to gather state information or perform operations, for example), but only specific management interfaces (CLI, NETCONF, and so on) are allowed to edit the configuration of the router in different modes. For example, only classic CLI and SNMP can be used to edit the configuration when in classic mode.

Default

configuration-mode model-driven

Parameters

classic

Enables editing of router configuration via classic CLI and SNMP management interfaces, but not using model-driven interfaces.

model-driven

Enables editing of router configuration via model-driven management interfaces (NETCONF with 'Nokia' YANG models, MD-CLI or gRPC), but not using classic interfaces.

mixed

Enables editing of router configuration using a mix of classic CLI and/or model-driven management interfaces (with some restrictions and limitations).

Platforms

7705 SAR Gen 2

7.101 configure

configure

Syntax

configure

Context

[Tree] (configure)

Full Context

configure

Description

Commands in this context edit the system configuration.

Platforms

7705 SAR Gen 2

7.102 confirm

confirm

Syntax

confirm

Context

[Tree] (candidate confirm)

Full Context

candidate confirm

Description

This command is used to stop an automatic reversion to the previous configuration after the **candidate commit confirmed** command was used. If the **confirm** command is not executed before the commit confirmed timeout period expires then the previous commit changes will be undone and the previous candidate configuration will be available for editing and a subsequent commit.

During the countdown the contents of the candidate will remain visible (candidate view) and changes to the candidate are blocked until the timeout is completed or the candidate confirm command is executed. Executing the **confirm** command clears the contents of the candidate and allows editing of the candidate.

Platforms

7705 SAR Gen 2

confirm

Syntax

[**no**] **confirm**

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>commit-options confirm)

Full Context

configure system management-interface cli md-cli environment commit-options confirm

Description

This command configures the requirement for a confirmed commit when committing configuration.

The **no** form of this command does not require a confirmed commit when committing configuration

Default

no confirm

Platforms

7705 SAR Gen 2

7.103 connect-retry

connect-retry

Syntax

connect-retry *seconds*

no connect-retry

Context

[Tree] (config>service>vprn>bgp>group connect-retry)

[Tree] (config>service>vprn>bgp>group>neighbor connect-retry)

[Tree] (config>service>vprn>bgp connect-retry)

Full Context

configure service vprn bgp group connect-retry

configure service vprn bgp group neighbor connect-retry

configure service vprn bgp connect-retry

Description

This command configures the BGP connect retry timer value in seconds.

When this timer expires, BGP tries to reconnect to the configured peer. This configuration parameter can be set at three levels: global level (applies to all peers), peer-group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

120 seconds

Parameters

seconds

Specifies the BGP connect retry timer value in seconds, expressed as a decimal integer.

Values 1 to 65535

Platforms

7705 SAR Gen 2

connect-retry

Syntax

connect-retry *seconds*

no connect-retry

Context

[Tree] (config>router>origin-validation>rpki-session connect-retry)

Full Context

configure router origin-validation rpki-session connect-retry

Description

This command configures the time in seconds to wait between one TCP connection attempt that fails and the next attempt. The default (with **no connect-retry**) is 120 seconds.

Default

no connect-retry

Parameters

seconds

Specifies time in seconds.

Values 1 to 65535

Platforms

7705 SAR Gen 2

connect-retry

Syntax

connect-retry *seconds*

no connect-retry

Context

[Tree] (config>router>bgp>group>neighbor connect-retry)

[Tree] (config>router>bgp connect-retry)

[Tree] (config>router>bgp>group connect-retry)

Full Context

configure router bgp group neighbor connect-retry

configure router bgp connect-retry

configure router bgp group connect-retry

Description

This command configures the BGP connect retry timer value in seconds.

When this timer expires, BGP tries to reconnect to the configured peer. This configuration parameter can be set at three levels: global level (applies to all peers), peer-group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

connect-retry 120

Parameters

seconds

The BGP Connect Retry timer value in seconds expressed as a decimal integer.

Values 1 to 65535

Platforms

7705 SAR Gen 2

7.104 connection

connection

Syntax

[**no**] **connection** *ip-address*

Context

[Tree] (debug>router>pcep>pcc connection)

Full Context

debug router pcep pcc connection

Description

This command debugs PCC connection events.

The **no** form of this command disables debugging.

Parameters

ip-address

Specifies the IP address.

Values ipv4-prefix: a.b.c.d

ipv6-prefix:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF] H
- d: [0 to 255] D

Platforms

7705 SAR Gen 2

7.105 connection-profile-vlan

connection-profile-vlan

Syntax

connection-profile-vlan *conn-prof-id* [**create**]

no connection-profile-vlan *conn-prof-id*

Context

[Tree] (config connection-profile-vlan)

Full Context

configure connection-profile-vlan

Description

Commands in this context configure the VLAN ranges that will be associated with a service SAP.

Each connection-profile-vlan must be explicitly configured.

Parameters

conn-prof-id

Specifies the connection-profile identifier. This value will be configured in the service along with the SAP when the user associates a VLAN bundle to a single SAP. For example, a SAP defined in a dot1q port 1/1/1 that matches all the VLANs defined in the connection-profile-vlan 1 will be created as '**sap 1/1/1:cp-1 create**'.

Values 1 to 8000

Platforms

7705 SAR Gen 2

7.106 connection-timeout

connection-timeout

Syntax

connection-timeout *seconds*

no connection-timeout

Context

[Tree] (config>system>management-interface>remote-management connection-timeout)

Full Context

configure system management-interface remote-management connection-timeout

Description

This command configures the amount of time that all remote managers cannot be reached before they are considered to be down.

If this command is also configured for a specific manager in the **config**>**system**> **management-interface**>**remote-management**>**manager** context, that configuration takes precedence.

The **no** form of this command reverts to the default.

Default

connection-timeout 60

Parameters

seconds

Specifies the connection timeout in seconds.

Values 1 to 3600

Platforms

7705 SAR Gen 2

connection-timeout

Syntax

connection-timeout *seconds*

no connection-timeout

Context

[Tree] (config>system>management-interface>remote-management>manager connection-timeout)

Full Context

configure system management-interface remote-management manager connection-timeout

Description

This command configures the amount of time that this remote manager cannot be reached before it is considered to be down.

This command takes precedence over the same command configured in the global context (**config**>**system**>**management-interface**>**remote-management**).

The **no** form of this command reverts to the default.

Default

connection-timeout 60

Parameters

seconds

Specifies the connection timeout in seconds.

Values 1 to 3600

Platforms

7705 SAR Gen 2

7.107 connectivity-association

connectivity-association

Syntax

connectivity-association *ca-name* [**create**]

no connectivity-association *ca-name*

Context

[Tree] (config>macsec connectivity-association)

Full Context

configure macsec connectivity-association

Description

This command configures a connectivity association. MACsec connectivity associations are applied to a port dot1x configuration to enable MACsec on that port.

The **no** form of this command removes the connectivity association.

Parameters

ca-name

The name of the connectivity association, a string up to 32 characters long.

create

Mandatory while creating an entry.

Platforms

7705 SAR Gen 2

7.108 connectivity-verify

connectivity-verify

Syntax
connectivity-verify

Context
[Tree] (config>redundancy>mc>peer>mcr>l3ring>node connectivity-verify)

Full Context
configure redundancy multi-chassis peer mc-ring l3-ring ring-node connectivity-verify

Description
Commands in this context configure a node connectivity check.

Platforms
7705 SAR Gen 2

7.109 connector

connector

Syntax
connector

Context
[Tree] (config>port connector)

Full Context
configure port connector

Description
Commands in this context configure connector parameters.

Platforms
7705 SAR Gen 2

7.110 consider-system-ip-in-gep

consider-system-ip-in-gep

Syntax

[**no**] **consider-system-ip-in-gep**

Context

[Tree] (config>router>ldp consider-system-ip-in-gep)

Full Context

configure router ldp consider-system-ip-in-gep

Description

When this command is enabled, the system interprets the presence or absence of the system IP and its associated action in the applied Global Export Policies in the same way as for other interfaces' IP addresses. In that case:

- if the system IP is not present, its FEC will not be exported or it will be withdrawn if it has been exported
- if the system IP is present with "accept", its FEC will be exported
- if the system IP is present with "deny", its FEC will not be exported or it will be withdrawn if it had been exported

Enabling or disabling this command leads to the applied Global Export Policies being reevaluated.

The **no** form of this command causes the system to not interpret the presence or absence of the system IP in applied Global Export Policies, and the FEC for the system IP is exported (default behavior).

Default

no consider-system-ip-in-gep

Platforms

7705 SAR Gen 2

7.111 console

console

Syntax

console

Context

[Tree] (config>system>management-interface>cli>md-cli>environment console)

Full Context

configure system management-interface cli md-cli environment console

Description

Commands in this context configure console parameters.

Platforms

7705 SAR Gen 2

console

Syntax

console

Context

[Tree] (config>system>security>user-template console)

[Tree] (config>system>security>user console)

Full Context

configure system security user-template console

configure system security user console

Description

This command creates the context to configure user profile membership for the console (either Telnet or CPM serial port user).

Platforms

7705 SAR Gen 2

7.112 console-speed

console-speed

Syntax

console-speed *baud-rate*

no console-speed

Context

[Tree] (bof console-speed)

Full Context

bof console-speed

Description

This command configures the console port baud rate.

When this command is issued while editing the BOF file used for the most recent boot, both the BOF file and the active configuration are changed immediately.

The **no** form of this command reverts to the default value.

Default

console-speed 115200

Parameters

baud-rate

Specifies the console port baud rate, expressed as a decimal integer.

Values 9600, 19200, 38400, 57600, 115200

Platforms

7705 SAR Gen 2

7.113 contact

contact

Syntax

contact *contact-information*

no contact *contact-information*

Context

[Tree] (config>service>cust contact)

Full Context

configure service customer contact

Description

This command configures contact information for a customer.

Include any customer-related contact information such as a technician’s name or account contract name.

The **no** form of this command removes the contact information from the customer ID.

Default

no contact

Parameters

contact-information

Specifies customer contact information entered as an ASCII character string up to 80 characters in length. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Any printable, seven bit ASCII characters may be used within the string.

Platforms

7705 SAR Gen 2

contact

Syntax

contact *contact-name*

no contact

Context

[Tree] (config>system contact)

Full Context

configure system contact

Description

This command creates a text string that identifies the contact name for the device.

Only one contact can be configured, if multiple contacts are configured the last one entered will overwrite the previous entry.

The **no** form of the command reverts to default.

Default

no contact

Parameters

contact-name

Specifies the contact name character string. The string can be up to 80 characters long. Any printable, seven-bit ASCII characters can be used within the string. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms
7705 SAR Gen 2

7.114 context

context

Syntax
[**no**] **context**

Context
[Tree] (config>system>management-interface>cli>md-cli>environment>prompt context)

Full Context
configure system management-interface cli md-cli environment prompt context

Description
This command displays the current command context in the prompt.

The **no** form of this command suppresses the current command context in the prompt.

Default
context

Platforms
7705 SAR Gen 2

7.115 continuous

continuous

Syntax
[**no**] **continuous**

Context
[Tree] (config>saa>test continuous)

Full Context
configure saa test continuous

Description

This command specifies whether the SAA test is continuous. Once a test is configured as continuous, it cannot be started or stopped with the **oam saa** *test-name* {**start** | **stop**} command.

This option is not applicable to all SAA test types. Support is included for the following types:

- **cpe-ping**
- **dns**
- **eth-cfm-loopback**
- **eth-cfm-two-way-delay**
- **eth-cfm-two-way-slm**
- **icmp-ping** (not applicable to **rapid** type)
- **lsp-ping**
- **mac-ping**
- **sdp-ping**
- **vccv-ping**
- **vprn-ping**

The **no** form of this command disables the continuous execution of the test.

Platforms

7705 SAR Gen 2

7.116 control-channel-status

control-channel-status

Syntax

[**no**] **control-channel-status**

Context

[Tree] (config>service>epipe>spoke-sdp control-channel-status)

[Tree] (config>service>vpls>spoke-sdp control-channel-status)

Full Context

configure service epipe spoke-sdp control-channel-status

configure service vpls spoke-sdp control-channel-status

Description

This command enables the configuration of static pseudowire status signaling on a spoke SDP for which signaling for its SDP is set to OFF.

A control-channel-status **no shutdown** is allowed only if all of the following are true:

- SDP signaling is off.
- The control-word is enabled (the control-word is disabled by default)
- The service type is Epipe, Apipe, VPLS, Cpipe, or IES/VPRN
- Mate SDP signaling is off (in vc-switched services)
- The pw-path-id is configured for this spoke SDP.

The **no** form of this command removes control channel status signaling from a spoke SDP. It can only be removed if control channel status is shut down.

Default

no control-channel-status

Platforms

7705 SAR Gen 2

control-channel-status

Syntax

control-channel-status

Context

[Tree] (config>service>ies>if>spoke-sdp control-channel-status)

Full Context

configure service ies interface spoke-sdp control-channel-status

Description

This command enables the configuration of static pseudowire status signaling on a spoke-SDP for which signaling for its SDP is set to OFF.

A control-channel-status no shutdown is allowed only if all of the following are true:

- SDP signaling is off.
- The control-word is enabled (the control-word is disabled by default)
- The service type is Epipe, Apipe, VPLS, Cpipe, or IES/VPRN
- Mate SDP signaling is off (in vc-switched services)
- The pw-path-id is configured for this spoke-SDP.

The **no** form of this command removes control channel status signaling from a spoke-SDP. It can only be removed if control channel status is shut down.

Default

no control-channel-status

Platforms

7705 SAR Gen 2

control-channel-status

Syntax

control-channel-status

Context

[Tree] (config>service>vprn>if>spoke-sdp control-channel-status)

Full Context

configure service vprn interface spoke-sdp control-channel-status

Description

This command enables the configuration of static pseudowire status signaling on a spoke SDP for which signaling for its SDP is set to OFF.

A control-channel-status no shutdown is allowed only if all of the following are true:

- SDP signaling is off.
- The control-word is enabled (the control-word is disabled by default)
- The service type is Epipe, Apipe, VPLS, Cpipe, or IES/VPRN
- Mate SDP signaling is off (in vc-switched services)
- The pw-path-id is configured for this spoke SDP.

The **no** form of this command removes control channel status signaling from a spoke SDP. It can only be removed if control channel status is shut down.

Default

no control-channel-status

Platforms

7705 SAR Gen 2

7.117 control-word

control-word

Syntax

[**no**] **control-word**

Context

[Tree] (config>service>epipe>bgp-evpn>mpls control-word)

[Tree] (config>service>vpls>bgp-evpn>mpls control-word)

Full Context

configure service epipe bgp-evpn mpls control-word

configure service vpls bgp-evpn mpls control-word

Description

This command enables the transmission and reception of the **control-word**. As defined in RFC 7432, the use of the control-word helps avoid frame disordering.

It is enabled or disabled for all EVPN-MPLS destinations at the same time.

Default

no control-word

Platforms

7705 SAR Gen 2

control-word

Syntax

[**no**] **control-word**

Context

[Tree] (config>service>epipe>spoke-sdp control-word)

Full Context

configure service epipe spoke-sdp control-word

Description

The control word command provides the option to add a control word as part of the packet encapsulation for pseudowire types for which the control word is optional. These are Ethernet pseudowires (Epipe). ATM N:1 cell mode pseudowires (apipe vc-types atm-vcc and atm-vpc) and VT pseudowire (apipe vc-type atm-cell).

The configuration for the two directions of the pseudowire must match because the control word negotiation procedures described in Section 6.2 of RFC 4447 are not supported. The C-bit in the pseudowire FEC sent in the label mapping message is set to 1 when the control word is enabled. Otherwise, it is set to 0.

The service will only come up if the same C-bit value is signaled in both directions. If a spoke-sdp is configured to use the control word but the node receives a label mapping message with a C-bit clear, the node releases the label with the an "Illegal C-bit" status code as per Section 6.1 of RFC 4447. As soon as the user also enabled the control the remote peer, the remote peer will withdraw its original label and will

send a label mapping with the C-bit set to 1 and the VLL service will be up in both nodes. The control word must be enabled to allow MPLS-TP OAM to be used on a static spoke-sdp in a Apipe, Epipe and Cpipe service.

Platforms

7705 SAR Gen 2

control-word

Syntax

[**no**] **control-word**

Context

[Tree] (config>service>vpls>spoke-sdp control-word)

Full Context

configure service vpls spoke-sdp control-word

Description

The control word command provides the option to add a control word as part of the packet encapsulation for pseudowire types for which the control word is optional. These are Ethernet pseudowires (Epipe). ATM N:1 cell mode pseudowires (apipe vc-types atm-vcc and atm-vpc) and VT pseudowire (apipe vc-type atm-cell).

The configuration for the two directions of the pseudowire must match because the control word negotiation procedures described in Section 6.2 of RFC 4447 are not supported. The C-bit in the pseudowire FEC sent in the label mapping message is set to 1 when the control word is enabled. Otherwise, it is set to 0.

The service will only come up if the same C-bit value is signaled in both directions. If a spoke-sdp is configured to use the control word but the node receives a label mapping message with a C-bit clear, the node releases the label with the an "Illegal C-bit" status code as per Section 6.1 of RFC 4447. As soon as the user also enabled the control the remote peer, the remote peer will withdraw its original label and will send a label mapping with the C-bit set to 1 and the VLL service will be up in both nodes. The control word must be enabled to allow MPLS-TP OAM to be used on a static spoke-sdp in a Apipe, Epipe and Cpipe service.

Platforms

7705 SAR Gen 2

control-word

Syntax

[**no**] **control word**

Context

[Tree] (config>service>vpls>spoke-sdp control-word)

[Tree] (config>service>vpls>mesh-sdp control-word)

Full Context

configure service vpls spoke-sdp control-word

configure service vpls mesh-sdp control-word

Description

This command enables the use of the control word on pseudowire packets in VPLS and enables the use of the control word individually on each mesh SDP or spoke-SDP. By default, the control word is disabled. When the control word is enabled, all VPLS packets, including the BPDU frames, are encapsulated with the control word when sent over the pseudowire. The T-LDP control plane behavior is the same as in the implementation of control word for VLL services. The configuration for the two directions of the Ethernet pseudowire should match. The **no** form of this command reverts the mesh SDP or spoke-SDP to the default behavior of not using the control word. The control word must be enabled to use MPLS-TP OAM on a static spoke-sdp terminating in a VPLS.

Default

no control word

Platforms

7705 SAR Gen 2

7.118 controlword

controlword

Syntax

[**no**] **controlword**

Context

[Tree] (config>service>pw-template controlword)

Full Context

configure service pw-template controlword

Description

This command enables the use of the control word on pseudowire packets in VPLS and VPWS and enables the use of the control word individually on each mesh-sdp or spoke-sdp. By default, the control word is disabled. When the control word is enabled, all VPLS/VPWS packets, including the BPDU frames, are encapsulated with the control word when sent over the pseudowire. The T-LDP control plane behavior

is the same as in the implementation of control word for VLL services. The configuration for the two directions of the Ethernet pseudowire should match.

The **no** form of the command reverts the mesh SDP or spoke-sdp to the default behavior of not using the control word.

Default

no controlword

Platforms

7705 SAR Gen 2

7.119 convergence

convergence

Syntax

convergence

Context

[Tree] (config>service>vprn>bgp convergence)

Full Context

configure service vprn bgp convergence

Description

Commands in this context configure route convergence delay.

Platforms

7705 SAR Gen 2

convergence

Syntax

convergence

Context

[Tree] (config>router>bgp convergence)

Full Context

configure router bgp convergence

Description

Commands in this context configure route convergence delay.

Platforms

7705 SAR Gen 2

7.120 convert-file

convert-file

Syntax

convert-file *filename* **to** *output-file-name* **format** {**secure** | **legacy**} [**force**]

Context

[Tree] (admin>certificate convert-file)

Full Context

admin certificate convert-file

Description

This command converts imported certificates and keys in the cf3:/system-pki directory between secure and legacy format.

Parameters

filename

Specifies an existing filename, up to 95 characters.

output-file-name

Specifies the output file name, up to 95 characters. If the output filename already exists, and the **force** keyword is not selected, the system prompts to proceed or abort.

format

Specifies the target format.

Values secure — Specifies the enhanced secure format

legacy — Specifies the legacy format

force

Forces the conversion even if there is an existing file with the same output filename.

Platforms

7705 SAR Gen 2

7.121 coordinates

coordinates

Syntax

coordinates *coordinates*

no coordinates

Context

[Tree] (config>system coordinates)

Full Context

configure system coordinates

Description

This command creates a text string that identifies the system coordinates for the device location. For example, the command **coordinates** "*37.390 -122.0550*" is read as latitude 37.390 north and longitude 122.0550 west.

Only one set of coordinates can be configured. If multiple coordinates are configured, the last one entered overwrites the previous entry.

The **no** form of the command reverts to the default value.

Parameters

coordinates

Specifies the coordinates describing the device location character string. The string may be up to 80 characters long. Any printable, seven-bit ASCII characters can be used within the string. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. If the coordinates are subsequently used by an algorithm that locates the exact position of this node then the string must match the requirements of the algorithm.

Platforms

7705 SAR Gen 2

7.122 copy

copy

Syntax

copy

Context

[Tree] (config>filter copy)

Full Context

configure filter copy

Description

This command copies existing filter list entries for a specific filter ID to another filter ID. The **copy** command is a configuration level maintenance tool used to create new filters using existing filters. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword. If **overwrite** is not specified, an error will occur if the destination policy ID exists.

Platforms

7705 SAR Gen 2

copy

Syntax

copy *source-file-url dest-file-url* [**force**] [**no-redirect**] [**client-tls-profile** *profile*] [**proxy** *proxy-url*]

Context

[Tree] (file copy)

Full Context

file copy

Description

This command copies a file or all files in a directory from a source URL to a destination URL. At least one of the specified URLs should be a local URL. The optional wildcard (*) can be used to copy multiple files that share a common (partial) prefix and/or (partial) suffix.

When a file is copied to a destination with the same file name, the original file is overwritten by the new file specified in the operation. The following prompt appears if the destination file already exists:

"Overwrite destination file (y/n)?”

For example:

To copy a file named `srcfile` in a directory called `test` on cf2 in slot B to a file called `destfile` in a directory called production on cf1 in slot A, the syntax is:

```
sr1>file cf2:\ # copy cf2-B/test/srcfile cf1-A/production/destfile
```

To FTP a file named `121201.cfg` in directory `mydir` stored on cf1 in slot A to a network FTP server with IP address 192.0.2.79 in a directory called `backup` with a destination file name of `121201.cfg`, the FTP syntax is:

```
copy cf1-A/mydir/121201.cfg 192.0.2.79/backup/121201.cfg
```

Parameters

source-file-url

Specifies the location of the source file or directory to be copied.

Values

local-url	[*cflash-id*/][*file-path*] up to 200 characters, including cflash-id directory length 99 chars max each
remote-url	[{ftp:// | tftp:// | http:// | https://}*login*:*pswd*@*remote-locn*/][*file-path*]
	up to 247 characters
	directory length up to 199 characters
remote-locn	[*hostname* | *ipv4-address* | [*ipv6-address*]]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D
	interface - up to 32 characters, for link local addresses 255
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

dest-file-url

Specifies the destination of the copied file or directory.

force

Specifies to force an immediate copy of the specified file(s). Executes the command without displaying a user prompt message. This command also automatically accepts HTTP redirects unless overridden by the **no-redirect** parameter.

profile

Specifies the TLS client profile configured under **config**>**system**>**security**>**tls**>**client-tls-profile** to use.

proxy-url

Specifies the URL of an HTTP proxy. For example, http://proxy.mydomain.com:8000. This URL must be an HTTP URL and not an HTTPS URL.

no-redirect

Specifies to automatically refuse any HTTP redirects without prompting the user.

Platforms

7705 SAR Gen 2

copy

Syntax

copy [*line*]

Context

[Tree] (candidate copy)

Full Context

candidate copy

Description

This command copies the selected CLI node (which includes all sub-branches) into a temporary buffer that can be used for a subsequent insert. The contents of the temporary buffer are deleted when the operator exits the candidate edit mode.

Parameters

line

Specifies which line to copy.

Values line, offset, **first**, **edit-point**, **last**

line — absolute line number

offset — relative line number to the current edit point. Prefixed with '+' or '-'.

first — keyword to indicate the first line

edit-point — keyword to indicate the current edit point

last — keyword to indicate the last line that is not 'exit'

Platforms

7705 SAR Gen 2

copy

Syntax

copy {**user** *source-user* | *profile source-profile*} **to** *destination* [**overwrite**]

Context

[Tree] (config>system>security copy)

Full Context

configure system security copy

Description

This command copies a profile or user from a source profile to a destination profile.

Parameters

source-profile

Specifies an existing profile to copy.

dest-profile

Specifies the copied profile is copied to the destination profile.

overwrite

Specifies that the destination profile configuration is overwritten with the copied source profile configuration. A profile is not overwritten if the **overwrite** command is not specified.

Platforms

7705 SAR Gen 2

7.123 copy-config

copy-config

Syntax

[**no**] **copy-config**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization copy-config)

Full Context

configure system security profile netconf base-op-authorization copy-config

Description

This command enables the NETCONF <copy-config> RPC.

The **no** form of this command disables the RPC.

Default

no copy-config

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

7.124 copy-traffic-class-upon-decapsulation

copy-traffic-class-upon-decapsulation

Syntax

[**no**] **copy-traffic-class-upon-decapsulation**

Context

[Tree] (config>ipsec>tnl-temp copy-traffic-class-upon-decapsulation)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel copy-traffic-class-upon-decapsulation)

[Tree] (config>router>if>ipsec>ipsec-tunnel copy-traffic-class-upon-decapsulation)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel copy-traffic-class-upon-decapsulation)

[Tree] (config>service>ies>interface>ipsec>ipsec-tunnel copy-traffic-class-upon-decapsulation)

Full Context

configure ipsec tunnel-template copy-traffic-class-upon-decapsulation

configure service vprn interface ipsec ipsec-tunnel copy-traffic-class-upon-decapsulation

configure router interface ipsec ipsec-tunnel copy-traffic-class-upon-decapsulation

configure service vprn interface sap ipsec-tunnel copy-traffic-class-upon-decapsulation

configure service ies interface ipsec ipsec-tunnel copy-traffic-class-upon-decapsulation

Description

This command copies the traffic class from the outer tunnel IP packet header to the payload IP packet header upon tunnel decapsulation (public to private direction).

The **no** form of this command disables the traffic copying.

Default
copy-traffic-class-upon-decapsulation

Platforms
7705 SAR Gen 2

7.125 core-connectivity

core-connectivity

Syntax
[**no**] **core-connectivity**

Context
[Tree] (debug>service>id>stp core-connectivity)

Full Context
debug service id stp core-connectivity

Description
This command enables STP debugging for core connectivity.

The **no** form of the command disables debugging.

Platforms
7705 SAR Gen 2

7.126 count

count

Syntax
count *number*

no count

Context
[Tree] (config>system>cron>sched count)

Full Context

configure system cron schedule count

Description

This command configures the total number of times a CRON "interval" schedule is run. For example, if the interval is set to 600 and the count is set to 4, the schedule runs 4 times at 600 second intervals.

Default

no count

Parameters

number

Specifies the number of times the schedule is run.

Values 1 to 65535

Default 65535

Platforms

7705 SAR Gen 2

7.127 cpe-check

cpe-check

Syntax

cpe-check *cpe-ip-address*

no cpe-check [*cpe-ip-address*]

Context

[Tree] (config>service>vprn>static-route-entry>indirect cpe-check)

[Tree] (config>service>vprn>static-route-entry>next-hop cpe-check)

Full Context

configure service vprn static-route-entry indirect cpe-check

configure service vprn static-route-entry next-hop cpe-check

Description

This command enables CPE-check and specifies the IP address of the target CPE device.

This option initiates a background ICMP ping test to the configured target IP address. The IP address can either be an IPv4 address for IPv4 static routes or an IPv6 address for IPv6 static routes. The target-ip-

address cannot be in the same subnet as the static route subnet itself to avoid possible circular references. This option is mutually exclusive with BFD support on a given static route.

Note:
A node that is sourcing CPE-check packets waits an additional full interval before taking action, which gives the CPE time to respond. For example, with a **drop-count** of 3 and an interval of 1s, three CPE-check packets are sent out and the node waits for the duration of another interval before acting on the loss. Failure declaration may take extra time depending on the load, interval, and other factors. In line with multitasking, multi-priority operating principles of the node, and the relative priority of **cpe-ping**, the node paces these minor events.

The **no** form of this command disables the **cpe-check** option.

Default

no cpe-check

Parameters

cpe-ip-address

Specifies the IP address of the CPE device.

Platforms

7705 SAR Gen 2

cpe-check

Syntax

cpe-check *cpe-ip-address*

no cpe-check [*cpe-ip-address*]

Context

[Tree] (config>router>static-route-entry>next-hop cpe-check)

[Tree] (config>router>static-route-entry>indirect cpe-check)

Full Context

configure router static-route-entry next-hop cpe-check

configure router static-route-entry indirect cpe-check

Description

This command enables CPE-check and specifies the IP address of the target CPE device.

This option initiates a background ICMP ping test to the configured target IP address. The IP address can either be an IPv4 address for IPv4 static routes or an IPv6 address for IPv6 static routes. The target-ip-address cannot be in the same subnet as the static route subnet itself to avoid possible circular references. This option is mutually exclusive with BFD support on a given static route.

Note:
A node that is sourcing CPE-check packets waits an additional full interval before taking action, which gives the CPE time to respond. For example, with a **drop-count** of 3 and an interval of 1s, three CPE-check packets are sent out and the node waits for the duration of another interval before acting on the loss. Failure declaration may take extra time depending on the load, interval, and other factors. In line with multitasking, multi-priority operating principles of the node, and the relative priority of **cpe-ping**, the node paces these minor events.

The **no** form of this command disables the **cpe-check** option.

Default

no cpe-check

Parameters

cpe-ip-address

Specifies the IP address of the CPE device.

Platforms

7705 SAR Gen 2

7.128 cpr-window-size

cpr-window-size

Syntax

cpr-window-size *window-size*

Context

[Tree] (config>port>dwdm>coherent cpr-window-size)

Full Context

configure port dwdm coherent cpr-window-size

Description

This command configures the window size used for carrier phase recovery.

Default

32

Parameters

window-size

Indicates the number of symbols used for carrier phase recovery algorithm of the receiver. When this parameter is changed, the link bounces because the receiver needs to be reconfigured.

Values 2, 4, 8, 16, 32, 64

Platforms

7705 SAR Gen 2

7.129 crc-monitor

crc-monitor

Syntax

crc-monitor

Context

[Tree] (config>port>ethernet crc-monitor)

Full Context

configure port ethernet crc-monitor

Description

This command configures Ethernet CRC Monitoring parameters.

Platforms

7705 SAR Gen 2

7.130 create

create

Syntax

[**no**] **create**

Context

[Tree] (environment create)

Full Context

environment create

Description

By default, the **create** command is required to create a new OS entity.

The **no** form of the command disables requiring the **create** keyword.

Default

create

Platforms

7705 SAR Gen 2

7.131 create-mpls-tunnel

create-mpls-tunnel

Syntax

[**no**] **create-mpls-tunnel**

Context

[Tree] (config>router>policy-options>policy-statement>entry>action create-mpls-tunnel)

Full Context

configure router policy-options policy-statement entry action create-mpls-tunnel

Description

This command enables the creation of an MPLS tunnel to the BGP next-hop. It is supported for the following address families:

- vpn-ipv4
- vpn-ipv6
- evpn
- label-ipv4
- label-ipv6
- ipv4
- ipv6

The **no** form of the command disables the creation of an MPLS tunnel.

Default

no create-mpls-tunnel

Platforms

7705 SAR Gen 2

7.132 create-subscription

create-subscription

Syntax

[**no**] **create-subscription**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization create-subscription)

Full Context

configure system security profile netconf base-op-authorization create-subscription

Description

This command enables the NETCONF <create-subscription> RPC in the default user profile.

The **base-op-authorization create-subscription** configuration is not pre-emptive, which means that it is checked only at the time of the initial subscription. Configuration changes to the **base-op-authorization** do not cancel any in-progress subscriptions and operators who successfully subscribed continue to receive messages.

The **no** form of this command disables the RPC.

Default

no create-subscription

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

7.133 create-udp-tunnel

create-udp-tunnel

Syntax

create-udp-tunnel

no create-udp-tunnel

Context

[Tree] (config>router>policy-options>policy-statement>entry>action create-udp-tunnel)

[Tree] (config>router>policy-options>policy-statement>default-action create-udp-tunnel)

Full Context

configure router policy-options policy-statement entry action create-udp-tunnel

configure router policy-options policy-statement default-action create-udp-tunnel

Description

This command instructs the router to create an MPLS-over-UDP tunnel upon receiving BGP routes that match the import policy.

Default

no create-udp-tunnel

Platforms

7705 SAR Gen 2

7.134 credential

credential

Syntax

credential

Context

[Tree] (config>ipsec>client-db>client credential)

Full Context

configure ipsec client-db client credential

Description

Commands in this context configure the parameters used to authenticate peers.

Platforms

7705 SAR Gen 2

7.135 credits

credits

Syntax

credits [**lowercase** *credits*] [**uppercase** *credits*] [**numeric** *credits*] [**special-character** *credits*]

no credits

Context

[Tree] (config>system>security>password>complexity-rules credits)

Full Context

configure system security password complexity-rules credits

Description

The maximum credits given for usage of the different character classes in the local passwords.

The **no** form of this command resets to default.

Default

no credits

Parameters

credits

Specifies the number of credits that can be used for each characters class.

Values 0 to 10

Platforms

7705 SAR Gen 2

7.136 crl-expiration-warning

crl-expiration-warning

Syntax

crl-expiration-warning *hours* [**repeat** *repeat-hours*]

no crl-expiration-warning

Context

[Tree] (config>system>security>pki crl-expiration-warning)

Full Context

configure system security pki crl-expiration-warning

Description

This command specifies when the systems issues a **BeforeExp** message before a CRL expires. For example, with **certificate-expiration-warning 5**, the system issues a **BeforeExp** message 5 hours before a CRL expires. An optional **repeat** *repeat-hour* parameter enables the system to repeat the **BeforeExp** message every hour until the CRL expires.

If the user only wants **AfterExp**, then **certificate-expiration-warning 0** can be used to achieve this.

BeforeExp and **AfterExp** warnings can be cleared in following cases:

- The CRL is reloaded by the **admin certificate reload** command. In this case, if the reloaded file is not expired, then **AfterExp** is cleared. And, if the reloaded file is outside of configured warning window, then the **BeforeExp** is also cleared.
- When the **ca-profile** is shutdown, then **BeforeExp** and **AfterExp** of corresponding certificates are cleared.
- When **no crl-expiration-warning** command is configured, then all existing **BeforeExp** and **AfterExp** are cleared.
- Users may change the configuration of the **crl-expiration-warning** so that certain CRL are no longer in the warning window. **BeforeExp** of corresponding CRL are cleared.
- If the system time changes so that the new time causes the CRL to no longer be in the warning window, then **BeforeExp** is cleared. If the new time causes an expired CRL to come non-expired, then **AfterExp** is cleared.

Default

no crl-expiration-warning

Parameters

hours

Specifies the amount of time before a CRL expires when system issues **BeforeExp**

Values 0 to 8760

repeat-hour

Specifies that the system repeats **BeforeExp** every repeat-hour

Values 0 to 8760

Platforms

7705 SAR Gen 2

7.137 crl-file

crl-file

Syntax

crl-file *filename*

no crl-file

Context

[Tree] (config>system>security>pki>ca-profile crl-file)

Full Context

configure system security pki ca-profile crl-file

Description

This command specifies the name of a file in cf3:\system-pki\crl as the Certification Revoke List file of the **ca-profile**.

Notes:

- The system performs following checks against configured crl-file when a **no shutdown** command is issued:
  - A valid cert-file of the ca-profile must be already configured.
  - Configured crl-file must be a DER formatted CRLv2 file.
  - All non-optional fields defined in section 5.1 of RFC 5280 must exist and conform to the RFC 5280 defined format.
  - Check the version field to see if its value is 0x1.
  - Delta CRL Indicator must not exist (delta CRL is not supported).
  - CRL's signature must be verified by using the cert-file of ca-profile.

  If any of above checks fail, the **no shutdown** command fails.
- Changing or removing the **crl-file** is only allowed when the **ca-profile** is in a **shutdown** state.

The **no** form of this command removes the filename from the configuration.

Parameters

filename

Specifies the name of CRL file stored in cf3:\system-pki\crl.

Platforms

7705 SAR Gen 2

7.138 crl-update

crl-update

Syntax

crl-update ca *ca-profile-name*

Context

[Tree] (admin>certificate crl-update)

Full Context

admin certificate crl-update

Description

This command manually triggers the Certificate Revocation List file (CRL) update for the specified ca-profile.

Using this command requires shutting down the **auto-crl-update**.

Parameters

ca-profile-name

Specifies the name of the Certificate Authority profile.

Platforms

7705 SAR Gen 2

7.139 crl-urls

crl-urls

Syntax

crl-urls

Context

[Tree] (config>system>security>pki>ca-prof>auto-crl-update crl-urls)

Full Context

configure system security pki ca-profile auto-crl-update crl-urls

Description

Commands in this context configure **crl-urls** parameters. The system allows up to eight URL entries to be configured and tries each URL in order and stop when a qualified CRL is successfully downloaded. A qualified CRL is a valid CRL signed by the CA and is more recent than the existing CRL.

If none of the configured URLs returns a qualified CRL, then:

- If the schedule-type is next-update-based, system will wait for configure retry-interval before it start from beginning of the list again.
- If the schedule-type is periodic, then system will wait till next periodic update time.

If the user wants to manually stop the download, shutting down of auto-crl-retrieval could be used to achieve this.

Platforms

7705 SAR Gen 2

7.140 cron

cron

Syntax

cron

Context

[Tree] (config>system cron)

Full Context

configure system cron

Description

This command creates the context to create scripts, script parameters and schedules which support the Service Assurance Agent (SAA) functions.

CRON features are saved to the configuration file on both primary and backup control modules. If a control module switchover occurs, CRON events are restored when the new configuration is loaded. If a control module switchover occurs during the execution of a cron script, the failover behavior will be determined by the contents of the script.

Platforms

7705 SAR Gen 2

cron

Syntax

cron

Context

[Tree] (config>system>security>cli-script>authorization cron)

Full Context

configure system security cli-script authorization cron

Description

Commands in this context configure authorization for the Cron job-scheduler.

Platforms

7705 SAR Gen 2

7.141 csnp-authentication

csnp-authentication

Syntax

[**no**] **csnp-authentication**

Context

[Tree] (config>service>vprn>isis csnp-authentication)

[Tree] (config>service>vprn>isis>level csnp-authentication)

Full Context

configure service vprn isis csnp-authentication

configure service vprn isis level csnp-authentication

Description

This command enables authentication of individual ISIS packets of complete sequence number PDUs (CSNP) type for the VPRN instance.

Platforms

7705 SAR Gen 2

csnp-authentication

Syntax

[**no**] **csnp-authentication**

Context

[Tree] (config>router>isis>level csnp-authentication)

[Tree] (config>router>isis csnp-authentication)

Full Context

configure router isis level csnp-authentication

configure router isis csnp-authentication

Description

This command enables authentication of individual IS-IS packets of complete sequence number PDUs (CSNP) type.

The **no** form of this command suppresses authentication of CSNP packets.

Default

csnp-authentication

Platforms

7705 SAR Gen 2

7.142 csnp-interval

csnp-interval

Syntax

csnp-interval *seconds*

no csnp-interval

Context

[Tree] (config>service>vprn>isis>if csnp-interval)

Full Context

configure service vprn isis interface csnp-interval

Description

This command configures the time interval, in seconds, to send complete sequence number (CSN) PDUs from the interface. IS-IS must send CSN PDUs periodically.

The **no** form of this command reverts to the default value.

Default

csnp-interval 10 — CSN PDUs are sent every 10 seconds for LAN interfaces.

csnp-interval 5 — CSN PDUs are sent every 5 seconds for point-to-point interfaces.

Parameters

seconds

The time interval, in seconds between successive CSN PDUs sent from this interface expressed as a decimal integer.

Values 1 to 65535

Platforms

7705 SAR Gen 2

csnp-interval

Syntax

csnp-interval *seconds*

no csnp-interval

Context

[Tree] (config>router>isis>interface csnp-interval)

Full Context

configure router isis interface csnp-interval

Description

This command configures the time interval, in seconds, to send complete sequence number (CSN) PDUs from the interface. IS-IS must send CSN PDUs periodically.

The **no** form of this command reverts to the default value.

Default

csnp-interval 10 — CSN PDUs are sent every 10 seconds for LAN interfaces.

csnp-interval 5 — CSN PDUs are sent every 5 seconds for point-to-point interfaces.

Parameters

seconds

Specifies the time interval, in seconds, between successive CSN PDUs sent from this interface expressed as a decimal integer.

Values 1 to 65535

Platforms

7705 SAR Gen 2

7.143 csnp-on-p2p

csnp-on-p2p

Syntax

[**no**] **csnp-on-p2p**

Context

[Tree] (config>router>isis csnp-on-p2p)

[Tree] (config>service>vprn>isis csnp-on-p2p)

Full Context

configure router isis csnp-on-p2p

configure service vprn isis csnp-on-p2p

Description

This command enables the periodic transmission of CSNP PDUs to point-to-point adjacent systems.

The **no** form of this command disables the periodic transmission of CSNP PDUs to point-to-point adjacent systems.

Default

csnp-on-p2p

Platforms

7705 SAR Gen 2

7.144 cspf

cspf

Syntax

[**no**] **cspf**

Context

[Tree] (debug>router>isis cspf)

Full Context

debug router isis cspf

Description

This command enables debugging for IS-IS cspf.

The **no** form of the command disables debugging.

Platforms

7705 SAR Gen 2

cspf

Syntax

cspf [*ip-address*]

no cspf

Context

[Tree] (debug>router>ospf cspf)

Full Context

debug router ospf cspf

Description

This command enables debugging for an OSPF constraint-based shortest path first (CSPF).

Parameters

ip-address

Specifies the IP address for the range used for CSPF.

Platforms

7705 SAR Gen 2

7.145 cspf-on-loose-hop

cspf-on-loose-hop

Syntax

[**no**] **cspf-on-loose-hop**

Context

[Tree] (config>router>mpls cspf-on-loose-hop)

Full Context

configure router mpls cspf-on-loose-hop

Description

This command enables the option to do CSPF calculations until the next loose hop or the final destination of LSP on LSR. On receiving a PATH message on LSR and processing of all local hops in the received ERO, if the next hop is loose, then the LSR node will first do a CSPF calculation until the next loose hop. On successful completion of CSPF calculation, ERO in PATH message is modified to include newly calculated intermediate hops and propagate it forward to the next hop. This allows setting up inter-area LSPs based on ERO expansion method.

Note:
The LSP may fail to set up if this option is enabled on an LSR that is not an area border router and receives a PATH message without proper next loose hop in ERO. The 'cspf-on-loose-hop' configuration is allowed to change dynamically and applied to new LSP setup after change.

Default

no cspf-on-loose-hop

Platforms

7705 SAR Gen 2

7.146 cspf-te

cspf-te

Syntax

cspf-te [**detail**]

no cspf-te

Context

[Tree] (debug>router>pcep>pcc cspf-te)

[Tree] (debug>router>pcep>pcc>conn cspf-te)

Full Context

debug router pcep pcc cspf-te

debug router pcep pcc connection cspf-te

Description

This command debugs Constrained Shortest Path First-Traffic Engineering (CSPF-TE) events.

The **no** form of this command disables debugging.

Parameters

detail

Keyword used to specify detailed information about all events.

Platforms

7705 SAR Gen 2

7.147 cumulative-factor

cumulative-factor

Syntax

[**no**] **cumulative-factor** *cumulative-factor*

Context

[Tree] (config>service>vpls>mac-move>secondary-ports cumulative-factor)

[Tree] (config>service>vpls>mac-move>primary-ports cumulative-factor)

[Tree] (config>service>template>vpls-template>mac-move>secondary-ports cumulative-factor)

[Tree] (config>service>template>vpls-template>mac-move>primary-ports cumulative-factor)

Full Context

configure service vpls mac-move secondary-ports cumulative-factor

configure service vpls mac-move primary-ports cumulative-factor

configure service template vpls-template mac-move secondary-ports cumulative-factor

configure service template vpls-template mac-move primary-ports cumulative-factor

Description

This command defines a factor defining how many mac-relearn measurement periods can be used to measure mac-relearn rate. The rate must be exceeded during the defined number of consecutive periods before the corresponding port is blocked by the mac-move feature. The cumulative-factor of primary ports must be higher than cumulative-factor of secondary ports.

Default

cumulative-factor 2 — secondary ports

cumulative-factor 3 — primary ports

Parameters

factor

Specifies the factor defining the number of mac-relearn measurement periods can be used to measure mac-relearn rate

Values 2 to 10

Platforms

7705 SAR Gen 2

7.148 current-hop-limit

current-hop-limit

Syntax

current-hop-limit *limit*

no current-hop-limit

Context

[Tree] (config>service>vprn>router-advert>if current-hop-limit)

Full Context

configure service vprn router-advertisement interface current-hop-limit

Description

This command configures the hop limit to be advertised.

The **no** form of this command returns the command to the default setting.

Default

current-hop-limit 64

Parameters

limit

Specifies the default value to be placed in the current hop limit field in router advertisement policies sent.

Values 0 to 255

Platforms

7705 SAR Gen 2

current-hop-limit

Syntax

current-hop-limit *number*

no current-hop-limit

Context

[Tree] (config>router>router-advert>if current-hop-limit)

Full Context

configure router router-advertisement interface current-hop-limit

Description

This command configures the current-hop-limit in the router advertisement messages. It informs the nodes on the subnet about the hop-limit when originating IPv6 packets.

Default

current-hop-limit 64

Parameters

number

Specifies the hop limit.

Values 0 to 255. A value of zero means there is an unspecified number of hops.

Platforms

7705 SAR Gen 2

7.149 custom-option

custom-option

Syntax

custom-option *option-number* **address** [*ip-address*]

custom-option *option-number* **address** *ipv6-address* [*ipv6-address*]

custom-option *option-number* **domain** [*domain-string*]

custom-option *option-number* **hex** *hex-string*

custom-option *option-number* **string** *ascii-string*

no custom-option *option-number*

Context

[Tree] (config>router>dhcp>server>pool>subnet>options custom-option)

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options custom-option)

[Tree] (config>service>vprn>dhcp>server>pool>options custom-option)

[Tree] (config>router>dhcp>server>pool>options custom-option)

Full Context

configure router dhcp local-dhcp-server pool subnet options custom-option

configure subscriber-mgmt local-user-db ipoe host options custom-option

configure service vprn dhcp local-dhcp-server pool options custom-option

configure router dhcp local-dhcp-server pool options custom-option

Description

This command configures specific DHCP options. The options defined here can overrule options in the local user database.

The **no** form of the removes the custom option parameters from the configuration.

Parameters

option-number

Specifies up to four option numbers that the DHCP server uses to send the identification strings to the DHCP client.

Values 1 to 254

ip-address

Specifies the IP address of a host.

Values a.b.c.d

ipv6-address

Specifies the IPv6 address of a host. Applicable to DHCP6 only.

Values

ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

domain-string

Specifies the domain name, up to 127 characters.

hex-string

Specifies the hex value of this option.

Values 0x0 to 0xFFFFFFFF (up to 254 hex nibbles)

ascii-string

Specifies the value of this option, up to 127 characters.

Platforms

7705 SAR Gen 2

custom-option

Syntax

custom-option *option-number* **address** [*ipv6-address*]

custom-option *option-number* **domain** [*domain-string*]

custom-option *option-number* **hex** *hex-string*

custom-option *option-number* **string** *ascii-string*

no custom-option *option-number*

Context

[Tree] (config>router>dhcp6>server>pool>options custom-option)

[Tree] (config>service>vprn>dhcp6>server>pool>options custom-option)

[Tree] (config>router>dhcp6>server>pool>prefix>options custom-option)

[Tree] (config>service>vprn>dhcp6>server>pool>prefix>options custom-option)

Full Context

configure router dhcp6 local-dhcp-server pool options custom-option

configure service vprn dhcp6 local-dhcp-server pool options custom-option

configure router dhcp6 local-dhcp-server pool prefix options custom-option

configure service vprn dhcp6 local-dhcp-server pool prefix options custom-option

Description

This command configures specific DHCP6 options. The options defined here can overrule options in the local user database.

The **no** form of the removes the custom option parameters from the configuration.

Parameters

option-number

Specifies up to four option numbers that the DHCP6 server uses to send the identification strings to the DHCP6 client.

Values 1 to 254

ipv6-address

Specifies the IPv6 address of a host.

Values	:ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)	
		x:x:x:x:x:x:d.d.d.d	
		x:	[0 to FFFF]H
		d:	[0 to 255]D

domain-string

Specifies the domain name, up to 127 characters.

hex-string

Specifies the hex value of this option.

Values 0x0 to 0xFFFFFFFF (up to 254 hex nibbles)

ascii-string

Specifies the value of this option, up to 127 characters.

Platforms

7705 SAR Gen 2

7.150 custom-record

custom-record

Syntax

[**no**] **custom-record**

Context

[Tree] (config>log>acct-policy custom-record)

Full Context

configure log accounting-policy custom-record

Description

Commands in this context configure the layout and setting for a custom accounting record associated with this accounting policy.

The **no** form of this command reverts the configured values to the defaults.

Platforms

7705 SAR Gen 2

7.151 customer

customer

Syntax

customer *customer-id* [**create**] [**name** *name*]

no customer *customer-id*

Context

[Tree] (config>service customer)

Full Context

configure service customer

Description

This command creates a customer ID and customer context used to associate information with a particular customer. Services can later be associated with this customer at the service level.

Each *customer-id* must be unique. The **create** keyword must follow each new **customer** *customer-id* entry.

Enter an existing **customer** *customer-id* (without the *create* keyword) to edit the customer's parameters.

An optional customer **name** can be specified and is tied to the **customer-name** in the customer context (setting either **customer-name** or **name** will cause the other to change as well).

The **no** form of this command removes a *customer-id* and all associated information. Before removing a *customer-id*, all references to that customer in all services must be deleted or changed to a different customer ID.

Default

customer 1 always exists on the system and cannot be deleted.

Parameters

customer-id

Specifies the ID number to be associated with the customer, expressed as an integer.

Values *customer-id*: 1 to 2147483647

customer-name: 64 characters maximum

create

This keyword is required when first creating the configuration context. Once the context is created, it is possible to navigate into the context without the **create** keyword.

name name

This parameter configures an optional customer name, up to 64 characters in length, which adds a name identifier to a given customer to then use that customer name in configuration references as well as display and use customer names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the SR OS platforms.

All services are required to assign a customer ID to initially create a customer. However, either the customer ID or the customer name can be used to identify and reference a given customer once it is initially created.

If a name is not specified at creation time, then SR OS assigns a string version of the *customer-id* as the name.

Values *name*: 64 characters maximum

Platforms

7705 SAR Gen 2

7.152 customer-id-range

customer-id-range

Syntax

customer-id-range start *customer-id* **end** *customer-id*

no customer-id-range

Context

[Tree] (config>service>md-auto-id customer-id-range)

Full Context

configure service md-auto-id customer-id-range

Description

This command specifies the range of IDs used by SR OS to automatically assign an ID to customers that are created in model-driven interfaces without an ID explicitly specified by the user or client.

A customer created with an explicitly-specified ID cannot use an ID in this range. In the classic CLI and SNMP, the ID range cannot be changed while objects exist inside the previous or new range. In MD interfaces, the range can be changed, which causes any previously existing objects in the previous ID range to be deleted and re-created using a new ID in the new range.

The **no** form of this command removes the range values.

See the **config>service md-auto-id** command for further details.

Default

no customer-id-range

Parameters

start *customer-id*

Specifies the lower value of the ID range. The value must be less than or equal to the **end** value.

Values 2 to 2147483647

end *customer-id*

Specifies the upper value of the ID range. The value must be greater than or equal to the **start** value.

Values 2 to 2147483647

Platforms

7705 SAR Gen 2

8 d Commands

8.1 d-path-ignore

d-path-ignore

Syntax

[**no**] **d-path-ignore**

Context

[Tree] (config>service>system>bgp-evpn>ad-per-evi-routes d-path-ignore)

Full Context

configure service system bgp-evpn ad-per-evi-routes d-path-ignore

Description

This command makes the router ignore the Domain PATH attribute (D-PATH) when BGP computes the best path selection for received routes.

The **no** form of this command considers the D-PATH length and value as a tiebreaker in determining the best-path selection. In accordance with *draft-sr-bess-evpn-dpath*, the router compares the D-PATH attribute received in AD per-EVI routes with the same key as follows:

- The routes with the shortest D-PATH are preferred; therefore, routes not tied for the shortest D-PATH are removed. Routes without D-PATH are considered zero-length D-PATH.
- The routes with the numerically lowest left-most Domain-ID are preferred; therefore, routes not tied for the numerically lowest left-most Domain-ID are removed from consideration.

Default

no d-path-ignore

Platforms

7705 SAR Gen 2

8.2 d-path-length-ignore

d-path-length-ignore

Syntax

[**no**] **d-path-length-ignore**

Context

[Tree] (config>service>vprn d-path-length-ignore)

[Tree] (config>router>bgp>path-selection d-path-length-ignore)

Full Context

configure service vprn d-path-length-ignore

configure router bgp best-path-selection d-path-length-ignore

Description

This command enables and disables the ability of the router to ignore D-PATH domain segment length during best-path selection. At the base router level (or **vprn>bgp** level for PE-CE routers), this command allows BGP to ignore the D-PATH domain segment length for best-path selection purposes. The D-PATH length is ignored when comparing two VPN routes or two IFL routes within the same RD. However, these VPN/IFL routes are processed in Main-BGP instance.

At the VPRN router level, this command allows the VPRN RTM to ignore the D-PATH domain segment length for best path selection purposes (for routes in VPRN). The user can control whether the D-PATH length is considered when two VPN routes with different RDs are compared.

Best-path selection for EVPN-IFF routes against other owners (for example, EVPN-IFL or IPVPN) still relies on RTM preference. When EVPN-IFF RTM preference matches the RTM preference of another BGP owner, the existing RTM selection applies and D-PATH is not considered, irrespective of the **d-path-length-ignore** configuration.

The **no** form of this command disables the ability to ignore the D-PATH domain segment length.

Default

no d-path-length-ignore

Platforms

7705 SAR Gen 2

8.3 dad-disable

dad-disable

Syntax

[**no**] **dad-disable**

Context

[Tree] (config>service>ies>if>ipv6 dad-disable)

[Tree] (config>service>vprn>if>ipv6 dad-disable)

[Tree] (config>router>if>ipv6 dad-disable)

Full Context

configure service ies interface ipv6 dad-disable

configure service vprn interface ipv6 dad-disable

configure router interface ipv6 dad-disable

Description

This command disables duplicate address detection (DAD) on the interface. When **dad-disable** is configured on the interface, the router does not perform a DAD check and all IPv6 addresses on the interface immediately enter a preferred state without checking for uniqueness on the interface. This command is useful when an interface enters a looped state during troubleshooting and becomes operationally disabled when the loop is detected; a manual intervention is required to clear the DAD violation.

The **no** form of this command enables duplicate address detection (DAD) on the interface.

Default

no dad-disable

Platforms

7705 SAR Gen 2

8.4 damp-peer-oscillations

damp-peer-oscillations

Syntax

damp-peer-oscillations [**idle-hold-time** *initial-wait second-wait max-wait*] [**error-interval** *error-interval*]

Context

[Tree] (config>service>vprn>bgp damp-peer-oscillations)

[Tree] (config>service>vprn>bgp>group damp-peer-oscillations)

[Tree] (config>service>vprn>bgp>group>neighbor damp-peer-oscillations)

Full Context

configure service vprn bgp damp-peer-oscillations

configure service vprn bgp group damp-peer-oscillations

configure service vprn bgp group neighbor damp-peer-oscillations

Description

This command controls how long a BGP peer session remains in the idle-state after some type of error causes the session to reset. In the idle state, BGP does not initiate or respond to attempts to establish a new session. Repeated errors that occur a short while after each session reset cause longer and longer hold times in the idle state. This command supports the DampPeerOscillations FSM behavior described in section 8.1 of RFC 4271, *A Border Gateway Protocol 4 (BGP-4).*

The default behavior, which applies when no damp-peer-oscillations is configured, is to immediately transition out of the idle-state after every reset.

Default

no damp-peer-oscillations

Parameters

initial-wait

Specifies the amount of time, in minutes, that a session remains in the idle-state after it has been stable for a while.

Values 0 to 2048

Default 0

second-wait

Specifies the period of time, in minutes, that is doubled after each repeated session failure that occurs within a relatively short span of time.

Values 0 to 2048

Default 5

max-wait

Specifies the maximum amount of time, in minutes, that a session remains in the idle-state after it has experienced repeated instability.

Values 0 to 2048

Default 60

error-interval

Specifies the interval of time, in minutes after a session reset, during which the session must be error-free in to reset the penalty counter and return to idle-hold-time to initial-wait.

Values 0 to 2048

Default 30

Platforms

7705 SAR Gen 2

damp-peer-oscillations

Syntax

damp-peer-oscillations [**idle-hold-time** *initial-wait second-wait max-wait*] [**error-interval** *error-interval*]

Context

[Tree] (config>router>bgp>group>neighbor damp-peer-oscillations)

[Tree] (config>router>bgp damp-peer-oscillations)

[Tree] (config>router>bgp>group damp-peer-oscillations)

Full Context

configure router bgp group neighbor damp-peer-oscillations

configure router bgp damp-peer-oscillations

configure router bgp group damp-peer-oscillations

Description

This command controls how long a BGP peer session remains in the idle-state after some type of error causes the session to reset. In the idle state, BGP does not initiate or respond to attempts to establish a new session. Repeated errors that occur a short while after each session reset cause longer and longer hold times in the idle state. This command supports the DampPeerOscillations FSM behavior described in section 8.1 of RFC 4271, *A Border Gateway Protocol 4 (BGP-4).*

The default behavior, which applies when no damp-peer-oscillations is configured, is to immediately transition out of the idle-state after every reset.

Default

no damp-peer-oscillations

Parameters

initial-wait

The amount of time, in minutes, that a session remains in the idle-state after it has been stable for a while.

Values 0 to 2048

Default 0

second-wait

A period of time, in minutes, that is doubled after each repeated session failure that occurs within a relatively short span of time.

Values 1 to 2048

Default 5

max-wait

The maximum amount of time, in minutes, that a session remains in the idle-state after it has experienced repeated instability.

Values 1 to 2048

Default 60

error-interval

The interval of time, in minutes after a session reset, during which the session must be error-free in order to reset the penalty counter and return from idle-hold-time to initial-wait.

Values 0 to 2048

Default 30

Platforms

7705 SAR Gen 2

8.5 dampening

dampening

Syntax

dampening

Context

[Tree] (config>port>ethernet dampening)

Full Context

configure port ethernet dampening

Description

Commands in this context configure exponential port dampening for an Ethernet port.

Exponential Port Dampening (EPD) reduces the number of physical link transitions reported to upper layer protocols, potentially reducing upper layer protocol churn caused by a faulty link. Penalties are added against a port whenever the port's physical link state transitions from a link up state to a link down state. When the penalties exceed a configurable threshold, port-up and port-down transitions are no longer advertised to upper layers and the port's operational state will remain down until the penalty amount drops below a configurable reuse threshold. Each transition of link up state to link down state increments the accumulated penalty value by 1000. The accumulated penalties for a port are reduced at an exponential decay rate according to a configurable half-life parameter.

Platforms

7705 SAR Gen 2

8.6 damping

damping

Syntax

[**no**] **damping**

Context

[Tree] (config>service>vprn>bgp damping)

[Tree] (config>service>vprn>bgp>group>neighbor damping)

[Tree] (config>service>vprn>bgp>group damping)

Full Context

configure service vprn bgp damping

configure service vprn bgp group neighbor damping

configure service vprn bgp group damping

Description

This command enables BGP route damping for learned routes which are defined within the route policy. Use damping to reduce the number of update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set via route policy definition.

The **no** form of this command used at the global level disables route damping.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters:

Half-life: 15 minutes

Max-suppress: 60 minutes

Suppress-threshold: 3000

Reuse-threshold: 750

Default

no damping — Learned route damping is disabled.

Platforms

7705 SAR Gen 2

damping

Syntax

[**no**] **damping**

Context

[Tree] (config>router>bgp damping)

[Tree] (config>router>bgp>group>neighbor damping)

[Tree] (config>router>bgp>group damping)

Full Context

configure router bgp damping

configure router bgp group neighbor damping

configure router bgp group damping

Description

This command enables BGP route damping for learned routes which are defined within the route policy. Use damping to reduce the number of update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set via route policy definition.

The **no** form of this command used at the global level reverts route damping.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters:

- Half-life: 15 minutes
- Max-suppress: 60 minutes
- Suppress-threshold: 3000
- Reuse-threshold: 750

Default

no damping

Platforms

7705 SAR Gen 2

damping

Syntax

[**no**] **damping** *name*

Context

[Tree] (config>router>policy-options damping)

Full Context

configure router policy-options damping

Description

This command creates a context to configure a route damping profile to use in route policy entries.

The **no** form of this command deletes the named route damping profile.

Default

no damping

Parameters

name

Specifies the damping profile name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

damping

Syntax

damping {*name* | **none**}

no damping

Context

[Tree] (config>router>policy-options>policy-statement>entry>action damping)

[Tree] (config>router>policy-options>policy-statement>default-action damping)

Full Context

configure router policy-options policy-statement entry action damping

configure router policy-options policy-statement default-action damping

Description

This command configures a damping profile used for routes matching the route policy statement entry.

If no damping criteria is specified, the default damping profile is used.

The **no** form of this command removes the damping profile associated with the route policy entry.

Default

no damping

Parameters

name

The damping profile name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@," "start@variable@end"," @variable@end", or "start@variable@".

The *name* specified must already be defined.

none

Disables route damping for the route policy.

Platforms

7705 SAR Gen 2

8.7 data

data

Syntax

data [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]

no data

Context

[Tree] (debug>router>pim data)

Full Context

debug router pim data

Description

This command enables debugging for PIM data exception.

The **no** form of this command disables PIM data exception debugging.

Parameters

grp-ip-address

Debugs information associated with the specified data exception.

Values multicast group address (ipv4, ipv6)

ip-address

Debugs information associated with the specified data exception.

Values source address (ipv4, ipv6)

detail

Debugs detailed IP data exception information.

Platforms

7705 SAR Gen 2

8.8 database-export

database-export

Syntax

database-export [**identifier** *id*] [**bgp-ls-identifier** *bgp-ls-id*] [**reachable-ls-only**]

no database-export

Context

[Tree] (config>router>isis database-export)

Full Context

configure router isis database-export

Description

This command configures the population of the extended Traffic Engineering Database (TE-DB) with the link-state information from a specific IGP instance.

This information includes the IGP, TE, and the SR information, prefix SID sub-TLV, adjacency SID sub-TLV, and router SR capability TLV.

The **no** form of this command disables database exportation.

Parameters

id

Specifies an entry ID to export. This parameter is used to uniquely identify the IGP instance in the BGP-LS NLRI when a router has interfaces participating in multiple IGP instances. This parameter defaults to the IGP instance ID assigned by SR OS. However, because the concept of instance ID defined in IS-IS (RFC 6822) is unique within a routing domain while the one specified for OSPF is local subnet significant (RFC 6549), the user can remove any overlap by configuring the **identifier** value to be unique within a specific IGP domain when this router sends the IGP link-state information using BGP-LS.

Values 0 to 18446744073709551615

bgp-ls-id

Specifies a BGP-LS ID to export. This parameter is used, along with the Autonomous System Number (ASN) to correlate the BGP-LS NLRI advertisements of multiple BGP-LS speakers of the same IGP domain. If an NRC-P network domain has multiple IGP domains, BGP-LS speakers within each IGP domain must be configured with the same unique {bgp-ls-identifier, asn} tuple.

The BGP-LS identifier is optional and is only sent in a BGP-LS NLRI if configured in the IGP instance of an IGP domain.

Note: If this IGP instance participates in traffic engineering with RSVP-TE or SR-TE, the **traffic-engineering** command is not strictly required because enabling the extended TE-DB populates this information automatically. However, Nokia recommends enabling it to make the configuration consistent with other routers in the network that do not require the enabling of the extended TE-DB.

Values 0 to 4294967295

reachable-ls-only

Keyword to specify that only reachable link-state information is encoded. When this keyword is configured, the router, acting as a BGP-LS producer, must withdraw all link-state objects it has advertised in BGP, in accordance with section 5.9 of RFC 9552. This withdrawal occurs when the node that originated the corresponding LSPs is determined to be unreachable in the IGP based on the failure of a reachability check for that node. This withdrawal operation assists network controllers in assessing a reachable IGP topology, even in networks with segmented areas. For backward compatibility, the default behavior remains unchanged.

Platforms

7705 SAR Gen 2

database-export

Syntax

database-export [**identifier** *id*] [**bgp-ls-identifier** *bgp-ls-id*] [**reachable-ls-only**]

no database-export

Context

[Tree] (config>router>ospf database-export)

[Tree] (config>router>ospf3 database-export)

Full Context

configure router ospf database-export

configure router ospf3 database-export

Description

This command enables the population of the extended Traffic Engineering Database (TE-DB) with the link-state information from a specific IGP instance.

This information includes the IGP, TE, and the SR information, prefix SID sub-TLV, adjacency SID sub-TLV, and router SR capability TLV.

The **no** form of this command disables database exportation.

Default

no database-export

Parameters

id

Specifies an entry ID to export. This parameter is used to uniquely identify the IGP instance in the BGP-LS NLRI when a router has interfaces participating in multiple IGP instances. This parameter defaults to the IGP instance ID assigned by SR OS. The concept of instance ID specified for OSPF is local subnet significant (RFC 6549). The user can remove the router specific overlap by configuring the **identifier** value to be unique within a specific IGP domain when this router sends the IGP link-state information using BGP-LS.

Values 0 to 18446744073709551615

bgp-ls-id

Specifies a BGP-LS ID to export. This parameter is used, along with the ASN, to correlate the BGP-LS NLRI advertisements of multiple BGP-LS speakers of the same IGP domain. If an NRC-P network domain has multiple IGP domains, BGP-LS speakers within each IGP domain must be configured with the same unique {bgp-ls-identifier, asn} tuple.

The BGP-LS identifier is optional and is only sent in a BGP-LS NLRI if configured in the IGP instance of an IGP domain.

Note: If this IGP instance participates in traffic engineering with RSVP-TE or SR-TE, the **traffic-engineering** command is not strictly required because enabling the extended TE-DB populates this information automatically. However, Nokia recommends enabling it to make the configuration consistent with other routers in the network that do not require the enabling of the extended TE-DB.

Values 0 to 4294967295

reachable-ls-only

Keyword to specify that only reachable link-state information is encoded. When this keyword is configured, the router, acting as a BGP-LS producer, must withdraw all link-state objects it has advertised in BGP, in accordance with section 5.9 of RFC 9552. This withdrawal occurs when the node that originated the corresponding LSPs is determined to be unreachable in the IGP based on the failure of a reachability check for that node. This withdrawal operation assists network controllers in assessing a reachable IGP topology, even in networks with segmented areas. For backward compatibility, the default behavior remains unchanged.

Platforms

7705 SAR Gen 2

8.9 database-export-exclude

database-export-exclude

Syntax

[**no**] **database-export-exclude**

Context

[Tree] (config>router>isis>level database-export-exclude)

Full Context

configure router isis level database-export-exclude

Description

This command allows the user to prune the IGP link-state information of a specific IS-IS level from being exported into the extended TE-DB.

The **no** form of this command returns to the default behavior inherited from the database-export command at the IS-IS instance level.

Default

no database-export-exclude

Platforms

7705 SAR Gen 2

database-export-exclude

Syntax

[**no**] **database-export-exclude**

Context

[Tree] (config>router>ospf>area database-export-exclude)

[Tree] (config>router>ospf3>area database-export-exclude)

Full Context

configure router ospf area database-export-exclude

configure router ospf3 area database-export-exclude

Description

This command allows the user to prune the IGP link-state information of a specific OSPF level or OSPF area from being exported into the extended TE-DB. The **no** form of this command returns to the default behavior inherited from the database-export command at the OSPF or OSPF3 instance level.

Default

no database-export-exclude

Platforms

7705 SAR Gen 2

8.10 day-of-month

day-of-month

Syntax

day-of-month {*day-number* [*..day-number*] **all**}

no day-of-month

Context

[Tree] (config>system>cron>sched day-of-month)

Full Context

configure system cron schedule day-of-month

Description

This command specifies which days of the month that the schedule will occur. Multiple days of the month can be specified. When multiple days are configured, each of them will cause the schedule to trigger. If a day-of-month is configured without configuring month, weekday, hour, and minute, the event will not execute.

Using the **weekday** command as well as the **day-of-month** command will cause the script to run twice. For example, consider that today is Monday January 1. If Tuesday January 5 is configured, the script will run on Tuesday (tomorrow) as well as January 5 (Friday).

The **no** form of this command removes the specified day-of-month from the list.

Default

no day-of-month

Parameters

day-number

Specifies the positive integers specify the day of the month counting from the first of the month. The negative integers specify the day of the month counting from the last day of the month. For example, configuring **day-of-month -5, 5** in a month that has 31 days will specify the schedule to occur on the 27th and 5th of that month.

Integer values must map to a valid day for the month in question. For example, February 30 is not a valid date.

Values 1 to 31, -31 to -1 (maximum 62 day-numbers)

all

Specifies all days of the month.

Platforms

7705 SAR Gen 2

8.11 db

db

Syntax

db [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]

no db

Context

[Tree] (debug>router>pim db)

Full Context

debug router pim db

Description

This command enables debugging for PIM database.

The **no** form of this command disables PIM database debugging.

Parameters

grp-ip-address

Debugs information associated with the specified database.

Values multicast group address (ipv4, ipv6) or zero

ip-address

Debugs information associated with the specified database.

Values source address (ipv4, ipv6)

detail

Debugs detailed IP database information.

Platforms

7705 SAR Gen 2

db

Syntax

db [**detail**]

no db

Context

[Tree] (debug>router>pcep>pcc db)

[Tree] (debug>router>pcep>pcc>conn db)

Full Context

debug router pcep pcc db

debug router pcep pcc connection db

Description

This command enables debugging for PCC or connection database events.

The **no** form of this command disables debugging.

Parameters

detail

Keyword used to specify detailed information about PCC or connection database events.

Platforms
7705 SAR Gen 2

8.12 ddm-events

ddm-events

Syntax
[**no**] **ddm-events**

Context
[Tree] (config>port ddm-events)

Full Context
configure port ddm-events

Description
This command enables Digital Diagnostic Monitoring (DDM) events for the port.

The **no** form of this command disables DDM events.

Platforms
7705 SAR Gen 2

8.13 de-1-out-profile

de-1-out-profile

Syntax
[**no**] **de-1-out-profile**

Context
[Tree] (config>qos>sap-ingress>fc de-1-out-profile)

Full Context
configure qos sap-ingress fc de-1-out-profile

Description

This command, when enabled on a parent forwarding class, applies a color profile mode to the packets stored in the queue associated with this forwarding class. The queue associated with the parent forwarding class must be of type **profile-mode**.

When this QoS policy is applied to the ingress of a Frame Relay VLL SAP, the system will treat the received FR frames with DE bit set as out-of-profile, regardless of their previous marking as the result of the default classification or on a match with an IP filter. It also adjusts the CIR of the ingress SAP queue to consider out-of-profile frames that were sent while the SAP queue was in the "< CIR" state of the bucket. This makes sure that the CIR of the SAP is achieved.

All received DE = 0 frames that are classified into this parent forwarding class or any of its subclasses have their profile unchanged by enabling this option. That is, the DE = 0 frame profile could be undetermined (default), in-profile, or out-of-profile as per previous classification. The DE = 0 frames that have a profile of undetermined will be evaluated by the system CIR marking algorithm and will be marked appropriately.

The **priority** option, if used, has no effect. All FR VLL DE = 1 frames have their priority automatically set to low while DE = 0 frames have their priority set to high. Furthermore, DE = 1 frames have the drop-preference bit set in the internal header. The internal settings of the priority bit and of the drop-preference bit of the frame is independent of the use of the profile mode.

All other capabilities of the Fpipe service are maintained. This includes remarking of the DE bit on egress SAP, and FR PW control word on egress network port for the packets that were classified into "out-of-profile" at ingress SAP.

This **de-1-out-profile** keyword has an effect when applied to the ingress of a SAP that is part of an Fpipe service. It can also be used on the ingress of an Epipe or VPLS SAP.

The **no** form of this command disables the color profile mode of operation on all SAPs to which this ingress QoS policy is applied.

Default

no de-1-out-profile

Platforms

7705 SAR Gen 2

8.14 de-mark

de-mark

Syntax

de-mark [**force** *de-value*]

no demark

Context

[Tree] (config>qos>sap-egress>fc de-mark)

Full Context

configure qos sap-egress fc de-mark

Description

This command is used to explicitly define the marking of the DE bit for **fc** *fc-name* according to the inplus-profile/in-profile or out-of-profile/exceed-profile status of the packet (*fc-name* may be used to identify the dot1p-value).

If no DE value is present, the default values are used for the marking of the DE bit; for example, 0 for inplus-profile or in-profile packets, 1 for out-of-profile or exceed-profile packets. For more information, refer to the *IEEE 802.1ad-2005 standard*.

In the PBB case, for a Backbone SAP (B-SAP) and for packets originated from a local I-VPLS/PBB-Epipe, the command dictates the marking of the DE bit for both the BVID and ITAG.

If this command is not used, the DE bit should be preserved if an ingress TAG exist; otherwise, set to zero.

If the DE value is included in the command line, this value is to be used for all the packets of this forwarding class regardless of their profile status.

The commands **de-mark-inner** and **de-mark-outer** take precedence over the **de-mark** command if both are specified in the same policy.

Parameters

de-value

Specifies the DE marking value.

Values 0 or 1

Platforms

7705 SAR Gen 2

de-mark

Syntax

de-mark [**force** *de-value*]

no de-mark

Context

[Tree] (config>qos>network>egress>fc de-mark)

Full Context

configure qos network egress fc de-mark

Description

This command is used to explicitly define the marking of the DE bit for **fc** *fc-name* according to the inplus-profile or in-profile and out-of-profile or exceed-profile status of the packet (*fc-name* may be used to identify the dot1p value).

Parameters

de-value

Specifies that this value is to be used for all the packets of this forwarding class regardless of their profile status.

If no DE value is present, the default values are used for the marking of the DE bit; that is, 0 for inplus-profile and in-profile packets, 1 for out-of-profile and exceed-profile packets. For more information, refer to the *IEEE 802.1ad-2005 standard*.

In the PBB case, use the following rules for a network port (B-SDP):

- the outer VID follows the rules for regular SDP
- for packets originating from a local I-VPLS/PBB-Epipe, this command dictates the marking of the DE bit for both the outer (link level) BVID and ITAG; if the command is not used, the DE bit is set to zero.
- for transit packets (B-SAP/B-SDP to B-SDP), the related ITAG bits are preserved, the same as for BVID.

Values 0, 1

Platforms

7705 SAR Gen 2

8.15 de-mark-inner

de-mark-inner

Syntax

de-mark-inner [**force** *de-value*]

no de-mark-inner

Context

[Tree] (config>qos>sap-egress>fc de-mark-inner)

Full Context

configure qos sap-egress fc de-mark-inner

Description

This command is used to explicitly define the marking of the DE bit in the inner VLAN tag for **fc** *fc-name* on a QinQ SAP, according to the in- and out-of-profile status of the packet.

If no DE value is present, the default values are used for the marking of the DE bit; for example, 0 for inplus-profile or in-profile packets, 1 for out-of-profile or exceed-profile packets. For more information, refer to the *IEEE 802.1ad-2005 standard*.

If the DE value is included in the command line, this value is used for all the inner tags of packets of this forwarding class, regardless of their profile status.

This command takes precedence over the **de-mark** command if both are specified in the same policy and over the default action.

The configuration of **qinq-mark-top-only** under the SAP egress takes precedence over the use of the **de-mark-inner** in the policy. That is, the inner VLAN tag is not remarked when **qinq-mark-top-only** is configured (the marking used for the inner VLAN tag is based on the current default, which is governed by the marking of the packet received at the ingress to the system).

If no de-mark commands are used, the DE bit is preserved if an ingress inner tag exists; otherwise, set to 0.

Remarking the inner DE bit is not supported based on the profile result of egress policing.

Parameters

de-value

Specifies the DE marking value.

Values 0 or 1

Platforms

7705 SAR Gen 2

8.16 de-mark-outer

de-mark-outer

Syntax

de-mark-outer [**force** *de-value*]

no de-mark-outer

Context

[Tree] (config>qos>sap-egress>fc de-mark-outer)

Full Context

configure qos sap-egress fc de-mark-outer

Description

This command is used to explicitly define the marking of the DE bit in the outer or single VLAN tag on a qinq or dot1q SAP, respectively, according to the in, out, or exceed-profile status of the packet.

If no DE value is present, the default values are used for the marking of the DE bit; for example, 0 for inplus-profile/in-profile packets, 1 for out-of-profile/exceed-profile packets. For more information, refer to the *IEEE 802.1ad-2005 standard*.

If the DE value is included in the command line, this value is used for all the outer or single tags of packets of this forwarding class, regardless of their profile status.

In the PBB case, for a Backbone SAP (B-SAP) and for packets originated from a local I-VPLS/PBB-Epipe, the command dictates the marking of the DE bit for both the BVID and ITAG.

This command takes precedence over the **de-mark** command if both are specified in the same policy and over the default action.

If no de-mark commands are used, the DE bit is preserved if an ingress outer or single tag exists; otherwise, set to 0.

Parameters

de-value

Specifies the DE marking value.

Values 0 or 1

Platforms

7705 SAR Gen 2

8.17 dead-interval

dead-interval

Syntax

dead-interval *seconds*

no dead-interval

Context

[Tree] (config>service>vprn>ospf3>area>if dead-interval)

[Tree] (config>service>vprn>ospf>area>sham-link dead-interval)

[Tree] (config>service>vprn>ospf>area>virtual-link dead-interval)

[Tree] (config>service>vprn>ospf3>area>virtual-link dead-interval)

[Tree] (config>service>vprn>ospf>area>if dead-interval)

Full Context

configure service vprn ospf3 area interface dead-interval

configure service vprn ospf area sham-link dead-interval

configure service vprn ospf area virtual-link dead-interval

configure service vprn ospf3 area virtual-link dead-interval

configure service vprn ospf area interface dead-interval

Description

This command configures the time, in seconds, that OSPF waits before declaring a neighbor router down. If no Hello packets are received from a neighbor for the duration of the dead interval, the router is assumed to be down. The minimum interval must be two times the Hello interval.

The **no** form of this command reverts to the default value.

Default

dead-interval 40

Parameters

seconds

The dead interval expressed as a decimal integer.

Values 2 to 65535 seconds

Platforms

7705 SAR Gen 2

dead-interval

Syntax

dead-interval *seconds*

no dead-interval

Context

[Tree] (config>router>ospf3>area>virtual-link dead-interval)

[Tree] (config>router>ospf3>area>interface dead-interval)

[Tree] (config>router>ospf>area>virtual-link dead-interval)

[Tree] (config>router>ospf>area>interface dead-interval)

Full Context

configure router ospf3 area virtual-link dead-interval

configure router ospf3 area interface dead-interval

configure router ospf area virtual-link dead-interval

configure router ospf area interface dead-interval

Description

This command configures the time, in seconds, that OSPF waits before declaring a neighbor router down. If no hello packets are received from a neighbor for the duration of the dead interval, the router is assumed to be down. The minimum interval must be two times the hello interval.

The **no** form of this command reverts to the default value.

Default

dead-interval 40

Parameters

seconds

The dead interval expressed in seconds.

Values 2 to 65535

Platforms

7705 SAR Gen 2

8.18 dead-timer

dead-timer

Syntax

dead-timer *seconds*

no dead-timer

Context

[Tree] (config>router>pcep>pcc dead-timer)

Full Context

configure router pcep pcc dead-timer

Description

This command configures the PCEP session dead timer value, which is the amount of time a PCEP speaker (PCC or PCE) will wait after the receipt of the last PCEP message before declaring its peer down.

The keep-alive mechanism is asymmetric, meaning that each PCEP speaker can propose a different dead timer value to its peer to use to detect session timeout.

The **no** form of the command returns the dead timer to the default value.

Default

dead-timer 120

Parameters

seconds

the dead timer value, in seconds

Values 1 to 255

Platforms

7705 SAR Gen 2

8.19 debounce

debounce

Syntax

debounce *occurrences* [**within** *seconds*]

no debounce

Context

[Tree] (config>log>event-trigger>event>trigger-entry debounce)

Full Context

configure log event-trigger event trigger-entry debounce

Description

This command configures when to trigger, for example after one or more event occurrences. The number of occurrences of an event can be bounded by a time window or left open.

The **no** form of this command removes the debounce configuration.

Parameters

occurrences

Specifies the number of times an event must occur for EHS to trigger a response.

Values 2 to 15

within *seconds*

Specifies the time window within which a specific event must occur a number of times equivalent to the specified *occurrences* for EHS to trigger a response.

Values 1 to 604800

Platforms

7705 SAR Gen 2

8.20 debug

debug

Syntax
debug

Context
[Tree] (debug)

Full Context
debug

Description
Commands in this context specify debugging options.

Platforms
7705 SAR Gen 2

8.21 debug-save

debug-save

Syntax
debug-save [*file-url*]

Context
[Tree] (admin debug-save)

Full Context
admin debug-save

Description
This command saves existing debug configuration (configuration done under the debug branch of CLI). Debug configurations are not saved by the **admin save** command and not preserved across a node reboot or CPM switchover. The **debug-save** command makes the debug configuration available for the operator to execute after a reboot by using the **exec** command or after a CPM switchover by using the **switchover-exec** command, if desired.

Parameters

file-url

Specifies the file URL location to save the debug configuration. If no file-url is specified then the debug configuration is saved at the same location as the standard configuration file (**bof>primary-config**/**bof>secondary-config**/**bof>tertiary-config**) with the same file name as the standard configuration file but with a .dbg suffix.

Values

file url	local-url | remote-url: 255 chars max
local-url	[*cflash-id*/][*file-path*] 200 chars max, including cflash-id *file-path* 199 chars max
remote-url	[{ftp:// | tftp://}*login*:*pswd@remote-locn*/][*file-path*]
	255 chars max
	directory length 99 chars max each
remote-locn	{hostname | ipv4-address | [ipv6-address]}
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D
	interface - 32 chars max, for link local addresses 255
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

8.22 def-recv-evpn-encap

def-recv-evpn-encap

Syntax

def-recv-evpn-encap {**mpls** | **vxlan**}

no def-recv-evpn-encap

Context

[Tree] (config>router>bgp>group>neighbor def-recv-evpn-encap)

[Tree] (config>router>bgp>group def-recv-evpn-encap)

Full Context

configure router bgp group neighbor def-recv-evpn-encap

configure router bgp group def-recv-evpn-encap

Description

This command defines how the BGP will treat a received EVPN route without RC5512 BGP encapsulation extended community. If no encapsulation is received, BGP will validate the route as MPLS or VXLAN depending on how this command is configured.

Default

no def-recv-evpn-encap

Parameters

mpls

Specifies that **mpls** is the default encapsulation value in the case where no RFC 5512 extended community is received in the incoming BGP-EVPN route.

vxlan

Specifies that **vxlan** is the default encapsulation value.

Platforms

7705 SAR Gen 2

8.23 default

default

Syntax

[**no**] **default**

Context

[Tree] (config>log>accounting-policy default)

Full Context

configure log accounting-policy default

Description

This command configures the default accounting policy to be used with all SAPs that do not have an accounting policy.

If no access accounting policy is defined on a SAP, accounting records are produced in accordance with the default access policy. If no default access policy is created, then no accounting records will be collected other than the records for the accounting policies that are explicitly configured.

If no network accounting policy is defined on a network port, accounting records will be produced in accordance with the default network policy. If no network default policy is created, then no accounting records will be collected other than the records for the accounting policies explicitly configured.

Only one access accounting policy ID can be designated as the default access policy. Likewise, only one network accounting policy ID can be designated as the default network accounting policy.

The record name must be specified prior to assigning an accounting policy as default.

If a policy is configured as the default policy, then a **no default** command must be issued before a new default policy can be configured.

The **no** form of this command removes the default policy designation from the policy ID. The accounting policy will be removed from all SAPs or network ports that do not have this policy explicitly defined.

Platforms

7705 SAR Gen 2

8.24 default-action

default-action

Syntax

default-action {**bypass-host-creation** | **drop**}

no default-action

Context

[Tree] (config>filter>dhcp-filter default-action)

Full Context

configure filter dhcp-filter default-action

Description

This command specifies the default action for DHCP filters when no entries match.

The **no** form of this command reverts to the default.

Parameters

bypass-host-creation

Specifies to bypass ESM host creation options.

drop

Specifies to drop and not process the DHCP message.

Platforms

7705 SAR Gen 2

default-action

Syntax

default-action bypass-host-creation [**na**] [**pd**]

default-action drop

no default-action

Context

[Tree] (config>filter>dhcp6-filter default-action)

Full Context

configure filter dhcp6-filter default-action

Description

This command specifies the default action for DHCP6 filters when no entries match.

The **no** form of this command reverts to the default.

Parameters

bypass-host-creation

Specifies to bypass ESM host creation options.

Values **na** — Bypasses the DHCP NA hosts creation.
pd — Bypasses the DHCP PD hosts creation.

drop

Specifies to drop and not process the DHCP6 message.

Platforms

7705 SAR Gen 2

default-action

Syntax

default-action {**drop** | **forward**}

no default-action

Context

[Tree] (config>service>vprn>log>filter default-action)

Full Context

configure service vprn log filter default-action

Description

The default action specifies the action that is applied to events when no action is specified in the event filter entries or when an event does not match the specified criteria.

When multiple **default-action** commands are entered, the last command overwrites the previous command.

The **no** form of this command reverts the default action to the default value (forward).

Default

default-action forward — The events which are not explicitly dropped by an event filter match are forwarded.

Parameters

drop

The events which are not explicitly forwarded by an event filter match are dropped.

forward

The events which are not explicitly dropped by an event filter match are forwarded.

Platforms

7705 SAR Gen 2

default-action

Syntax

default-action fc *fc-name* **profile** {**in** | **out**}

Context

[Tree] (config>qos>network>ingress default-action)

Full Context

configure qos network ingress default-action

Description

This command defines or edits the default action to be taken for packets that have an undefined DSCP or MPLS EXP bit set. The **default-action** command specifies the forwarding class to which such packets are assigned.

Multiple default-action commands will overwrite each previous default-action command.

Default

default-action fc be profile out

Parameters

fc-name

Specifies the forwarding class name. All packets with DSCP value or MPLS EXP or dot1p bits that are not defined will be placed in this forwarding class.

Values be, l2, af, l1, h2, ef, h1, nc

profile {in | out}

All packets that are assigned to this forwarding class will be considered in-profile or out-of-profile based on this command. In case of congestion, the in-profile packets are preferentially queued over the out-of-profile packets.

Values in, out

Platforms

7705 SAR Gen 2

default-action

Syntax

default-action {**drop** | **forward**}

Context

[Tree] (config>filter>ip-filter default-action)

[Tree] (config>filter>ipv6-filter default-action)

Full Context

configure filter ip-filter default-action

configure filter ipv6-filter default-action

Description

This command defines the default action to be applied to packets not matching any entry in this ACL filter policy or to packets for that match a PBF/PBR filter entry for which the PBF/PBR target is down and **pbr-down-action-override** per-entry is set to **filter-default-action**.

Default

default-action drop

Parameters

drop

Specifies the default action is to drop a packet.

forward

Specifies the default action is to forward a packet.

Platforms

7705 SAR Gen 2

default-action

Syntax

default-action {**drop** | **forward**}

no default-action

Context

[Tree] (config>log>filter default-action)

Full Context

configure log filter default-action

Description

The default action specifies the action that is applied to events when no action is specified in the event filter entries or when an event does not match the specified criteria.

When multiple **default-action** commands are entered, the last command overwrites the previous command.

The **no** form of this command reverts the default action to the default value (forward).

Default

default-action forward

Parameters

drop

The events which are not explicitly forwarded by an event filter match are dropped.

forward

The events which are not explicitly dropped by an event filter match are forwarded.

Platforms

7705 SAR Gen 2

default-action

Syntax

default-action {**permit** | **deny** | **deny-host-unreachable**}

Context

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter default-action)

[Tree] (config>system>security>mgmt-access-filter>ip-filter default-action)

[Tree] (config>system>security>mgmt-access-filter>mac-filter default-action)

Full Context

configure system security management-access-filter ipv6-filter default-action

configure system security management-access-filter ip-filter default-action

configure system security management-access-filter mac-filter default-action

Description

This command creates the default action for management access in the absence of a specific management access filter match.

The **default-action** is applied to a packet that does not satisfy any match criteria in any of the management access filters. Whenever management access filters are configured, the **default-action** must be defined.

Parameters

permit

Specifies that packets not matching the configured selection criteria in any of the filter entries will be permitted.

deny

Specifies that packets not matching the selection criteria be denied and that an ICMP host unreachable message will not be issued.

deny-host-unreachable

Specifies that packets not matching the selection criteria be denied access and that an ICMP host unreachable message will be issued.

The **deny-host-unreachable** only applies to ip-filter and ipv6filter.

Platforms

7705 SAR Gen 2

default-action

Syntax

default-action {**deny-all** | **permit-all** | **none** | **read-only-all**}

Context

[Tree] (config>system>security>profile default-action)

Full Context

configure system security profile default-action

Description

This command specifies the default action to be applied when no match conditions are met in the list of **profile entry match** commands. It does not apply in any way to other parts of the profile such as **grpc rpc-authorization** or **netconf base-op-authorization**.

When a user is a member of multiple profiles, profiles are evaluated in the order that they are configured. Evaluation stops if there is a match, or when the default action of the profile is **deny-all**, **permit-all**, or **read-only-all**. When the profile default action is **none** and if no match conditions are met in the profile, the next profile is evaluated. When the default action of the last profile is **none** and no explicit match is found, the command is denied.

Parameters

deny-all

Sets the default of the profile to deny access to all commands.

permit-all

Sets the default of the profile to permit access to all commands.

Note:
In classic CLI but not in MD-CLI the **permit-all** parameter does not change access to **security** commands. Specific entries must be created in a command authorization profile in order to give access to **security** commands. The system populated "administrative" profile contains rules to access **security** commands.

none

Sets the default of the profile to no-action. This option is useful to assign multiple profiles to a user.

read-only-all

Sets the default of the profile to allow read-only access to all commands.

Platforms

7705 SAR Gen 2

default-action

Syntax

default-action {**accept** | **next-entry** | **next-policy** | **drop** | **reject**}

no default-action

Context

[Tree] (config>router>policy-options>policy-statement default-action)

Full Context

configure router policy-options policy-statement default-action

Description

Commands in this context configure actions to apply to routes that do not match any entries of a route policy statement.

The **no** form of this command deletes the **default-action** context for the policy statement.

Default

no default-action

Parameters

accept

Specifies that routes not matched by any entry should be allowed or accepted. This parameter provides a context for modifying route properties.

next-entry

Specifies that routes not matched by any entry should be evaluated by the next sequential entry in the policy chain, after route properties are possibly modified by the default action of the current policy.

next-policy

Specifies that routes not matched by any entry should be evaluated by the next sequential policy in the policy chain, after route properties are possibly modified by the default action of the current policy.

drop

Specifies that routes not matched by any entry should be disallowed or rejected. This parameter provides a context for modifying route properties.

reject

Specifies that routes not matched by any entry should be disallowed or rejected. This parameter does not provide a context for modifying route properties.

Platforms

7705 SAR Gen 2

8.25 default-domain

default-domain

Syntax

default-domain *dns-name*

no default-domain

Context

[Tree] (config>service>vprn>dns default-domain)

Full Context

configure service vprn dns default-domain

Description

This command configures the DNS domain name to be added in DNS retries when a DNS query is not replied or an empty DNS reply is received.

The **no** form of this command prevents DNS retries when the DNS query is not replied or an empty DNS reply is received.

Parameters

dns-name

Specifies the name of the default domain, up to 255 characters. Allowed values for characters are alphabetical (A-Z), numeric (0-9), the minus sign (-), and the period (.). For example, "3gpp-network.org”.

Platforms

7705 SAR Gen 2

8.26 default-fc

default-fc

Syntax

default-fc *fc-name*

no default-fc

Context

[Tree] (config>qos>sap-ingress default-fc)

Full Context

configure qos sap-ingress default-fc

Description

This command configures the default forwarding class for the policy. If an ingress packet does not match a higher priority (more explicit) classification command, the default forwarding class or subclass if associated with the packet. Unless overridden by an explicit forwarding class classification rule, all packets received on an ingress SAP using this ingress QoS policy are classified to the default forwarding class. Optionally, the default ingress enqueuing priority for the traffic can be overridden as well.

The default forwarding class is best effort (be). The **default-fc** settings are displayed in the **show configuration** and **save** output regardless of inclusion of the **detail** keyword.

Default

default-fc "be”

Parameters

fc-name

Specify the forwarding class name for the queue. The value specified for *fc-name* must be one of the predefined forwarding classes in the system.

The subclass-name parameter is optional and used with the fc-name parameter to define a preexisting subclass. The fc-name and subclass-name parameters must be separated by a period (dot). If subclass-name does not exist in the context of fc -name, an error will occur. If subclass-name is removed using the **no fc** *fc-name.subclass-name* **force** command, the **default-fc** command will automatically drop the *subclass-name* and only use *fc-name* (the parent forwarding class for the subclass) as the forwarding class.

Values

fc: *class*[.*subclass*]

class: be, l2, af, l1, h2, ef, h1, nc

subclass: 29 characters max

Platforms

7705 SAR Gen 2

8.27 default-instance

default-instance

Syntax

[**no**] **default-instance**

Context

[Tree] (config>service>vprn>isis>if default-instance)

Full Context

configure service vprn isis interface default-instance

Description

This command enables a non-MI capable router to establish an adjacency and operate with an SR OS in a non-zero instance. If the router does not receive IID-TLVs, it establishes an adjacency in a single instance. Instead of establishing an adjacency in the standard instance 0, the router establishes an adjacency in the configured non-zero instance. The router then operates in the configured non-zero instance so that it appears to be in the standard instance 0 to its neighbor. This feature is supported on point-to-point interfaces, broadcast interfaces are not supported.

The **no** form of this command disables the functionality so that the router can only establish adjacencies in the standard instance 0.

Default

no default-instance

Platforms

7705 SAR Gen 2

default-instance

Syntax

[**no**] **default-instance**

Context

[Tree] (config>router>isis>interface default-instance)

Full Context

configure router isis interface default-instance

Description

This command enables a non-MI capable router to establish an adjacency and operate with a router in a non-zero instance. If the router does not receive IID-TLVs, it will establish an adjacency in a single instance. Instead of establishing an adjacency in the standard instance 0, the router will establish an adjacency in the configured non-zero instance. The router will then operate in the configured non-zero instance so that it appears to be in the standard instance 0 to its neighbor. This feature is supported on point-to-point interfaces, broadcast interfaces are not supported.

This feature must be configured on the router connected to non-MI capable routers and on all other SR OS routers in the area, so that they receive non-MI LSPs in the correct instance and not in the base instance.

The **no** form of this command disables the functionality so that the router can only establish adjacencies in the standard instance 0.

Default

no default-instance

Platforms

7705 SAR Gen 2

8.28 default-ipv4-multicast-metric

default-ipv4-multicast-metric

Syntax

default-ipv4-multicast-metric *metric*

no default-ipv4-multicast-metric

Context

[Tree] (config>service>vprn>isis>level default-ipv4-multicast-metric)

Full Context

configure service vprn isis level default-ipv4-multicast-metric

Description

This command configures the default metric to be used for the IS-IS interface in the IPv4 multicast topology (MT3).

The **no** form of this command deletes the specified default metric and reverts to using the system default of 10.

Default

default-ipv4-multicast-metric 10

Parameters

metric

Specifies the default metric for interfaces in the IPv4 multicast topology (MT3).

Values 1 to 16777215

Platforms

7705 SAR Gen 2

default-ipv4-multicast-metric

Syntax

default-ipv4-multicast-metric *metric*

no default-ipv4-multicast-metric

Context

[Tree] (config>router>isis>level default-ipv4-multicast-metric)

Full Context

configure router isis level default-ipv4-multicast-metric

Description

This command configures the default metric to be used for the IS-IS interface in the IPv4 multicast topology (MT3).

The **no** form of this command deletes the specified default metric and reverts to using the system default of 10.

Default

default-ipv4-multicast-metric 10

Parameters

metric

Specifies the default metric for interfaces in the IPv4 multicast topology (MT3).

Values 1 to 16777215

Platforms

7705 SAR Gen 2

8.29 default-ipv6-multicast-metric

default-ipv6-multicast-metric

Syntax

default-ipv6-multicast-metric *metric*

no default-ipv6-multicast-metric

Context

[Tree] (config>router>isis>level default-ipv6-multicast-metric)

Full Context

configure router isis level default-ipv6-multicast-metric

Description

This command configures the default metric to be used for the IS-IS interface in the IPv6 multicast topology (MT4).

The **no** form of this command deletes the specified default metric and reverts to using the system default of 10.

Default

default-ipv6-multicast-metric 10

Parameters

metric

Specifies the default metric for interfaces in the IPv4 multicast topology (MT4).

1 to 16777215

Platforms

7705 SAR Gen 2

8.30 default-ipv6-unicast-metric

default-ipv6-unicast-metric

Syntax

default-ipv6-unicast-metric *ipv6 metric*

no default-ipv6-unicast-metric

Context

[Tree] (config>service>vprn>isis>level default-ipv6-unicast-metric)

Full Context

configure service vprn isis level default-ipv6-unicast-metric

Description

This command specifies the default metric for IPv6 unicast.

Default

default-ipv6-unicast-metric 10

Parameters

ipv6-metric

Specifies the default metric for IPv6 unicast.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

default-ipv6-unicast-metric

Syntax

default-ipv6-unicast-metric *ipv6 metric*

no default-ipv6-unicast-metric

Context

[Tree] (config>router>isis>level default-ipv6-unicast-metric)

Full Context

configure router isis level default-ipv6-unicast-metric

Description

This command specifies the default metric for IPv6 unicast.

The **no** form of this command reverts to the default value.

Default

default-ipv6-unicast-metric 10

Parameters

ipv6-metric

Specifies the default metric for IPv6 unicast.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

8.31 default-label-preference

default-label-preference

Syntax

default-label-preference [**ebgp** *ebgp label preference*] [**ibgp** *ibgp label preference*]

no default-label-preference

Context

[Tree] (config>router>bgp>group>neighbor default-label-preference)

[Tree] (config>router>bgp>group default-label-preference)

[Tree] (config>router>bgp default-label-preference)

Full Context

configure router bgp group neighbor default-label-preference

configure router bgp group default-label-preference

configure router bgp default-label-preference

Description

This command specifies a route-table preference value to use for EBGP or IBGP routes carrying labeled-unicast prefixes and received from peers covered by the context of the command. Route-table preference comes into play when the route-table has multiple routes for the same IP prefix. In this case the route with the numerically lowest preference value is usually the route that is activated and installed into the IP FIB. By default all BGP routes have a route-table preference value of 170.

This command overrides the preference value assigned by the **label-preference** command; that other command does not distinguish between EBGP and IBGP routes. Overriding happens even when the default-label-preference value is inherited from a higher level of configuration and competes with an explicitly configured label-preference value at a lower level of configuration in the BGP hierarchy.

Note:
The preference value assigned by the **default-label-preference** command can always be overwritten by a route policy entry that accepts the route with a **preference** command in the action.

The **no** form of the command lets BGP route-table preference for labeled-unicast routes to be controlled by other means.

Default

no default-label-preference

Parameters

ebgp label preference

Specifies the EBGP default preference label value.

Values 0 to 255

ibgp label preference

Specifies the IBGP default preference label value.

Values 0 to 255

Platforms

7705 SAR Gen 2

8.32 default-metric

default-metric

Syntax

default-metric *ipv4 metric*

no default-metric

Context

[Tree] (config>service>vprn>isis>level default-metric)

Full Context

configure service vprn isis level default-metric

Description

This command specifies the configurable default metric used for all IS-IS interfaces on this level. This value is not used if a metric is configured for an interface.

Default

default-metric 10

Parameters

ipv4 metric

Specifies the default metric for IPv4 unicast.

Values 1 to 16777214

Platforms

7705 SAR Gen 2

default-metric

Syntax

default-metric *metric*

no default-metric

Context

[Tree] (config>service>vprn>ospf3>area>stub default-metric)

[Tree] (config>service>vprn>ospf>area>stub default-metric)

Full Context

configure service vprn ospf3 area stub default-metric

configure service vprn ospf area stub default-metric

Description

This command configures the metric used by the area border router (ABR) for the default route into a stub area. The default metric should only be configured on an ABR of a stub area. An ABR generates a default route if the area is a **stub** area.

The **no** form of this command reverts to the default value.

Default

default-metric 1

Parameters

metric

The metric expressed as a decimal integer for the default route cost to be advertised into the stub area.

Values 1 to 16777214

Platforms

7705 SAR Gen 2

default-metric

Syntax

default-metric *ipv4 metric*

no default-metric

Context

[Tree] (config>router>isis>level default-metric)

Full Context

configure router isis level default-metric

Description

This command specifies the configurable default metric used for all IS-IS interfaces on this level. This value is not used if a metric is configured for an interface.

The **no** form of this command reverts to the default value.

Default

default-metric 10

Parameters

ipv4 metric

Specifies the default metric for IPv4 unicast.

Values 1 to 16777214

Platforms

7705 SAR Gen 2

default-metric

Syntax

default-metric *metric*

no default-metric

Context

[Tree] (config>router>ospf3>area>stub default-metric)

[Tree] (config>router>ospf>area>stub default-metric)

Full Context

configure router ospf3 area stub default-metric

configure router ospf area stub default-metric

Description

This command configures the metric used by the area border router (ABR) for the default route into a stub area.

The default metric should only be configured on an ABR of a stub area.

An ABR generates a default route if the area is a **stub** area.

The **no** form of this command reverts to the default value.

Default

default-metric 1

Parameters

metric

Specifies the metric expressed as a decimal integer for the default route cost to be advertised into the stub area.

Values 1 to 16777214

Platforms

7705 SAR Gen 2

8.33 default-path

default-path

Syntax

default-path *path-name*

Context

[Tree] (config>router>mpls>lsp-template default-path)

Full Context

configure router mpls lsp-template default-path

Description

A default path binding must be provided before the LSP template can be used for signaling LSP. The LSP template must be shutdown to modify default-path binding.

Parameters

path-name

Configures the default path binding

Platforms

7705 SAR Gen 2

8.34 default-preference

default-preference

Syntax

default-preference [**ebgp** *ebgp preference*] [**ibgp** *ibgp preference*]

no default-preference

Context

[Tree] (config>router>bgp>group default-preference)

[Tree] (config>router>bgp>group>neighbor default-preference)

[Tree] (config>router>bgp default-preference)

Full Context

configure router bgp group default-preference

configure router bgp group neighbor default-preference

configure router bgp default-preference

Description

This command specifies a route-table preference value to use for EBGP or IBGP routes carrying unlabeled prefixes and received from peers covered by the context of the command. Route-table preference comes into play when the route-table has multiple routes for the same IP prefix. In this case, the route with the numerically lowest preference value is usually the route that is activated and installed into the IP FIB. By default all BGP routes have a route-table preference value of 170.

This command overrides the preference value assigned by the **preference** command; that other command does not distinguish between EBGP and IBGP routes. Overriding happens even when the default-preference value is inherited from a higher level of configuration and competes with an explicitly configured preference value at a lower level of configuration in the BGP hierarchy.

Note:
The preference value assigned by the **default-preference** command can always be overwritten by a route policy entry that accepts the route with a **preference** command in the action.

The **no** form of the command lets BGP route-table preference to be controlled by other means.

Default

no default-preference

Parameters

ebgp preference

Specifies the EBGP default preference value.

Values 0 to 255

ibgp preference

Specifies the IBGP default preference value.

Values 0 to 255

Platforms

7705 SAR Gen 2

8.35 default-priority

default-priority

Syntax

default-priority {**high** | **low**}

no default-priority

Context

[Tree] (config>qos>sap-ingress default-priority)

Full Context

configure qos sap-ingress default-priority

Description

This command configures the default enqueuing priority for all packets received on an ingress SAP using this policy. To change the default priority for the policy, the **fc-name** must be defined whether it is being changed or not.

Default

default-priority low

Parameters

high

Setting the enqueuing parameter to **high** for a packet increases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

low

Setting the enqueuing parameter to **low** for a packet decreases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

Platforms

7705 SAR Gen 2

8.36 default-result

default-result

Syntax

default-result {**revoked** | **good**}

no default-result

Context

[Tree] (config>service>vprn>if>sap>ipsec-gw>cert>status-verify default-result)

[Tree] (config>ipsec>trans-mode-prof>dyn>cert>status-verify default-result)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>dyn>cert>status-verify default-result)

[Tree] (config>service>ies>if>sap>ipsec-gw>cert>status-verify default-result)

[Tree] (config>service>vprn>if>sap>ipsec-tun>dyn>cert>status-verify default-result)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>dyn>cert>status-verify default-result)

[Tree] (config>router>if>ipsec>ipsec-tun>dyn>cert>status-verify default-result)

Full Context

configure service vprn interface sap ipsec-gw cert status-verify default-result

configure ipsec ipsec-transport-mode-profile dynamic-keying cert status-verify default-result

configure service vprn interface ipsec ipsec-tunnel dynamic-keying cert status-verify default-result

configure service ies interface sap ipsec-gw cert status-verify default-result

configure service vprn interface sap ipsec-tunnel dynamic-keying cert status-verify default-result

configure service ies interface ipsec ipsec-tunnel dynamic-keying cert status-verify default-result

configure router interface ipsec ipsec-tunnel dynamic-keying cert status-verify default-result

Description

This command specifies the default certificate revocation status that is used result when both the primary and secondary CSV methods fail to verify the status.

Default

default-result revoked

Parameters

good

Specifies that the certificate is considered as acceptable.

revoked

Specifies that the certificate is considered as revoked.

Platforms

7705 SAR Gen 2

default-result

Syntax

default-result *certificate-tls-status*

no default-result

Context

[Tree] (config>system>security>tls>client-tls-profile>status-verify default-result)

[Tree] (config>system>security>tls>server-tls-profile>status-verify default-result)

Full Context

configure system security tls client-tls-profile status-verify default-result

configure system security tls server-tls-profile status-verify default-result

Description

This command configures the default result of the entity certificate verification in the TLS client or server profile. This command overwrites the EE certificate revocation verification for the TLS client or server profile.

The **no** form of this command leaves the configured default result unchanged.

Default

default-result revoked

Parameters

certificate-tls-status

Specifies the certificate status.

Values **good** — Keyword to specify that the certificate is acceptable.

revoked — Keyword to specify that the certificate is considered revoked.

Platforms

7705 SAR Gen 2

8.37 default-route-tag

default-route-tag

Syntax

default-route-tag *tag*

no default-route-tag

Context

[Tree] (config>service>vprn>bgp-evpn>mpls default-route-tag)

[Tree] (config>service>epipe>bgp-evpn>mpls default-route-tag)

[Tree] (config>service>vpls>bgp-evpn>mpls default-route-tag)

Full Context

configure service vprn bgp-evpn mpls default-route-tag

configure service epipe bgp-evpn mpls default-route-tag

configure service vpls bgp-evpn mpls default-route-tag

Description

This command configures a route tag that EVPN and IP-VPN use when sending a route to the BGP application (for the corresponding service and BGP instance). If the corresponding BGP EVPN instance is enabled, the command cannot be changed. Additionally, EVPN services can add tags to routes with **proxy-arp/nd>evpn-route-tag** or the route table tag (added using the import policy). Only one tag is passed from EVPN to the BGP for matching on export policies. In case of a conflict with other route tags pushed by EVPN, the default route tag has the least priority.

The following are examples of the conflict priority handling:

- If a service is configured with both **default-route-tag** *X* and **proxy-arp>evpn-route-tag** *Y*, EVPN uses route tag *Y* when sending EVPN proxy-arp routes to the BGP RIB for advertisement.
- If a given IP-prefix route is tagged in the route-table with tag *A* and the R-VPLS, in which the route is advertised, uses *B* as the **default-route-tag**, then EVPN keeps tag *A* when sending the route to the BGP RIB.

The **default-route-tag** configuration is only supported on EVPN and IP-VPN service routes. The route tag for ES and AD per-ES routes is always zero.

The **no** form of this command removes the **default-route-tag** (that is, it sets the route tag to zero).

Default

no default-route-tag

Parameters

tag

Specifies the route tag.

Values 1 to 255

Platforms

7705 SAR Gen 2

default-route-tag

Syntax

default-route-tag *tag*

no default-route-tag

Context

[Tree] (config>service>vprn>isis default-route-tag)

Full Context

configure service vprn isis default-route-tag

Description

This command configures the route tag for default route for the router or VPRN service.

Parameters

tag

Assigns a default tag.

Values 1 — 4294967295

Platforms

7705 SAR Gen 2

default-route-tag

Syntax

default-route-tag *tag*

no default-route-tag

Context

[Tree] (config>router>isis default-route-tag)

Full Context

configure router isis default-route-tag

Description

This command configures the route tag for default route.

Parameters

tag

Assigns a default tag.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

8.38 default-route-target

default-route-target

Syntax

[**no**] **default-route-target**

Context

[Tree] (config>router>bgp>group default-route-target)

[Tree] (config>router>bgp>group>neighbor default-route-target)

Full Context

configure router bgp group default-route-target

configure router bgp group neighbor default-route-target

Description

This command originates the default RTC route (zero prefix length) towards the selected peers.

Default

no default-route-target

Platforms

7705 SAR Gen 2

8.39 default-router

default-router

Syntax

default-router *ip-address* [*ip-address*]

no default-router

Context

[Tree] (config>router>dhcp>server>pool>subnet>options default-router)

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options default-router)

Full Context

configure router dhcp local-dhcp-server pool subnet options default-router

configure subscriber-mgmt local-user-db ipoe host options default-router

Description

This command configures the IP address of the default router for a DHCP client. Up to four IP addresses can be specified.

The **no** form of this command removes the address(es) from the configuration.

Parameters

ip-address

Specifies up to four default router IP addresses. Each address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

Values a.b.c.d

Platforms

7705 SAR Gen 2

8.40 default-sap

default-sap

Syntax

[**no**] **default-sap**

Context

[Tree] (config>service>vpls>sap>managed-vlan-list default-sap)

Full Context

configure service vpls sap managed-vlan-list default-sap

Description

This command adds a default SAP to the managed VLAN list.

The **no** form of this command removes the default SAP to the managed VLAN list.

Platforms

7705 SAR Gen 2

8.41 default-secure-service

default-secure-service

Syntax

default-secure-service *service-id ip-int-name*

no default-secure-service

Context

[Tree] (config>service>ies>if>sap>ipsec-gw default-secure-service)

[Tree] (config>service>vprn>if>sap>ipsec-gw default-secure-service)

Full Context

configure service ies interface sap ipsec-gw default-secure-service

configure service vprn interface sap ipsec-gw default-secure-service

Description

This command specifies a service ID or service name of the default security service used by this SAP IPsec gateway.

Parameters

service-id

Specifies a default secure service.

Values *service-id*: 1 to 2147483647*svc-name:* An existing service name up to 64 characters.

ip-int-name

The name of private IPsec tunnel interface.

Platforms
7705 SAR Gen 2

8.42 default-tunnel-template

default-tunnel-template

Syntax
default-tunnel-template *ipsec-template-identifier*
no default-tunnel-template

Context
[Tree] (config>service>ies>if>sap>ipsec-gw default-tunnel-template)
[Tree] (config>service>vprn>if>sap>ipsec-gw default-tunnel-template)

Full Context
configure service ies interface sap ipsec-gw default-tunnel-template
configure service vprn interface sap ipsec-gw default-tunnel-template

Description
This command configures a default tunnel policy template for the gateway.

Platforms
7705 SAR Gen 2

8.43 delay

delay

Syntax
delay

Context
[Tree] (config>router>if>if-attribute delay)

Full Context
configure router interface if-attribute delay

Description

Commands in this context configure or apply delay interface attributes such as static delay.

Platforms

7705 SAR Gen 2

delay

Syntax

delay *interval*

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>progress-indicator delay)

Full Context

configure system management-interface cli md-cli environment progress-indicator delay

Description

This command sets the delay before the progress indicator is displayed in the MD-CLI.

Default

delay 500

Parameters

interval

Specifies the delay interval, in milliseconds.

Values 1 to 10000

Platforms

7705 SAR Gen 2

8.44 delay-event

delay-event

Syntax

delay-event {**forward** | **backward** | **round-trip**} **lowest-bin** *bin-number* **threshold***raise-threshold* [**clear** *clear-threshold*]

no delay-event {**forward** | **backward** | **round-trip**}

Context

[Tree] (config>oam-pm>bin-group>bin-type delay-event)

Full Context

configure oam-pm bin-group bin-type delay-event

Description

This command sets the bin number, the threshold and the direction that is monitored to determine if a delay metric threshold crossing event has occurred or has cleared. It requires a bin number, a rising threshold value and a direction. If the *clear-threshold* value is not specified, the traffic crossing alarm is stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and without regard to any previous window. When a raise threshold is reached, the log event is generated. Each unique threshold can only be raised once for the threshold within measurement interval. If the optional clear threshold is specified, the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to following measurement intervals. If a threshold crossing event is raised another is not raised until a measurement interval completes, and the clear threshold has not been exceeded. A clear event is raised under that condition. In general, alarms are generated when there is a state change. The thresholds configured are applied to the count in specified bin and all higher number bins.

The **no** form of this command removes thresholding for this delay metric. The complete command must be configured in order to remove the specific threshold.

Parameters

forward

Specifies the threshold is applied to the forward direction bin.

backward

Specifies the threshold is applied to the backward direction bin.

round-trip

Specifies the threshold is applied to the roundtrip direction bin.

bin-number

Specifies the number of the bin that the threshold is applied to. This bin and all higher bins are monitored to determine if the sum total results in these bins have reached or crossed the configured threshold.

Values 0 to 9

raise-threshold

Specifies the rising numerical value in the range that determines when the event is to be generated, when value reached.

Values 1 to 864000

clear-threshold

Specifies an optional numerical value in the range threshold used to indicate stateful behavior that allows the operator to configure a lower value than the rising threshold that determines when the clear event should be generated. Clear is generated when the end of measurement interval count is less than or equal to the configured value. If this option is

not configured the behavior is stateless. Zero means no results can exist in the lower bin or any higher.

Values 0 to 863999

Default Clear threshold disabled

Platforms

7705 SAR Gen 2

8.45 delay-event-exclusion

delay-event-exclusion

Syntax

delay-event-exclusion {**forward** | **backward** | **round-trip**} **lowest-bin** *bin-number*

no delay-event-exclusion {**forward** | **backward** | **round-trip**}

Context

[Tree] (config>oam-pm>bin-group>bin-type delay-event-exclusion)

Full Context

configure oam-pm bin-group bin-type delay-event-exclusion

Description

This optional command allows results from probes that map to the specified bin and higher bins to be excluded from the TCA count. The TCA count is used to determine if a threshold has been reached by the event monitoring function. Individual counters are incremented in the bin, but the counts in the specified bin and higher bins are not included in the TCA threshold computation. A **delay-event** must be configured in the same direction, and the **lowest-bin** configured as part of the **delay-event-exclusion** command must be higher than the lowest bin specified by the corresponding **delay-event** command.

The bin group allows this optional command to be added, modified, or deleted while tests are actively referencing the bin group. The bin group does not need to be shut down during **delay-event-exclusion** configuration. If the values are modified while the active tests are executing, all configured TCAs for the specified direction within the bin group enters a pending (p) state until the start of the next measurement interval. Any existing stateful TCAs that were raised are cleared without creating a log event, and no further processing for the affected TCAs occur in the active window. Depending on timing, the pending state may continue past the adjacent measurement interval until the start of the following measurement interval.

The **no** form of this command does not exclude any values from the configured TCA threshold.

Default

no delay-event-exclusion forward

no delay-event-exclusion backward

no delay-event-exclusion round-trip

Parameters

forward

Specifies the forward direction bin.

backward

Specifies the backward direction bin.

round-trip

Specifies the round-trip direction bin.

bin-number

Specifies the number of the lowest bin that the exclusion is applied to. This bin and all higher bins are excluded from the **delay-event** (TCA) count. If no bin numbers are configured, this command is ignored.

Values 1 to 9

Platforms

7705 SAR Gen 2

8.46 delay-events

delay-events

Syntax

[**no**] **delay-events**

Context

[Tree] (config>oam-pm>session>meas-intvl>event-mon delay-events)

Full Context

configure oam-pm session meas-interval event-mon delay-events

Description

This command enables the monitoring of all configured delay events. Adding this functionality starts the monitoring of the configured delay events at the start of the next measurement interval. If the function is removed using the **no** command, all monitoring of configured delay events, logging, and recording of new events for that session are suspended. Any existing events at the time of the shut down are maintained until the active measurement window in which the removal was performed has completed. The state of this monitoring function can be changed without having to shut down all the tests in the session.

The **no** form of this command disables the monitoring of all configured delay events.

Platforms

7705 SAR Gen 2

8.47 delay-metric-limit

delay-metric-limit

Syntax

delay-metric-limit *delay-metric-limit*

no delay-metric-limit

Context

[Tree] (config>router>mpls>lsp>secondary delay-metric-limit)

[Tree] (config>router>mpls>lsp>primary delay-metric-limit)

Full Context

configure router mpls lsp secondary delay-metric-limit

configure router mpls lsp primary delay-metric-limit

Description

This command configures the upper limit of the delay metric used by the local CSPF in the LSP path computation. The configured limit is used only if the configured metric type is **delay**. (The metric type is configured using the **configure router mpls lsp metric-type** or **configure router mpls lsp-template metric-type** command.)

The **no** form of this command causes the computation to select the lowest latency path if the configured metric type is **delay**.

Default

no delay-metric-limit

Parameters

delay-metric-limit

Specifies the limit, in microseconds.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

8.48 delay-normalization

delay-normalization

Syntax

[**no**] **delay-normalization**

Context

[Tree] (config>router>isis>interface delay-normalization)

Full Context

configure router isis interface delay-normalization

Description

Commands in this context configure delay normalization for the interface within the IGP instance.

When configured, the normalized delay is used by the respective TLVs within the IGP link-state packets.

The **no** form of this command removes the delay normalization configuration. When unconfigured, the measured delay is used by the respective TLVs within the IGP link-state packets.

Default

no delay-normalization

Platforms

7705 SAR Gen 2

delay-normalization

Syntax

[**no**] **delay-normalization**

Context

[Tree] (config>router>ospf>area>interface delay-normalization)

Full Context

configure router ospf area interface delay-normalization

Description

Commands in this context configure delay normalization for the interface within the IGP instance.

When configured, the normalized delay is used by the respective TLVs within the IGP link-state packets.

The **no** form of this command removes the delay normalization configuration. When unconfigured, the measured delay is used by the respective TLVs within the IGP link-state packets.

Default

no delay-normalization

Platforms

7705 SAR Gen 2

8.49 delay-on-boot

delay-on-boot

Syntax

delay-on-boot *delay*

no delay-on-boot

Context

[Tree] (config>system>grpc delay-on-boot)

Full Context

configure system grpc delay-on-boot

Description

This command configures the delay timer for gRPC connections. When the timer expires, gRPC becomes operational and connections are accepted. This delay prevents automation from managing the system while it is still converging. This delay prevents automation from managing the system while it is still converging.

The **no** form of this command specifies that connections are accepted after the system boots and gRPC becomes operational.

Default

no delay-on-boot

Parameters

delay

Specifies the delay, in seconds.

Values 1 to 3600

Platforms

7705 SAR Gen 2

delay-on-boot

Syntax

delay-on-boot *delay*

no delay-on-boot

Context

[Tree] (config>system>netconf>listen delay-on-boot)

Full Context

configure system netconf listen delay-on-boot

Description

This command configures the delay timer for NETCONF connections. When the timer expires, NETCONF becomes operational and connections are accepted. This delay prevents automation from managing the system while it is still converging.

The **no** form of this command specifies that connections are accepted after the system boots and NETCONF becomes operational.

Default

no delay-on-boot

Parameters

delay

Specifies the delay, in seconds.

Values 1 to 3600

Platforms

7705 SAR Gen 2

delay-on-boot

Syntax

delay-on-boot *delay*

no delay-on-boot

Context

[Tree] (config>system>grpc-tunnel delay-on-boot)

Full Context

configure system grpc-tunnel delay-on-boot

Description

This command configures the delay timer for gRPC tunnels. When the timer expires, gRPC tunnels become operational and connections are accepted. This delay prevents automation from trying to initiate gRPC tunnels while it is still converging.

The **no** form of this command specifies that gRPC tunnels are initiated after the system boots and gPRC becomes operational.

Default

no delay-on-boot

Parameters

delay

Specifies the delay, in seconds.

Values 1 to 3600

Platforms

7705 SAR Gen 2

delay-on-boot

Syntax

delay-on-boot *delay*

no delay-on-boot

Context

[Tree] (config>system>management-interface>remote-management delay-on-boot)

Full Context

configure system management-interface remote-management delay-on-boot

Description

This command configures the delay timer for remote management connections over gRPC. When the timer expires, remote management becomes operational and connections are accepted. This delay prevents automation from managing the system while it is still converging.

The **no** form of this command specifies that remote management connections are accepted after the system boots and gRPC becomes operational.

Default

no delay-on-boot

Parameters

delay

Specifies the delay, in seconds.

Values 1 to 3600

Platforms

7705 SAR Gen 2

delay-on-boot

Syntax

delay-on-boot *delay*

no delay-on-boot

Context

[Tree] (config>system>telemetry>persistent-subscriptions delay-on-boot)

Full Context

configure system telemetry persistent-subscriptions delay-on-boot

Description

This command configures the delay timer for gRPC telemetry persistent subscriptions. When the timer expires, gRPC telemetry persistent subscriptions become operational and connections are initiated. This delay prevents the system from trying to establish gRPC persistent subscriptions while it is still converging.

The **no** form of this command specifies that gRPC telemetry persistent subscriptions are initiated after the system boots and gRPC is operational.

Default

no delay-on-boot

Parameters

delay

Specifies the delay, in seconds.

Values 1 to 3600

Platforms

7705 SAR Gen 2

8.50 delay-tolerance-interval

delay-tolerance-interval

Syntax

delay-tolerance-interval [*delay-tolerance-interval*]

no delay-tolerance-interval

Context

[Tree] (config>router>isis>if>delay-normalization delay-tolerance-interval)

Full Context

configure router isis interface delay-normalization delay-tolerance-interval

Description

This command configures the interval, in microseconds, used by the IGP between two delay values.

The **no** form of this command reverts to the default.

Default

10 usec

Parameters

delay-tolerance-interval

Specifies the delay tolerance interval, in microseconds.

Values 1 to 10000000

Platforms

7705 SAR Gen 2

delay-tolerance-interval

Syntax

delay-tolerance-interval [*delay-tolerance-interval*]

no delay-tolerance-interval

Context

[Tree] (config>router>ospf>area>if>delay-normalization delay-tolerance-interval)

Full Context

configure router ospf area interface delay-normalization delay-tolerance-interval

Description

This command configures the interval used by the IGP between two delay values on the interface.

The **no** form of this command reverts to the default.

Default

10 usec

Parameters

delay-tolerance-interval

Specifies the delay tolerance interval, in microseconds.

Values 1 to 10000000

Platforms

7705 SAR Gen 2

8.51 delegated-prefix-length

delegated-prefix-length

Syntax

delegated-prefix-length [**minimum** *prefix-length*] [**maximum** *prefix-length*]

no delegated-prefix-length

Context

[Tree] (config>service>vprn>dhcp6>server>pool delegated-prefix-length)

[Tree] (config>router>dhcp6>server>pool delegated-prefix-length)

Full Context

configure service vprn dhcp6 local-dhcp-server pool delegated-prefix-length

configure router dhcp6 local-dhcp-server pool delegated-prefix-length

Description

This command configures the delegated prefix length that is used if the DHCPv6 client does not specify a prefix length hint.

The DHCPv6 client prefix length hint is limited by the range specified by the **minimum** and **maximum** parameters. If the hint is smaller than the minimum, the allocated prefix length is equal to the minimum length. If the hint is larger than the maximum, the allocated prefix length is equal to the maximum length.

The **no** form of this command reverts to the default.

Default

delegated-prefix-length 64 minimum 48 maximum 64

Parameters

prefix-length

Specifies the minimum or maximum allowed prefix length, in bits.

Values 48 to 127

Platforms

7705 SAR Gen 2

8.52 delete

delete

Syntax

delete *file-url* [**force**] [**no-redirect**] [**client-tls-profile** *profile*] [**proxy** *proxy-url*]

Context

[Tree] (file delete)

Full Context

file delete

Description

This command deletes the specified file.

The optional wildcard (*) can be used to delete multiple files that share a common (partial) prefix and/or (partial) suffix. When the wildcard is entered, the following prompt displays for each file that matches the wildcard:

"Delete file <filename> (y/n)?”

Parameters

file-url

Specifies the file name to delete.

Values		
	local-url	[*cflash-id*/][*file-path*] up to 200 characters, including cflash-id directory length up to 99 each

remote-url	[{ftp:// | tftp:// | http:// | https://}*login*:*pswd*@*remote-locn*/][*file-path*]
	up to 247 characters
	directory length up to 99 characters each
remote-locn	[hostname | ipv4-address | [ipv6-address]]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D
	interface - up to 32 characters, for link local addresses 255
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

force

Forces an immediate deletion of the specified file(s). The command **file delete * force** deletes all the wildcard matching files without displaying a user prompt message. This command also automatically accepts HTTP redirects unless overridden by the **no-redirect** parameter.

profile

Specifies the TLS client profile configured under **config**>**system**>**security**>**tls**>**client-tls-profile** to use.

proxy-url

Specifies the URL of an HTTP proxy. For example, http://proxy.mydomain.com:8000. This URL must be an HTTP URL and not an HTTPS URL.

no-redirect

Specifies to automatically refuse any HTTP redirects without prompting the user.

Platforms

7705 SAR Gen 2

delete

Syntax

delete [*line*]

Context

[Tree] (candidate delete)

Full Context

candidate delete

Description

This command deletes the selected CLI node (which includes all sub-branches). The deleted lines are also copied into a temporary buffer that can be used for a subsequent insert.

Parameters

line

Indicates which line to delete.

Values line, offset, **first**, **edit-point**, **last**

line	absolute line number
offset	relative line number to current edit point. Prefixed with '+' or '-'
first	keyword - first line
edit-point	keyword - current edit point
last	keyword - last line that is not 'exit'

Platforms

7705 SAR Gen 2

delete

Syntax

delete [{*checkpoint-id* | **rescue** | **latest-rb**}]

Context

[Tree] (admin>rollback delete)

Full Context

admin rollback delete

Description

This command deletes a rollback checkpoint and causes the suffixes to be adjusted (decremented) for all checkpoints older that the one that was deleted (to close the hole in the list of checkpoint files and create room to create another checkpoint).

If **config redundancy rollback-sync** is enabled, a rollback delete will also delete the equivalent checkpoint on the standby CF and shuffle the suffixes on the standby CF.

It is not advised to manually delete a rollback checkpoint (for example, using a **file delete** command). If a rollback checkpoint file is manually deleted without using the **admin rollback delete** command then the suffixes of the checkpoint files are not shuffled, nor is the equivalent checkpoint file deleted from the standby CF. This manual deletion creates a hole in the checkpoint file list until enough new checkpoints have been created to roll the hole off the end of the list.

Parameters

checkpoint-id

An ID indicating a specific rollback checkpoint. A checkpoint-id of 1 indicates the rollback checkpoint file (at the configured rollback location) with *.rb.1 as the suffix, 2 for file *.rb.2, and so on.

Values 1 to 9

latest-rb

Specifies the most recently created rollback checkpoint (corresponds to the file-url.rb rollback checkpoint file).

rescue

Deletes the rescue checkpoint. No checkpoint suffix numbers are changed.

Platforms

7705 SAR Gen 2

8.53 delete-config

delete-config

Syntax

[**no**] **delete-config**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization delete-config)

Full Context

configure system security profile netconf base-op-authorization delete-config

Description

This command enables the NETCONF <delete-config> RPC.

The **no** form of this command disables the RPC.

Default

no delete-config

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

8.54 delivery-service

delivery-service

Syntax

delivery-service *service-id*

delivery-service name *service-name*

no delivery-service

Context

[Tree] (config>service>vprn>if>sap>ip-tunnel delivery-service)

[Tree] (config>service>ies>if>sap>ip-tunnel delivery-service)

Full Context

configure service vprn interface sap ip-tunnel delivery-service

configure service ies interface sap ip-tunnel delivery-service

Description

This command sets the delivery service for GRE encapsulated packets associated with a particular GRE tunnel. This is the IES or VPRN service where the GRE encapsulated packets are injected and terminated. The delivery service may be the same service that owns the private tunnel SAP associated with the GRE tunnel. The GRE tunnel does not come up until a valid delivery service is configured.

The **no** form of this command deletes the delivery-service from the GRE tunnel configuration.

Parameters

service-id

Identifies the service used to originate and terminate the GRE encapsulated packets belonging to the GRE tunnel.

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **delivery-service name** *service-name* variant can be used in all configuration modes.

Values {*id* | *svc-name*}

id:	1 to 2147483647
svc-name:	up to 64 characters (*svc-name* is an alias for input only. The *svc-name* gets replaced with an id automatically by SR OS in the configuration).

service-name

Identifies the service used to originate and terminate the GRE encapsulated packets belonging to the GRE tunnel.

Values 1 to 64 characters

Platforms

7705 SAR Gen 2

8.55 delta-in-use-limit

delta-in-use-limit

Syntax

delta-in-use-limit *limit*

no delta-in-use-limit

Context

[Tree] (config>vrrp>policy delta-in-use-limit)

Full Context

configure vrrp policy delta-in-use-limit

Description

This command sets a lower limit on the virtual router in-use priority that can be derived from the delta priority control events.

Each *vrrp-priority-id* places limits on the delta priority control events to define the in-use priority of the virtual router instance. Setting this limit prevents the sum of the delta priority events from lowering the in-use priority value of the associated virtual router instances below the configured value.

The limit has no effect on explicit priority control events. Explicit priority control events are controlled by setting the in-use priority to any value between 1 and 254.

Only non-owner virtual router instances can be associated with VRRP priority control policies and their priority control events.

Once the total sum of all delta events is calculated and subtracted from the base **priority** of the virtual router instance, the result is compared to the **delta-in-use-limit** value. If the result is less than the limit, the

delta-in-use-limit value is used as the virtual router in-use priority value. If an explicit priority control event overrides the delta priority control events, the **delta-in-use-limit** has no effect.

Setting the limit to a higher value than the default of 1 limits the effect of the delta priority control events on the virtual router instance base **priority** value. This allows for multiple priority control events while minimizing the overall effect on the in-use priority.

Changing the *in-use-priority-limit* causes an immediate re-evaluation of the in-use priority values for all virtual router instances associated with this *vrrp-policy-id* based on the current sum of all active delta control policy events.

The **no** form of the command reverts to the default value.

Default

delta-in-use-limit 1 — Specifies the lower limit of 1 for the in-use priority, as modified, by delta priority control events.

Parameters

limit

Specifies the lower limit of the in-use priority base, as modified by priority control policies. The *in-use-priority-limit* has the same range as the non-owner virtual router instance base-priority parameter. If the result of the total delta priority control events minus the virtual router instances base-priority, is less than the *in-use-priority-limit*, the *in-use-priority-limit* value is used as the virtual router instances in-use priority value.

Setting the *in-use-priority-limit* to a value equal to or larger than the virtual router instance *base-priority* prevents the delta priority control events from having any effect on the virtual router instance in-use priority value.

Values 1 to 254

Platforms

7705 SAR Gen 2

8.56 depleted-event

depleted-event

Syntax

[**no**] **depleted-event**

Context

[Tree] (config>router>dhcp6>server>pool>prefix>thresholds>minimum-free depleted-event)

[Tree] (config>service>vprn>dhcp6>server>pool>prefix>thresholds>minimum-free depleted-event)

Full Context

configure router dhcp6 local-dhcp-server pool prefix thresholds minimum-free depleted-event

configure service vprn dhcp6 local-dhcp-server pool prefix thresholds minimum-free depleted-event

Description

This command enables the system to send out a warning when the prefix with a configured length is no longer available in the provisioned prefix.

For example, if the prefix 2001:0:0:ffe0::/50 is created at the pool level using the **pd** and **wan-host** parameters, and the threshold for the prefix length is 64, configuring this command enables the system to send out a warning when there is no available /64 that can be allocated out of 2001:0:0:ffe0::/50.

The **no** form of this command disables the warnings.

Platforms

7705 SAR Gen 2

depleted-event

Syntax

[**no**] **depleted-event**

Context

[Tree] (config>service>vprn>dhcp6>server>pool>thresholds>minimum-free depleted-event)

[Tree] (config>router>dhcp6>server>pool>thresholds>minimum-free depleted-event)

Full Context

configure service vprn dhcp6 local-dhcp-server pool thresholds minimum-free depleted-event

configure router dhcp6 local-dhcp-server pool thresholds minimum-free depleted-event

Description

This command enables the system to send out warnings when the prefix with the configured length is no longer available in the pool.

The **no** form of this command disables the warnings.

Platforms

7705 SAR Gen 2

8.57 description

description

Syntax

description *tiny-description-string*

no description

Context

[Tree] (config>ipsec>static-sa description)

Full Context

configure ipsec static-sa description

Description

This command configures a text description that is stored in the configuration file. The text string is associated with a configuration context to identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Parameters

tiny-description-string

Specifies the description character string. Allowed values are any string, up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

description

Syntax

description *short-description-string*

no description

Context

[Tree] (config>card>fp>ingress>access>queue-group description)

[Tree] (config>filter>dhcp-filter description)

[Tree] (config>filter>ipv6-filter>entry description)

[Tree] (config>qos>match-list>ipv6-prefix-list description)

[Tree] (config>subscr-mgmt>loc-user-db description)
[Tree] (config>router>dhcp>server description)
[Tree] (config>filter>redirect-policy description)
[Tree] (config>system>persistence>dhcp-server description)
[Tree] (config>system>script-control>script description)
[Tree] (config>filter>redirect-policy>destination description)
[Tree] (config>router>rip>group description)
[Tree] (config>filter>ip-exception>entry description)
[Tree] (config>service>ies>if>sap>ip-tunnel description)
[Tree] (config>qos>network-queue description)
[Tree] (config>service>vprn>bgp>group description)
[Tree] (config>service>vprn>spoke-sdp description)
[Tree] (config>filter>log description)
[Tree] (config>router>static-route-entry>indirect description)
[Tree] (config>router>dhcp6>server>pool description)
[Tree] (config>service>vprn>if>sap description)
[Tree] (config>card>fp>ingress>network>queue-group description)
[Tree] (config>service>epipe description)
[Tree] (config>system>persistence>nat-fwd description)
[Tree] (config>service>vprn>bgp description)
[Tree] (config>router>ripng description)
[Tree] (config>service>vprn>static-route-entry>black-hole description)
[Tree] (config>router>origin-validation>rpki-session description)
[Tree] (config>ipsec>tnl-temp description)
[Tree] (config>port>ethernet>access>egr>qgrp description)
[Tree] (config>filter>ip-exception description)
[Tree] (config>service>vprn>static-route-entry>ipsec-tunnel description)
[Tree] (config>service>vprn>bgp>group>neighbor description)
[Tree] (config>service>vprn>if>sap>ipsec-tunnel description)
[Tree] (config>service>ies>if>dhcp description)
[Tree] (config>router>bgp description)
[Tree] (config>port-xc>pxc description)
[Tree] (config>ipsec>trans-mode-prof description)
[Tree] (config>ipsec>client-db description)
[Tree] (config>router>fad>flex-algo description)
[Tree] (config>service>vprn>rip>group description)

[Tree] (config>port>ethernet>network>egr>qgrp description)
[Tree] (config>system>cron>sched description)
[Tree] (config>service>epipe>spoke-sdp description)
[Tree] (config>subscr-mgmt>rip-plcy description)
[Tree] (config>router>ripng>group description)
[Tree] (config>system>persistence>ancp description)
[Tree] (config>router>dhcp6>server description)
[Tree] (config>router>bgp>group description)
[Tree] (config>router>ripng>group>neighbor description)
[Tree] (config>router>if>dhcp description)
[Tree] (config>qos>policer-control-policy description)
[Tree] (config>service>epipe>endpoint description)
[Tree] (config>service>vprn>if>sap>ip-tunnel description)
[Tree] (config>service>vprn>ripng>group>neighbor description)
[Tree] (config>filter>ipv6-exception>entry description)
[Tree] (config>filter>match-list>protocol-list description)
[Tree] (config>router>rip description)
[Tree] (config>system>management-interface>remote-management>manager description)
[Tree] (config>service>vprn>static-route-entry>indirect description)
[Tree] (config>router>rip>group>neighbor description)
[Tree] (config>qos>policer-control-policy>tier>arbiter description)
[Tree] (config>service>vprn>static-route-entry>next-hop description)
[Tree] (config>filter>match-list>port-list description)
[Tree] (config>filter>ip-filter>entry description)
[Tree] (config>qos>match-list>ip-prefix-list description)
[Tree] (config>service>vprn>static-route-entry>grt description)
[Tree] (config>system>grpc-tunnel>tunnel description)
[Tree] (config>router>dhcp>server>pool description)
[Tree] (config>isa>tunnel-group description)
[Tree] (config>service>vprn>if>dhcp description)
[Tree] (config>redundancy>multi-chassis>peer description)
[Tree] (config>isa>tunnel-mem-pool description)
[Tree] (config>service>vprn>ripng>group description)
[Tree] (config>ipsec>ike-policy description)
[Tree] (config>service>ies description)
[Tree] (config>service>vprn>ripng description)

[Tree] (config>router>route-next-hop-policy>template description)

[Tree] (config>filter>match-list>ip-prefix-list description)

[Tree] (config>service>vprn>rip>group>neighbor description)

[Tree] (config>filter>ipv6-exception description)

[Tree] (config>router>bgp>group>neighbor description)

[Tree] (config>filter>ipv6-filter description)

[Tree] (config>router>static-route-entry>next-hop description)

[Tree] (config>service>ies>if>spoke-sdp description)

[Tree] (config>port>ethernet>access>ing>qgrp description)

[Tree] (config>system>grpc-tunnel>destination-group description)

[Tree] (config>router>network-domains>network-domain description)

[Tree] (config>service>ies>if>ipv6>dhcp6-relay description)

[Tree] (config>service>vprn>rip description)

[Tree] (config>filter>match-list>ipv6-prefix-list description)

[Tree] (config>filter>ip-filter description)

[Tree] (config>service>vpls>endpoint description)

[Tree] (config>router>static-route-entry>black-hole description)

Full Context

configure card fp ingress access queue-group description

configure filter dhcp-filter description

configure filter ipv6-filter entry description

configure qos match-list ipv6-prefix-list description

configure subscriber-mgmt local-user-db description

configure router dhcp local-dhcp-server description

configure filter redirect-policy description

configure system persistence dhcp-server description

configure system script-control script description

configure filter redirect-policy destination description

configure router rip group description

configure filter ip-exception entry description

configure service ies interface sap ip-tunnel description

configure qos network-queue description

configure service vprn bgp group description

configure service vprn spoke-sdp description

configure filter log description

configure router static-route-entry indirect description

configure router dhcp6 local-dhcp-server pool description
configure service vprn interface sap description
configure card fp ingress network queue-group description
configure service epipe description
configure system persistence nat-port-forwarding description
configure service vprn bgp description
configure router ripng description
configure service vprn static-route-entry black-hole description
configure router origin-validation rpki-session description
configure ipsec tunnel-template description
configure port ethernet access egress queue-group description
configure filter ip-exception description
configure service vprn static-route-entry ipsec-tunnel description
configure service vprn bgp group neighbor description
configure service vprn interface sap ipsec-tunnel description
configure service ies interface dhcp description
configure router bgp description
configure port-xc pxc description
configure ipsec ipsec-transport-mode-profile description
configure ipsec client-db description
configure router flexible-algorithm-definitions flex-algo description
configure service vprn rip group description
configure port ethernet network egress queue-group description
configure system cron schedule description
configure service epipe spoke-sdp description
configure subscriber-mgmt rip-policy description
configure router ripng group description
configure system persistence ancp description
configure router dhcp6 local-dhcp-server description
configure router bgp group description
configure router ripng group neighbor description
configure router interface dhcp description
configure qos policer-control-policy description
configure service epipe endpoint description
configure service vprn interface sap ip-tunnel description
configure service vprn ripng group neighbor description

configure filter ipv6-exception entry description
configure filter match-list protocol-list description
configure router rip description
configure system management-interface remote-management manager description
configure service vprn static-route-entry indirect description
configure router rip group neighbor description
configure qos policer-control-policy tier arbiter description
configure service vprn static-route-entry next-hop description
configure filter match-list port-list description
configure filter ip-filter entry description
configure qos match-list ip-prefix-list description
configure service vprn static-route-entry grt description
configure system grpc-tunnel tunnel description
configure router dhcp local-dhcp-server pool description
configure isa tunnel-group description
configure service vprn interface dhcp description
configure redundancy multi-chassis peer description
configure isa tunnel-member-pool description
configure service vprn ripng group description
configure ipsec ike-policy description
configure service ies description
configure service vprn ripng description
configure router route-next-hop-policy template description
configure filter match-list ip-prefix-list description
configure service vprn rip group neighbor description
configure filter ipv6-exception description
configure router bgp group neighbor description
configure filter ipv6-filter description
configure router static-route-entry next-hop description
configure service ies interface spoke-sdp description
configure port ethernet access ingress queue-group description
configure system grpc-tunnel destination-group description
configure router network-domains network-domain description
configure service ies interface ipv6 dhcp6-relay description
configure service vprn rip description
configure filter match-list ipv6-prefix-list description

configure filter ip-filter description

configure service vpls endpoint description

configure router static-route-entry black-hole description

Description

This command configures a text description that is stored in the configuration file. The text string is associated with a configuration context to identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default

no description

Parameters

short-description-string

Specifies the description entered as a character string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

description

Syntax

description *short-description-string*

no description

Context

[Tree] (config>saa>test description)

[Tree] (config>qos>network description)

[Tree] (config>oam-pm>session description)

[Tree] (config>mirror>mirror-dest description)

[Tree] (config>aaa>radius-srv-plcy description)

[Tree] (config>service>vpls>mesh-sdp description)

[Tree] (config>service>vprn>nat>outside>pool>address-range description)

[Tree] (config>service>vpls>sap description)

[Tree] (config>qos>network>egress>ip-criteria>entry description)

[Tree] (config>service>vprn>twamp-light>reflector description)

[Tree] (config>service>vprn>if>sap>ip-tunnel description)

[Tree] (config>qos>network>ingress>ip-criteria>entry description)

[Tree] (config>service>vpls>spoke-sdp description)
[Tree] (config>router>policy-options>policy-statement description)
[Tree] (config>qos>sap-egress>ip-criteria>entry description)
[Tree] (config>qos>network>egress>ipv6-criteria>entry description)
[Tree] (config>qos>sap-ingress>policer description)
[Tree] (config>qos>network>ingress>ipv6-criteria>entry description)
[Tree] (config>service>vprn>radius-server>server description)
[Tree] (config>router>radius-server>server description)
[Tree] (config>service>vprn>twamp-light>reflector>prefix description)
[Tree] (config>test-oam>twamp>server>prefix description)
[Tree] (config>service>vprn>ip-mirror-interface>spoke-sdp description)
[Tree] (config>router>twamp-light>reflector description)
[Tree] (cfg>qos>qgrps>ing>qgrp>policer description)
[Tree] (config>qos>sap-egress description)
[Tree] (config>service>vprn description)
[Tree] (config>qos>sap-ingress description)
[Tree] (config>vrrp>policy description)
[Tree] (config>service>nat>nat-policy description)
[Tree] (config>oam-pm>bin-group description)
[Tree] (config>router>policy-options>policy-statement>entry description)
[Tree] (config>qos>sap-ingress>mac-criteria>entry description)
[Tree] (config>service>vprn>nat>outside>pool description)
[Tree] (config>mirror>mirror-dest>endpoint description)
[Tree] (config>router>nat>outside>pool description)
[Tree] (config>qos>sap-ingress>ipv6-criteria>entry description)
[Tree] (config>qos>sap-ingress>ip-criteria>entry description)
[Tree] (cfg>qos>qgrps>egr>qgrp>policer description)
[Tree] (config>qos>sap-egress>policer description)
[Tree] (config>router>twamp-light>reflector>prefix description)
[Tree] (cfg>qos>qgrps>ing>qgrp description)
[Tree] (config>isa>nat-group description)
[Tree] (config>service>mac-list description)
[Tree] (config>router>nat>outside>pool>address-range description)
[Tree] (cfg>qos>qgrps>egr>qgrp description)
[Tree] (config>service>vpls>sap>dhcp description)
[Tree] (config>service>vpls description)

[Tree] (config>qos>sap-egress>ipv6-criteria>entry description)

[Tree] (config>macsec>connectivity-association description)

[Tree] (config>service>vpls>split-horizon-group description)

Full Context

configure saa test description

configure qos network description

configure oam-pm session description

configure mirror mirror-dest description

configure aaa radius-server-policy description

configure service vpls mesh-sdp description

configure service vprn nat outside pool address-range description

configure service vpls sap description

configure qos network egress ip-criteria entry description

configure service vprn twamp-light reflector description

configure service vprn interface sap ip-tunnel description

configure qos network ingress ip-criteria entry description

configure service vpls spoke-sdp description

configure router policy-options policy-statement description

configure qos sap-egress ip-criteria entry description

configure qos network egress ipv6-criteria entry description

configure qos sap-ingress policer description

configure qos network ingress ipv6-criteria entry description

configure service vprn radius-server server description

configure router radius-server server description

configure service vprn twamp-light reflector prefix description

configure test-oam twamp server prefix description

configure service vprn ip-mirror-interface spoke-sdp description

configure router twamp-light reflector description

configure qos queue-group-templates ingress queue-group policer description

configure qos sap-egress description

configure service vprn description

configure qos sap-ingress description

configure vrrp policy description

configure service nat nat-policy description

configure oam-pm bin-group description

configure router policy-options policy-statement entry description

configure qos sap-ingress mac-criteria entry description

configure service vprn nat outside pool description

configure mirror mirror-dest endpoint description

configure router nat outside pool description

configure qos sap-ingress ipv6-criteria entry description

configure qos sap-ingress ip-criteria entry description

configure qos queue-group-templates egress queue-group policer description

configure qos sap-egress policer description

configure router twamp-light reflector prefix description

configure qos queue-group-templates ingress queue-group description

configure isa nat-group description

configure service mac-list description

configure router nat outside pool address-range description

configure qos queue-group-templates egress queue-group description

configure service vpls sap dhcp description

configure service vpls description

configure qos sap-egress ipv6-criteria entry description

configure macsec connectivity-association description

configure service vpls split-horizon-group description

Description

This command configures a text description that is stored in the configuration file. The text string is associated with a configuration context to identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Parameters

short-description-string

Specifies the description character string. Allowed values are any string, up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

description

Syntax

description *short-description-string*

no description

Context

[Tree] (config>qos>scheduler-policy>tier>scheduler description)

[Tree] (config>log>filter>entry description)

[Tree] (config>system>security>keychain description)

[Tree] (config>log>filter description)

[Tree] (config>service>cust description)

[Tree] (config>system>security>pki>ca-profile description)

[Tree] (config>system>telemetry>destination-group description)

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry description)

[Tree] (config>system>security>user>public-keys>ecdsa>ecdsa-key description)

[Tree] (config>qos>port-scheduler-policy description)

[Tree] (config>qos>scheduler-policy description)

[Tree] (config>service>pw-template>split-horizon-group description)

[Tree] (config>log>event-trigger>event>trigger-entry description)

[Tree] (config>log>accounting-policy description)

[Tree] (config>log>snmp-trap-group description)

[Tree] (config>log>syslog description)

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter>entry description)

[Tree] (config>connection-profile-vlan description)

[Tree] (config>service>sdp description)

[Tree] (config>system>security>mgmt-access-filter>ip-filter>entry description)

[Tree] (config>log>event-handling>handler description)

[Tree] (config>grp-encryp>encryp-keygrp description)

[Tree] (config>log>log-id description)

[Tree] (config>service>cust>multi-service-site description)

[Tree] (config>log>event-trigger>event description)

[Tree] (config>log>file-id description)

[Tree] (config>system>security>dist-cpu-protection>policy description)

[Tree] (config>system>telemetry>persistent-subscriptions>subscription description)

[Tree] (config>log>event-handling>handler>action-list>entry description)

[Tree] (config>system>telemetry>sensor-groups>sensor-group description)

[Tree] (config>system>security>user>public-keys>rsa>rsa-key description)

Full Context

configure qos scheduler-policy tier scheduler description

configure log filter entry description

configure system security keychain description

configure log filter description
configure service customer description
configure system security pki ca-profile description
configure system telemetry destination-group description
configure system security management-access-filter mac-filter entry description
configure system security user public-keys ecdsa ecdsa-key description
configure qos port-scheduler-policy description
configure qos scheduler-policy description
configure service pw-template split-horizon-group description
configure log event-trigger event trigger-entry description
configure log accounting-policy description
configure log snmp-trap-group description
configure log syslog description
configure system security management-access-filter ipv6-filter entry description
configure connection-profile-vlan description
configure service sdp description
configure system security management-access-filter ip-filter entry description
configure log event-handling handler description
configure group-encryption encryption-keygroup description
configure log log-id description
configure service customer multi-service-site description
configure log event-trigger event description
configure log file-id description
configure system security dist-cpu-protection policy description
configure system telemetry persistent-subscriptions subscription description
configure log event-handling handler action-list entry description
configure system telemetry sensor-groups sensor-group description
configure system security user public-keys rsa rsa-key description

Description

This command configures a text description that is stored in the configuration file. The text string is associated with a configuration context to identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default

No description is associated with the configuration context.

Parameters

short-description-string

Specifies the description character string. Allowed values are any string, up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

description

Syntax

description *medium-description-string*

no description

Context

[Tree] (config>service>epipe>sap description)

[Tree] (config>service>ies>interface>sap description)

Full Context

configure service epipe sap description

configure service ies interface sap description

Description

This command configures a text description that is stored in the configuration file. The text string is associated with a configuration context to identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default

no description

Parameters

medium-description-string

Specifies the description character string. Allowed values are any string, up to 160 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

description

Syntax

description *long-description-string*

no description

Context

[Tree] (config>service>vprn>isis>link-group description)

[Tree] (config>router>isis>link-group description)

Full Context

configure service vprn isis link-group description

configure router isis link-group description

Description

This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default

no description

Parameters

long-description-string

Specifies the description character string. Allowed values are any string, up to 255-256 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

description

Syntax

description *long-description-string*

no description

Context

[Tree] (config>port description)

[Tree] (config>router>if description)

[Tree] (config>service>vpls>sap>dhcp6 description)

[Tree] (config>service>ies>interface description)

[Tree] (config>service>vpls>interface description)

[Tree] (config>service>vprn>if description)

[Tree] (config>service>vprn>nw-if description)

[Tree] (config>service>vprn>ip-mirror-interface description)

[Tree] (config>lag description)

Full Context

configure port description

configure router interface description

configure service vpls sap dhcp6 description

configure service ies interface description

configure service vpls interface description

configure service vprn interface description

configure service vprn network-interface description

configure service vprn ip-mirror-interface description

configure lag description

Description

This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Parameters

long-description-string

Specifies the description character string. Allowed values are any string up to 255-256 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

8.58 designated-role

designated-role

Syntax

designated-role {**standby** | **active**}

no designated-role

Context

[Tree] (config>redundancy>multi-chassis>ipsec-domain designated-role)

Full Context

configure redundancy multi-chassis ipsec-domain designated-role

Description

This command sets the designated role for the tunnel group in the IPsec domain.

The **no** form of this command reverts to the default value.

Default

designated-role standby

Parameters

standby

Sets the designated role to standby.

active

Sets the designated role to active.

Platforms

7705 SAR Gen 2

8.59 dest-ip

dest-ip

Syntax

[**no**] **dest-ip** *ip-address*

Context

[Tree] (config>service>ies>if>sap>ip-tunnel dest-ip)

[Tree] (config>service>vprn>if>sap>ip-tunnel dest-ip)

Full Context

configure service ies interface sap ip-tunnel dest-ip

configure service vprn interface sap ip-tunnel dest-ip

Description

This command configures a private IPv4 or IPv6 address of the remote tunnel endpoint. A tunnel can have up to 16 **dest-ip** commands. At least one **dest-ip** address is required in the configuration of a tunnel. A tunnel does not come up operationally unless all **dest-ip** addresses are reachable (part of a local subnet).

Unnumbered interfaces are not supported.

The **no** form of this command deletes the destination IP of the tunnel.

Parameters

ip-address

Specifies the private IPv4 or IPv6 address of the remote IP tunnel endpoint. If this remote IP address is not within the subnet of the IP interface associated with the tunnel then the tunnel will not come up.

Values

<ip-address>	ipv4-address	a.b.c.d
	ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x - [0 to FFFF]H
		d - [0 to 255]D

Platforms

7705 SAR Gen 2

8.60 dest-mac

dest-mac

Syntax

dest-mac {**nearest-bridge** | **nearest-non-tpmr** | **nearest-customer**}

Context

[Tree] (config>port>ethernet>lldp dest-mac)

Full Context

configure port ethernet lldp dest-mac

Description

This command configures destination MAC address parameters.

Default

dest-mac nearest-bridge

Parameters

nearest-bridge

Specifies to use the nearest bridge.

nearest-non-tpmr

Specifies to use the nearest non-Two-Port MAC Relay (TPMR).

nearest-customer

Specifies to use the nearest customer.

Platforms

7705 SAR Gen 2

dest-mac

Syntax

dest-mac {**nearest-bridge** | **nearest-non-tpmr** | **nearest-customer**}

Context

[Tree] (config>lag>lldp-member-template dest-mac)

Full Context

configure lag lldp-member-template dest-mac

Description

This command configures the destination MAC address parameters.

Default

dest-mac nearest-bridge

Parameters

nearest-bridge

Keyword to specify that the nearest bridge should be used.

nearest-non-tpmr

Keyword to specify that the nearest non-Two-Port MAC Relay (TPMR) should be used.

nearest-customer

Keyword to specify that the nearest customer should be used.

Platforms

7705 SAR Gen 2

8.61 dest-mac-address

dest-mac-address

Syntax

dest-mac-address *mac-address* [**create**]

no dest-mac-address *mac-address*

Context

[Tree] (config>macsec>mac-policy dest-mac-address)

Full Context

configure macsec mac-policy dest-mac-address

Description

This command specifies the destination MAC address.

The **no** form of this command removes the MAC address.

Parameters

mac-address

Specifies the value of the MAC address policy.

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

create

Mandatory to create the configuration.

Platforms

7705 SAR Gen 2

8.62 dest-mac-rewrite

dest-mac-rewrite

Syntax

dest-mac-rewrite *ieee-address*

no dest-mac-rewrite

Context

[Tree] (config>service>vpls>sap>egress dest-mac-rewrite)

Full Context

configure service vpls sap egress dest-mac-rewrite

Description

This commands enables the overwriting of a destination MAC address to an operator-configured value for all unicast packets egressing the specified SAP. The command is intended to be deployed with L2 PBF SAP redirect when a remote end of the SAP interface is an L3 interface with a MAC address different from the MAC address of the non-PBF-ed L3 interface. See Filter Policy in the *7705 SAR Gen 2 Router Configuration Guide* for more details.

The **no** form disables the option.

Default

no dest-mac-rewrite

Parameters

ieee-address

Specifies the MAC address

Values 1xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

Cannot be all zeros

Platforms

7705 SAR Gen 2

8.63 dest-udp-port

dest-udp-port

Syntax

dest-udp-port *udp-port-number*

no dest-udp-port

Context

[Tree] (config>oam-pm>session>ip dest-udp-port)

Full Context

configure oam-pm session ip dest-udp-port

Description

This command defines the destination UDP port on outbound TWAMP Light packets sent from the session controller. The destination UDP port must match the UDP port value configured on the TWAMP Light reflector that is responding to this specific TWAMP Light test.

The **no** form of this command removes the destination UDP port setting.

Parameters

udp-port-number

Specifies the UDP source port.

Values 1 to 65535

Platforms

7705 SAR Gen 2

8.64 destination

destination

Syntax

destination *ip-address*

no destination

Context

[Tree] (config>oam-pm>session>ip destination)

Full Context

configure oam-pm session ip destination

Description

This command defines the destination IP address that is assigned to the TWAMP Light packets. The destination address must be included in the prefix list on the session reflector within the configured context in order to allow the reflector to process the inbound TWAMP Light packets.

The **no** form of this command removes the destination parameters.

Parameters

ip-address

Specifies the IP address of the IP peer to which the packet is directed.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

Platforms

7705 SAR Gen 2

destination

Syntax

destination *ip-address* [**create**]

no destination *ip-address*

Context

[Tree] (config>filter>redirect-policy destination)

Full Context

configure filter redirect-policy destination

Description

This command defines a destination in a redirect policy. More than one destination can be configured. Whether a destination IPv4/IPv6 address will receive redirected packets depends on the effective priority value after evaluation.

The most preferred destination is programmed in hardware as action forward next-hop. If all destinations are down (as determined by the supported tests), action forward is programmed in hardware. All

destinations within a given policy must be either IPv4 or (exclusive) IPv6. The redirect policy with IPv4 destinations configured can only be used by IPv4 filter policies. The redirect policy with IPv6 destinations configured can only be used by IPv6 filter policies.

Default

no destination

Parameters

ip-address

Specifies the IPv4 address (in dotted decimal notation) or IPv6 address to send the redirected traffic to.

Values IPv4 address: ip-address: a.b.c.d

IPv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0..FFFF]H

d: [0..255]D

Platforms

7705 SAR Gen 2

destination

Syntax

destination memory *num-entries*

destination syslog *syslog-id*

no destination

Context

[Tree] (config>filter>log destination)

Full Context

configure filter log destination

Description

This command configures the destination for filter log entries for the filter log ID.

Filter logs can be sent to either memory (**memory**) or to an existing Syslog server definition (**syslog**).

If the filter log destination is **memory**, the maximum number of entries in the log must be specified.

The **no** form of the command deletes the filter log association.

Default

destination memory 1000

Parameters

memory *num-entries*

Specifies the destination of the filter log ID is a memory log. The *num-entries* value is the maximum number of entries in the filter log expressed as a decimal integer.

Values 10 to 50000

syslog *syslog-id*

Specifies the destination of the filter log ID is a Syslog server. The *syslog-id* parameter is the number of the Syslog server definition.

Values 1 to 10

Platforms

7705 SAR Gen 2

destination

Syntax

destination {*ip-address* | *fqdn*} **port** *port* [**create**]

no destination {*ip-address* | *fqdn*} **port** *port*

Context

[Tree] (config>system>grpc-tunnel>destination-group destination)

[Tree] (config>system>telemetry>destination-group destination)

Full Context

configure system grpc-tunnel destination-group destination

configure system telemetry destination-group destination

Description

This command configures a destination IP address and port for a specific destination within a destination group. Up to two destinations can be defined within a destination group. Each destination is an IPv4 address, an IPv6 address, or the Fully Qualified Domain Name (FQDN).

The **no** form of this command removes the destination from the destination group.

Parameters

ip-address

Specifies the IPv4 address (in dotted decimal notation) or IPv6 address.

Values IPv4 address: ip-address: a.b.c.d

IPv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0..FFFF]H

d: [0..255]D

fqdn

Specifies the FQDN.

port

Specifies the TCP destination port number.

Values 1 to 65535

create

Keyword used to create a destination.

Platforms

7705 SAR Gen 2

8.65 destination-group

destination-group

Syntax

destination-group *name* [**create**]

no destination-group *name*

Context

[Tree] (config>system>grpc-tunnel destination-group)

[Tree] (config>system>telemetry destination-group)

Full Context

configure system grpc-tunnel destination-group

configure system telemetry destination-group

Description

Commands in this context configure commands for destination groups.

The **no** form of this command removes the destination group name.

Parameters

name

Specifies the destination group name, up to 32 characters.

create

Keyword used to create a destination group.

Platforms

7705 SAR Gen 2

destination-group

Syntax

destination-group *name*

no destination-group

Context

[Tree] (config>system>grpc-tunnel>tunnel destination-group)

Full Context

configure system grpc-tunnel tunnel destination-group

Description

This command assigns the specified destination group to a gRPC tunnel.

The **no** form of this command removes the specified destionation group from the gRPC tunnel.

Default

no destination-group

Parameters

name

Specifies the destination group name, up to 32 characters

Platforms

7705 SAR Gen 2

destination-group

Syntax

destination-group *name*

no destination-group

Context

[Tree] (config>system>telemetry>persistent-subscriptions>subscription destination-group)

Full Context

configure system telemetry persistent-subscriptions subscription destination-group

Description

This command assigns an existing destination group to the specified persistent subscription. The assigned **destination-group** must already exist before the configured persistent subscription can be activated.

The **no** form of this command removes the destination group name from the persistent subscription.

Parameters

name

Specifies the destination group name, up to 32 characters.

Platforms

7705 SAR Gen 2

8.66 destination-prefix

destination-prefix

Syntax

destination-prefix *ip-prefix/length* [**nat-policy** *nat-policy-name*]

no destination-prefix *ip-prefix/length*

Context

[Tree] (config>router>nat>inside destination-prefix)

[Tree] (config>service>vprn>nat>inside destination-prefix)

Full Context

configure router nat inside destination-prefix

configure service vprn nat inside destination-prefix

Description

This command configures a destination prefix. An (internal) static route will be created for this prefix. All traffic that hits this route will be subject to NAT. The system will not allow a destination-prefix to be configured if the configured nat-policy refers to an IP pool that resides in the same service (as this would result in a routing loop).

Parameters

ip-prefix

Specifies the IP prefix; host bits must be zero (0).

Values a.b.c.d

length

Specifies the prefix length.

Values 0 to 32

nat-policy-name

Specifies the NAT policy name, up to 32 characters.

Platforms

7705 SAR Gen 2

8.67 detail-level

detail-level

Syntax

detail-level {**low** | **medium** | **high**}

no detail-level

Context

[Tree] (debug>router>ip>dhcp6 detail-level)

[Tree] (debug>router>ip>dhcp detail-level)

[Tree] (debug>router>local-dhcp-server detail-level)

Full Context

debug router ip dhcp6 detail-level

debug router ip dhcp detail-level

debug router local-dhcp-server detail-level

Description

This command debugs the DHCP tracing detail level.

Parameters

low

Displays a low detail level for DHCP debugging.

medium

Displays a medium detail level for DHCP debugging.

high

Displays a high detail level for DHCP debugging.

Platforms

7705 SAR Gen 2

detail-level

Syntax

detail-level {**low** | **medium** | **high**}

no detail-level

Context

[Tree] (debug>router>radius detail-level)

Full Context

debug router radius detail-level

Description

This command specifies the output detail level of command **debug router radius**.

Default

detail-level medium

Parameters

low

Specifies that the output includes packet type, server address, length, radius-server-policy name.

medium

Specifies all output in low level including the RADIUS attributes in the packet.

high

Specifies all output in medium level including the hex packet dump.

Platforms

7705 SAR Gen 2

detail-level

Syntax

detail-level {**low** | **medium** | **high**}

no detail-level

Context

[Tree] (debug>service>id>igmp-snooping detail-level)

Full Context

debug service id igmp-snooping detail-level

Description

This command enables and configures the IGMP tracing detail level.

The **no** form of this command disables the IGMP tracing detail level.

Platforms

7705 SAR Gen 2

detail-level

Syntax

detail-level {**low** | **medium** | **high**}

no detail-level

Context

[Tree] (debug>service>id>mld detail-level)

Full Context

debug service id mld-snooping detail-level

Description

This command enables and configures the MLD tracing detail level.

The **no** form of this command disables the MLD tracing detail level.

Platforms

7705 SAR Gen 2

detail-level

Syntax

detail-level {**low** | **medium** | **high**}

no detail-level

Context

[Tree] (debug>service>id>dhcp detail-level)

Full Context

debug service id dhcp detail-level

Description

This command configures the DHCP tracing detail level.

The **no** form of the command disables debugging.

Parameters

low

Displays a low detail level for DHCP debugging.

medium

Displays a medium detail level for DHCP debugging.

high

Displays a high detail level for DHCP debugging.

Platforms

7705 SAR Gen 2

8.68 detect

detect

Syntax

detect num-moves *num-moves* **window** *minutes* [**trusted-mac-move-factor** *factor*]

Context

[Tree] (config>service>vpls>bgp-evpn>mac-duplication detect)

Full Context

configure service vpls bgp-evpn mac-duplication detect

Description

This command modifies the behavior of the **mac-duplication** command, which is always enabled by default. It monitors the number of moves of a MAC address for a period of time (window).

Default

detect num-moves 5 window 3 trusted-mac-move-factor 1

Parameters

num-moves

Identifies the number of MAC moves in a VPLS service. The counter is incremented when a specified MAC is locally relearned in the FDB or flushed from the FDB due to the reception of a better remote EVPN route for that MAC.

Values 3 to 10

Default 5

minutes

Specifies the length of the window in minutes.

Values 1 to 15

Default 3

factor

Specifies the multiplying value used to calculate a MAC duplication event. The *num-moves* value is multiplied by this value to determine the number of moves needed to declare a trusted MAC as duplicate.

For example, if *num-moves*=5 and *factor*=3, five moves within the window is enough to declare a non-trusted MAC as duplicate. However, 15 moves are needed to declare a trusted MAC as duplicate.

By default, the value of *factor* is 1, which means the factor for a trusted MAC is the same as for a non-trusted MAC. This provides a backwards compatible solution upon upgrade of the node.

Values 1 to 10

Default 1

Platforms

7705 SAR Gen 2

8.69 detection-time

detection-time

Syntax

detection-time *seconds*

no detection-time

Context

[Tree] (config>sys>security>dist-cpu-protection>policy>protocol>dynamic-parameters detection-time)

Full Context

configure system security dist-cpu-protection policy protocol dynamic-parameters detection-time

Description

When a dynamic enforcing policer is instantiated, it remains allocated until at least a contiguous conforming period of detection-time passes.

Default

detection-time 30

Parameters

seconds

Specifies the detection time.

Values 1 to 128000

Platforms

7705 SAR Gen 2

detection-time

Syntax

detection-time *seconds*

Context

[Tree] (config>sys>security>dist-cpu-protection>policy>static-policer detection-time)

Full Context

configure system security dist-cpu-protection policy static-policer detection-time

Description

When a policer is declared as in an "exceed" state, it remains as exceeding until a contiguous conforming period of **detection-time** passes. The **detection-time** only starts after the exceed-action hold-down is complete. If the policer detects another exceed during the detection count down then a hold-down is once again triggered before the policer re-enters the detection time (that is, the countdown timer starts again at the configured value). During the hold-down (and the detection-time), the policer is considered as in an "exceed" state.

Default

detection-time 30

Parameters

seconds

Specifies the detection time.

Values 1 to 128000

Platforms
7705 SAR Gen 2

8.70 deterministic

deterministic

Syntax
deterministic

Context
[Tree] (config>router>nat>inside deterministic)
[Tree] (config>service>vprn>nat>inside deterministic)

Full Context
configure router nat inside deterministic
configure service vprn nat inside deterministic

Description
Commands in this context configure deterministic NAT.

Platforms
7705 SAR Gen 2

8.71 deterministic-med

deterministic-med

Syntax
[**no**] **deterministic-med**

Context
[Tree] (config>service>vprn>bgp>best-path-selection deterministic-med)

Full Context
configure service vprn bgp best-path-selection deterministic-med

Description

This command controls how the BGP decision process compares routes on the basis of MED. When **deterministic-med** is configured, BGP groups paths that are equal up to the MED comparison step based on neighbor AS, and then compares the best path from each group to arrive at the overall best path. This change to the BGP decision process makes best path selection completely deterministic in all cases. Without **deterministic-med**, the overall best path selection is sometimes dependent on the order of the route arrival because of the rule that MED cannot be compared in routes from different neighbor AS.

Default

no deterministic-med

Platforms

7705 SAR Gen 2

deterministic-med

Syntax

[**no**] **deterministic-med**

Context

[Tree] (config>router>bgp>best-path-selection deterministic-med)

Full Context

configure router bgp best-path-selection deterministic-med

Description

This command controls how the BGP decision process compares routes on the basis of MED. When **deterministic-med** is configured, BGP groups paths that are equal up to the MED comparison step based on neighbor AS, and then compares the best path from each group to arrive at the overall best path. This change to the BGP decision process makes best path selection completely deterministic in all cases. Without **deterministic-med**, the overall best path selection is sometimes dependent on the order of the route arrival because of the rule that MED cannot be compared in routes from different neighbor AS.

Default

no deterministic-med

Platforms

7705 SAR Gen 2

8.72 device-label

device-label

Syntax

device-label *name*

no device-label

Context

[Tree] (config>system>management-interface>remote-management device-label)

Full Context

configure system management-interface remote-management device-label

Description

This command configures the metadata label that is supplied to all remote managers. This label can be used to group devices (network-nodes) that serve a common purpose or role.

If this command is also configured for a specific remote manager in the **config**>**system**> **management-interface**>**remote-management**>**manager** context, that configuration takes precedence.

The **no** form of this command causes an empty string to be used.

Parameters

name

Specifies the device-label name, up to 64 characters.

Platforms

7705 SAR Gen 2

device-label

Syntax

device-label *name*

no device-label

Context

[Tree] (config>system>management-interface>remote-management>manager device-label)

Full Context

configure system management-interface remote-management manager device-label

Description

This command configures the metadata label that is supplied to this remote manager. This label can be used to group devices (network-nodes) with a common purpose/role.

This command takes precedence over the same command configured in the global context (**config**>**system**>**management-interface**>**remote-management**).

The **no** form of this command causes the device-label name to be inherited from the global context (**config**>**system**>**management-interface**>**remote-management)**.

Parameters

name

Specifies the device-label name, up to 64 characters.

Platforms

7705 SAR Gen 2

8.73 device-name

device-name

Syntax

device-name *name*

no device-name

Context

[Tree] (config>system>management-interface>remote-management device-name)

Full Context

configure system management-interface remote-management device-name

Description

This command configures a device name that is supplied to all remote managers. This name identifies the specified SR OS node in the network.

If this command is also configured for a specific manager in the **config**>**system**>**management-interface**>**remote-management**> **manager** context, that configuration takes precedence.

The **no** form of this command causes the system to use the default device name (system-name).

Default

system-name

Parameters

name

Specifies the device name, up to 64 characters.

Platforms

7705 SAR Gen 2

device-name

Syntax

device-name *name*

no device-name

Context

[Tree] (config>system>management-interface>remote-management>manager device-name)

Full Context

configure system management-interface remote-management manager device-name

Description

This command configures a device name that is supplied to the specific manager. This name identifies the specified SR OS node in the network.

This command takes precedence over the same command configured in the global context **(config>system>management-interface>remote-management**).

The **no** form of this command causes the device name to be inherited from the global context (**config>system>management-interface>remote-management)**.

Default

system-name

Parameters

name

Specifies the device name, up to 64 characters.

Platforms

7705 SAR Gen 2

8.74 dh-group

dh-group

Syntax

dh-group {**1** | **2** | **5** | **14** | **15** | **19** | **20** | **21**}

Context

[Tree] (config>ipsec>ike-transform dh-group)

Full Context

configure ipsec ike-transform dh-group

Description

This command specifies the Diffie-Hellman group to be used in this IKE transform instance.

Default

dh-group 2 (1024-bit — More Modular Exponential (MODP))

Parameters

dh-group {1 | 2 | 5 | 14 | 15 | 19 | 20 | 21}

Specifies which Diffie-Hellman group to use for calculating session keys. More bits provide a higher level of security, but require more processing. Three groups are supported with IKE-v1:

Group 1: 768 bits

Group 2: 1024 bits

Group 5: 1536 bits

Group 14: 2048 bits

Group 15: 3072 bits

Group 19: P-256 ECC Curve, 256 bits

Group 20: P-384 ECC Curve, 384 bits

Group 21: P-512 ECC Curve, 512 bits

Platforms

7705 SAR Gen 2

8.75 dhcp

dhcp

Syntax

dhcp

Context

[Tree] (config>service>ies>if dhcp)

[Tree] (config>service>vprn>if dhcp)

[Tree] (config>service>vpls>spoke-sdp dhcp)

[Tree] (config>service>vprn dhcp)

[Tree] (config>service>vpls>sap dhcp)

[Tree] (config>service>vpls>mesh-sdp dhcp)

Full Context

configure service ies interface dhcp

configure service vprn interface dhcp

configure service vpls spoke-sdp dhcp

configure service vprn dhcp

configure service vpls sap dhcp

configure service vpls mesh-sdp dhcp

Description

Commands in this context configure DHCP parameters.

Platforms

7705 SAR Gen 2

dhcp

Syntax

[**no**] **dhcp** [**interface** *ip-int-name*]

[**no**] **dhcp mac** *ieee-address*

[**no**] **dhcp sap** *sap-id*

Context

[Tree] (debug>router>ip dhcp)

Full Context

debug router ip dhcp

Description

This command enables DHCP debugging.

The **no** form of this command disables debugging.

Parameters

ip-int-name

Specifies the name of the IP interface, up to 32 characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

ieee-address

Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

sap-id

Specifies the physical port identifier portion of the SAP definition.

Platforms

7705 SAR Gen 2

dhcp

Syntax

[**no**] **dhcp**

Context

[Tree] (debug>service>id dhcp)

Full Context

debug service id dhcp

Description

Commands in this context perform DHCP debugging.

The **no** form of the command disables DHCP debugging.

Platforms

7705 SAR Gen 2

dhcp

Syntax

[**no**] **dhcp**

Context

[Tree] (config>service>vprn>if>sap>ipsec-gw dhcp)

[Tree] (config>service>ies>if>sap>ipsec-gw dhcp)

Full Context

configure service vprn interface sap ipsec-gw dhcp

configure service ies interface sap ipsec-gw dhcp

Description

Commands in this context configure DHCPv4-based address assignment for IKEv2 remote-access tunnels.

The system acts as a DHCPv4 client on behalf of the IPsec client, and also a relay agent to relay DHCPv4 packets to the DHCPv4 server.

DHCPv4 DORA(Discovery/Offer/Request/Ack) exchange happens during IKEv2 remote-access tunnel setup. The system also supports standard renew.

In order to use this feature, the **relay-proxy** must be enabled on the corresponding interface (either the private interface or the interface that has the gi-address as the interface address).

Default

no dhcp

Platforms

7705 SAR Gen 2

dhcp

Syntax

dhcp

Context

[Tree] (config>router>if dhcp)

Full Context

configure router interface dhcp

Description

Commands in this context configure DHCP parameters.

Platforms

7705 SAR Gen 2

dhcp

Syntax

dhcp [**include-user-class**] [**timeout** *timeout*]

dhcp client-id [**string** *ascii-string*] [**hex** *hex-string*] [**include-user-class**] [**timeout** *timeout*]

no dhcp

Context

[Tree] (bof>autoconfigure>ipv4 dhcp)

Full Context

bof autoconfigure ipv4 dhcp

Description

This command configures the IPv4 DHCP client for OOB management. The OOB management IPv4 can be set using a DHCP server offer.

The **no** form of this command disables IPv4 DHCP OOB management.

Default

no dhcp

Parameters

include-user-class

Specifies to include Option 77 user class data in the offer.

client-id

Specifies to include the client ID for IPv4 Option 61 for auto-discovery. The identifier is opaque and is in string format. By default, this is the chassis serial number.

timeout

Specifies the DHCP offer timeout, in seconds.

Values 1 to 65535

Default 30

ascii-string

Specifies the string format for this option, up to 127 characters.

hex-string

Specifies the hexadecimal format for this option, up to 254 hex nibbles.

Values 0x0 to 0xFFFFFFFF

Platforms

7705 SAR Gen 2

dhcp

Syntax

dhcp [**include-user-class**] [**timeout** *timeout*]

dhcp client-id *duid-type* [**string** *ascii-string*] [**hex** *hex-string*] [**include-user-class**] [**timeout** *timeout*]

no dhcp

Context

[Tree] (bof>autoconfigure>ipv6 dhcp)

Full Context

bof autoconfigure ipv6 dhcp

Description

This command configures the IPv6 DHCP client for out-of-band (OOB) management. The OOB management IPv6 can be set using a DHCP server offer.

The **no** form of this command disables IPv6 DHCP client OOB management.

Default

no dhcp

Parameters

include-user-class

Specifies to include Option 15 user class data in the offer.

client-id

Specifies to include the client ID for IPv6 DHCP Option 1 for auto-discovery. The identifier is opaque and is in string format. By default, this is the chassis serial number.

seconds

Specifies the DHCP client ID timeout, in seconds.

Values 1 to 65535

duid-type

Specifies the type code of the server DUID.

Values duid-link-local, duid-enterprise

ascii-string

Specifies the string format for this option, up to 124 characters.

hex-string

Specifies the hexadecimal format for this option, up to 248 hex nibbles.

Values 0x0 to 0xFFFFFFFF

timeout

Specifies the DHCP offer timeout, in seconds.

Values 1 to 65535

Default 30

Platforms

7705 SAR Gen 2

8.76 dhcp-filter

dhcp-filter

Syntax

dhcp-filter *filter-id* [**create**]

no dhcp-filter

Context

[Tree] (config>filter dhcp-filter)

Full Context

configure filter dhcp-filter

Description

Commands in this context create and configure the specified DHCP filter policy.

Parameters

filter-id

Specifies the DHCP filter policy ID expressed as a decimal integer.

Values 1 to 65535

create

Keyword required to create the configuration context.

Platforms

7705 SAR Gen 2

8.77 dhcp-leasetime-threshold

dhcp-leasetime-threshold

Syntax

dhcp-leasetime-threshold [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no dhcp-leasetime-threshold

Context

[Tree] (config>system>persistence>options dhcp-leasetime-threshold)

Full Context

configure system persistence options dhcp-leasetime-threshold

Description

This command configures Dynamic Data Persistence (DDP) compact flash access optimization for DHCP leases.

The DHCP lease-time threshold controls the eligibility of a DHCP lease for persistency updates when no data other than the lease expiry time is to be updated. When the offered lease time of the DHCP lease is less than the configured threshold, the lease is flagged to skip persistency updates and will be installed with its full lease time upon a persistency recovery after a reboot.

The **dhcp-leasetime-threshold** command controls persistency updates for DHCPv4 and DHCPv6 leases for a DHCP relay or proxy and DHCPv4 leases for DHCP snooping (enabled with **subscriber-mgmt**) and a DHCP server (enabled with **dhcp-server**).

The **no** form of the command disables the DHCP lease time threshold.

Default

no dhcp-leasetime-threshold

Parameters

days

Specifies the threshold in days.

Values 0 to 7305

hours

Specifies the threshold in hours.

Values 0 to 23

minutes

Specifies the threshold in minutes.

Values 0 to 59

seconds

Specifies the threshold in seconds.

Values 0 to 59

Platforms

7705 SAR Gen 2

8.78 dhcp-server

dhcp-server

Syntax

dhcp-server

Context

[Tree] (config>system>persistence dhcp-server)

Full Context

configure system persistence dhcp-server

Description

This command configures DHCP server persistence parameters.

Platforms

7705 SAR Gen 2

8.79 dhcp-user-db

dhcp-user-db

Syntax

dhcp-user-db *local-user-db-name*

no dhcp-user-db

Context

[Tree] (config>service>vpls>sap dhcp-user-db)

Full Context

configure service vpls sap dhcp-user-db

Description

This command enabled access to LUDB for DHCPv4 hosts under the capture SAP. The name of this local user database must match the name of local user database configured under the **config>service>vprn/ies>sub-if>group-if>dhcp** context.

Parameters

local-user-db

Specifies the name of the local user database name up to 32 characters.

Platforms

7705 SAR Gen 2

8.80 dhcp6

dhcp6

Syntax

dhcp6

Context

[Tree] (config>service>vprn dhcp6)

[Tree] (config>service>vpls>sap dhcp6)

Full Context

configure service vprn dhcp6

configure service vpls sap dhcp6

Description

Commands in this context configure DHCPv6 parameters.

Platforms

7705 SAR Gen 2

dhcp6

Syntax

dhcp6

Context

[Tree] (config>system dhcp6)

Full Context

configure system dhcp6

Description

Commands in this context configure system-wide DHCPv6 parameters.

Platforms

7705 SAR Gen 2

dhcp6

Syntax

[**no**] **dhcp6** [*ip-int-name*]

[**no**] **dhcp6 client-identifier duid** *duid-hex-string* [**mask** *mask-hex-string*]

[**no**] **dhcp6 client-identifier link-layer-address** *lla-hex-string*

[**no**] **dhcp6 interface** *ip-int-name*

[**no**] **dhcp6 sap** *sap-id*

Context

[Tree] (debug>router>ip dhcp6)

Full Context

debug router ip dhcp6

Description

This command enables DHCPv6 debugging with optional interface, SAP, and client-identifier match criteria to filter the debug output.

The **no** form of this command disables debugging.

Parameters

ip-int-name

Specifies the name of an existing IP interface, up to 32 characters. Up to four DHCPv6 interface match criteria can be specified per routing instance.

client-identifier

Specifies a client identifier option match criteria.

duid duid-hex-string

Specifies a hexadecimal value for an opaque match on the client DUID in the client identifier option. When the actual length of the client DUID is longer than the length of

the specified hex-string, only the left most octets are matched. Up to four DHCPv6 client-identifier match criteria can be specified per routing instance.

Values 0x0 to 0xFFFFFFFF (maximum 260 hex nibbles)

mask mask-hex-string

Specifies an optional substring match criteria. When a mask is specified, both hex-string lengths must be equal.

Values 0x0 to 0xFFFFFFFF (maximum 260 hex nibbles)

link-layer-address lla-hex-string

Specifies a hexadecimal value for a link layer address field match of a type 1 (DUID-LLT) or type 3 (DUID-LL) client DUID in the client identifier option. When the actual length of the link layer address field is longer than the length of the specified hex-string, only the left most octets are matched. Up to four DHCPv6 client-identifier match criteria can be specified per routing instance.

Values 0x0 to 0xFFFFFFFF (maximum 252 hex nibbles)

sap-id

Specifies an existing SAP identifier. Up to four DHCPv6 SAP match criteria can be specified per routing instance.

Platforms

7705 SAR Gen 2

dhcp6

Syntax

[**no**] **dhcp6**

Context

[Tree] (config>service>vprn>if>sap>ipsec-gw dhcp6)

[Tree] (config>service>ies>if>sap>ipsec-gw dhcp6)

Full Context

configure service vprn interface sap ipsec-gw dhcp6

configure service ies interface sap ipsec-gw dhcp6

Description

Commands in this context configure DHCPv6-based address assignment for IKEv2 remote-access tunnels.

The system acts as a DHCPv6 client on behalf of the IPsec client, and also acts as a relay agent to relay DHCPv6 packets to the DHCPv6 server.

DHCPv6 exchange happens during IKEv2 remote-access tunnel setup. The system also supports standard renew.

Default

no dhcp6

Platforms

7705 SAR Gen 2

8.81 dhcp6-filter

dhcp6-filter

Syntax

dhcp6-filter *filter-id* [**create**]

no dhcp6-filter

Context

[Tree] (config>filter dhcp6-filter)

Full Context

configure filter dhcp6-filter

Description

Commands in this context create and configure the specified DHCPv6 filter policy.

The **no** form of this command reverts to the default.

Parameters

filter-id

Specifies the DHCPv6 filter policy ID expressed as a decimal integer.

Values 1 to 65535

create

Keyword required to create the configuration context.

Platforms

7705 SAR Gen 2

8.82 dhcp6-relay

dhcp6-relay

Syntax

[**no**] **dhcp6-relay**

Context

[Tree] (config>service>vprn>if>ipv6 dhcp6-relay)

[Tree] (config>service>ies>if>ipv6 dhcp6-relay)

Full Context

configure service vprn interface ipv6 dhcp6-relay

configure service ies interface ipv6 dhcp6-relay

Description

Commands in this context configure DHCPv6 relay parameters for the interface.

The **no** form of this command disables DHCPv6 relay.

Platforms

7705 SAR Gen 2

8.83 dhcp6-user-db

dhcp6-user-db

Syntax

dhcp6-user-db *local-user-db*

no dhcp6-user-db

Context

[Tree] (config>service>vpls>sap dhcp6-user-db)

Full Context

configure service vpls sap dhcp6-user-db

Description

This command enabled access to LUDB for DHCPv6 hosts under the capture SAP. The name of this LUDB must match the name of the LUDB configured under the **config>service>vprn/ies>sub-if>grp-if>dhcp** hierarchy.

The **no** form of this command reverts to the default.

Parameters

local-user-db

Specifies the name of the local-user-database, up to 32 characters.

Platforms

7705 SAR Gen 2

8.84 diffserv-te

diffserv-te

Syntax

diffserv-te [**mam** | **rdm**]

no diffserv-te

Context

[Tree] (config>router>rsvp diffserv-te)

Full Context

configure router rsvp diffserv-te

Description

This command enables Diff-Serv TE on the node.

When this command is enabled, IS-IS and OSPF starts advertising available bandwidth for each TE class configured under the diffserv-te node. This command only takes effect if the operator has already enabled TE at the IS-IS, OSPF, or both routing protocol levels:

configure router isis traffic-engineering

and/or:

configure router ospf traffic-engineering

IGP advertises for each RSVP interface in the system the available bandwidth in each TE class in the unreserved bandwidth TE parameter for that class. In addition, IGP continues to advertise the existing Maximum Reservable Link Bandwidth TE parameter to mean the maximum bandwidth that can be booked on a given interface by all classes. The value advertised is adjusted with the link **subscription** *percentage* factor configured in the **configure router rsvp interface** context.

The user configures the following parameters for the operation of Diff-Serv:

- Definition of TE classes, TE Class = {Class Type (CT), LSP priority}.
- Mapping of the system forwarding classes to the Diff-Serv Class Type (CT).
- Configuration of the percentage of RSVP interface bandwidth each CT shares, that is, the Bandwidth Constraint (BC).

When Diff-Serv TE is enabled, the system automatically enables the Max Allocation Model (MAM) Admission Control Policy. MAM represents the bandwidth constraint model for the admission control of an LSP reservation to a link.

Each CT shares a percentage of the Maximum Reservable Link Bandwidth through the user-configured Bandwidth Constraint (BC) for this CT. The Maximum Reservable Link Bandwidth is the link bandwidth multiplied by the RSVP interface subscription factor.

The sum of all BC values across all CTs does not exceed the Maximum Reservable Link Bandwidth. In other words, the following rule is enforced:

SUM (BCc) =< Max-Reservable-Bandwidth, 0 <= c <= 7

An LSP of class-type CTc, setup priority p, holding priority h (h=<p), and bandwidth B is admitted into a link if the following condition is satisfied:

B <= Unreserved Bandwidth for TE-Class[i]

where TE-Class [i] maps to < CTc , p > in the definition of the TE classes on the node. The bandwidth reservation is effected at the holding priority, that is, in TE-class [j] = <CTc, h>. Thus, the reserved bandwidth for CTc and the unreserved bandwidth for the TE classes using CTc are updated as follows:

Reserved(CTc) = Reserved(CTc) + B

Unreserved TE-Class [j] = BCc - SUM (Reserved(CTc,q)) for 0<= q <= h

Unreserved TE-Class [i] = BCc - SUM (Reserved(CTc,q)) for 0<= q <= p

The same is done to update the unreserved bandwidth for any other TE class making use of the same CTc. These new values are advertised to the rest of the network at the next IGP-TE flooding.

The Russian Doll Model (RDM) LSP admission control policy allows bandwidth sharing across Class Types. It provides a hierarchical model by which the reserved bandwidth of a CT is the sum of the reserved bandwidths of the numerically equal and higher CTs.

The RDM model is defined using the following equations:

SUM (Reserved (CTc)) <= BCb,

where the SUM is across all values of c in the range b <= c <= (MaxCT - 1), and BCb is the bandwidth constraint of CTb.

BC0= Max-Reservable-Bandwidth, so that

SUM (Reserved(CTc)) <= Max-Reservable-Bandwidth,

where the SUM is across all values of c in the range 0 <= c <= (MaxCT - 1).

When Diff-Serv is disabled on the node, this model degenerates into a single default CT internally with eight preemption priorities and a non-configurable BC equal to the Maximum Reservable Link Bandwidth. This would behave exactly like CT0 with eight preemption priorities and BC= Maximum Reservable Link Bandwidth if Diff-Serv was enabled.

The enabling or disabling of Diff-Serv TE on the system requires the RSVP and MPLS protocol be shutdown.

The **no** form of this command reverts to the default value.

Default

no diffserv-te

Parameters

mam

Defines the default admission control policy for Diff-Serv LSPs.

rdm

Defines Russian doll model for the admission control policy of Diff-Serv LSPs.

Platforms

7705 SAR Gen 2

8.85 digital-coherent-optics

digital-coherent-optics

Syntax

[**no**] **digital-coherent-optics**

Context

[Tree] (config>port>transceiver digital-coherent-optics)

Full Context

configure port transceiver digital-coherent-optics

Description

This command specifies if a digital coherent optics module is used for this port.

The **no** form of this command specifies that the optical module used in this port is not a digital coherent optics module.

Default

no digital-coherent-optics

Platforms

7705 SAR Gen 2

8.86 dir

dir

Syntax

dir [*file-url*] [**sort-order** { **d** | **n** | **s**}] [**reverse**]

Context

[Tree] (file dir)

Full Context

file dir

Description

This command displays a list of files and subdirectories in a directory.

Parameters

file-url

Specifies the path or directory name.

Use the *file-url* with the optional wildcard (*) to reduce the number of files to list.

sort-order {d | n | s}

Specifies the sort order.

Values d — date
n — name
s — size

reverse

Reverses the sort order.

Default Lists all files in the current working directory.

local-url	[*cflash-id*/][*file-path*] up to 200 characters, including cflash-id directory length up to 99 each
remote-url	[{ftp:// | tftp://}*login*:*pswd*@*remote-locn*/][*file-path*] up to 247 characters directory length up to 99 characters each
remote-locn	[hostname | ipv4-address | [ipv6-address]]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]

	x:x:x:x:x:x:d.d.d.d[*-interface*] *x* - [0 to FFFF]H *d* - [0 to 255]D interface - up to 32 characters, for link local addresses 255
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

Output

The following output is an example of directory information.

Output Example

```
A:cses-E12>file cf3:\ # dir
  - dir [<file-url>] [sort-order { d | n | s}] [reverse]

 <file-url>           : <local-url> | <remote-url>
                        local-url      - [<cflash-id>/][<file-path>]
                                         200 chars max, including cflash-id
                                         directory length 99 chars max each
                        remote-url     - [ftp://<login>:<pswd>@<remote-locn>/
                                         ][<file-path>]
                                         255 chars max
                                         directory length 99 chars max each
                        remote-locn    - [ <hostname> | <ipv4-address> |
                                         [<ipv6-address>]]
                        ipv4-address   - a.b.c.d
                        ipv6-address   - x:x:x:x:x:x:x:x[-interface]
                                         x:x:x:x:x:x:d.d.d.d[-interface]
                                         x - [0..FFFF]H
                                         d - [0..255]D
                                         interface - 32 chars max, for link
                                         local addresses
                        cflash-id      - cf1:|cf1-A:|cf1-B:|cf2:|cf2-A:|
                                         cf2-B:|cf3:|cf3-A:|cf3-B:
 < d | n | s>         : Sort order: d - date, n - name, s - size
 <reverse>            : keyword - reverse order
A:cses-E12>file cf3:\ # dir
```

8.87 direction

direction

Syntax

direction *ipsec-direction*

no direction

Context

[Tree] (config>ipsec>static-sa direction)

Full Context

configure ipsec static-sa direction

Description

This command configures the direction for an IPsec manual SA.

The **no** form of this command reverts to the default value.

Default

direction bidirectional

Parameters

ipsec-direction

Identifies the direction to which this static SA entry can be applied.

Values inbound, outbound, bidirectional

Platforms

7705 SAR Gen 2

direction

Syntax

direction

Context

[Tree] (config>system>security>keychain direction)

Full Context

configure system security keychain direction

Description

This command specifies the data type that indicates the TCP stream direction to apply the keychain.

Platforms

7705 SAR Gen 2

8.88 disable-4byte-asn

disable-4byte-asn

Syntax

[**no**] **disable-4byte-asn**

Context

[Tree] (config>service>vprn>bgp disable-4byte-asn)

[Tree] (config>service>vprn>bgp>group>neighbor disable-4byte-asn)

[Tree] (config>service>vprn>bgp>group disable-4byte-asn)

Full Context

configure service vprn bgp disable-4byte-asn

configure service vprn bgp group neighbor disable-4byte-asn

configure service vprn bgp group disable-4byte-asn

Description

This command disables the use of 4-byte ASNs. It can be configured at all 3 level of the hierarchy so it can be specified down to the per peer basis.

If this command is enabled 4-byte ASN support should not be negotiated with the associated remote peer(s).

The **no** form of this command resets the behavior to the default which is to enable the use of 4-byte ASN.

Platforms

7705 SAR Gen 2

disable-4byte-asn

Syntax

[**no**] **disable-4byte-asn**

Context

[Tree] (config>router>bgp disable-4byte-asn)

[Tree] (config>router>bgp>group>neighbor disable-4byte-asn)

[Tree] (config>router>bgp>group disable-4byte-asn)

Full Context

configure router bgp disable-4byte-asn

configure router bgp group neighbor disable-4byte-asn

configure router bgp group disable-4byte-asn

Description

This command disables the support of 4-byte ASNs. It can be configured at all three levels of the hierarchy so it can be specified down to the per-peer basis.

If this command is enabled, 4-byte ASN support should not be negotiated with the associated remote peers.

The **no** form of this command resets the behavior to the default which is to enable the support of 4-byte ASN.

Default

no disable-4byte-asn

Platforms

7705 SAR Gen 2

8.89 disable-aging

disable-aging

Syntax

[**no**] **disable-aging**

Context

[Tree] (config>service>template>vpls-template disable-aging)

[Tree] (config>service>vpls>spoke-sdp disable-aging)

[Tree] (config>service>vpls>sap disable-aging)

[Tree] (config>service>vpls disable-aging)

Full Context

configure service template vpls-template disable-aging

configure service vpls spoke-sdp disable-aging

configure service vpls sap disable-aging

configure service vpls disable-aging

Description

This command disables MAC address aging across a VPLS service, on a VPLS service SAP or spoke-SDP, or VXLAN instance with static binds. Learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the VPLS forwarding database (FDB).

The **disable-aging** command turns off aging for local and remote learned MAC addresses. When **no disable-aging** is specified for a VPLS, aging can be disabled for specific SAPs, spoke-SDPs, and VXLAN instances (or any combination) by entering the **disable-aging** command at the appropriate level.

When the **disable-aging command** is entered at the VPLS level, the aging state of individual SAPs or SDPs or VXLAN instance is ignored.

The **no** form of this command enables aging on the VPLS service.

Default

no disable-aging

Except for VXLAN instances, where the **disable-aging** is the default mode

Platforms

7705 SAR Gen 2

disable-aging

Syntax

[**no**] **disable-aging**

Context

[Tree] (config>service>pw-template disable-aging)

Full Context

configure service pw-template disable-aging

Description

This command disables MAC address aging across a service.

The **no** form of this command enables aging.

Default

no disable-aging

Platforms

7705 SAR Gen 2

8.90 disable-capability-negotiation

disable-capability-negotiation

Syntax

[**no**] **disable-capability-negotiation**

Context

[Tree] (config>service>vprn>bgp>group>neighbor disable-capability-negotiation)

[Tree] (config>service>vprn>bgp>group disable-capability-negotiation)

Full Context

configure service vprn bgp group neighbor disable-capability-negotiation

configure service vprn bgp group disable-capability-negotiation

Description

This command disables the exchange of capabilities. When command is enabled and after the peering is flapped, any new capabilities are not negotiated and strictly supports IPv4 routing exchanges with that peer.

The **no** form of this command removes this command from the configuration and restores the normal behavior.

Default

no disable-capability-negotiation

Platforms

7705 SAR Gen 2

disable-capability-negotiation

Syntax

[**no**] **disable-capability-negotiation**

Context

[Tree] (config>router>bgp>group disable-capability-negotiation)

[Tree] (config>router>bgp>group>neighbor disable-capability-negotiation)

Full Context

configure router bgp group disable-capability-negotiation

configure router bgp group neighbor disable-capability-negotiation

Description

This command disables capability negotiation. When the command is enabled and after the peering is flapped, any new capabilities are not negotiated and will strictly support IPv4 routing exchanges with that peer.

The **no** form of this command removes this command from the configuration and restores the normal behavior.

Default

no disable-capability-negotiation

Platforms

7705 SAR Gen 2

8.91 disable-client-reflect

disable-client-reflect

Syntax

[**no**] **disable-client-reflect**

Context

[Tree] (config>service>vprn>bgp>group disable-client-reflect)

[Tree] (config>service>vprn>bgp>group>neighbor disable-client-reflect)

[Tree] (config>service>vprn>bgp disable-client-reflect)

Full Context

configure service vprn bgp group disable-client-reflect

configure service vprn bgp group neighbor disable-client-reflect

configure service vprn bgp disable-client-reflect

Description

This command disables the reflection of routes by the route reflector to the group or neighbor. This only disables the reflection of routes from other client peers. Routes learned from non-client peers are still reflected to all clients.

The **no** form re-enables client reflection of routes.

Default

no disable-client-reflect

Platforms

7705 SAR Gen 2

disable-client-reflect

Syntax

[**no**] **disable-client-reflect**

Context

[Tree] (config>router>bgp>group disable-client-reflect)

[Tree] (config>router>bgp disable-client-reflect)

[Tree] (config>router>bgp>group>neighbor disable-client-reflect)

Full Context

configure router bgp group disable-client-reflect

configure router bgp disable-client-reflect

configure router bgp group neighbor disable-client-reflect

Description

This command determines whether routes received from neighbors considered to be RR clients are reflected to other clients.

The **no** form enables client reflection of routes.

Default

no disable-client-reflect

Platforms

7705 SAR Gen 2

8.92 disable-communities

disable-communities

Syntax

disable-communities [**standard**] [**extended**] [**large**]

no disable-communities

Context

[Tree] (config>service>vprn>bgp>group>neighbor disable-communities)

[Tree] (config>service>vprn>bgp>group disable-communities)

[Tree] (config>service>vprn>bgp disable-communities)

Full Context

configure service vprn bgp group neighbor disable-communities

configure service vprn bgp group disable-communities

configure service vprn bgp disable-communities

Description

This command configures BGP to disable sending standard, extended, or large communities to specific peers.

By default, all communities that are attached to a BGP route (any address family) are not stripped from the route when it is advertised to any type of peer: IBGP, EBGP or confed-EBGP.

Default

no disable-communities

Parameters

standard

Specifies that standard 4-byte communities should be removed.

extended

Specifies that 8-byte extended communities (of all types) should be removed.

large

Specifies that 12-byte large communities should be removed.

Platforms

7705 SAR Gen 2

disable-communities

Syntax

disable-communities [**standard**] [**extended**] [**large**]

no disable-communities

Context

[Tree] (config>router>bgp>group disable-communities)

[Tree] (config>router>bgp disable-communities)

[Tree] (config>router>bgp>group>neighbor disable-communities)

Full Context

configure router bgp group disable-communities

configure router bgp disable-communities

configure router bgp group neighbor disable-communities

Description

This command configures BGP to disable sending standard, extended, or large communities to specific peers.

By default, all communities that are attached to a BGP route (any address family) are not stripped from the route when it is advertised to any type of peer: IBGP, EBGP, or confed-EBGP.

Default

no disable-communities

Parameters

standard

Advertise the Communities attribute to peers.

extended

Advertise the Extended Communities attribute to peers.

large

Advertise the Large Communities attribute to peers.

Platforms

7705 SAR Gen 2

8.93 disable-explicit-null

disable-explicit-null

Syntax

[**no**] **disable-explicit-null**

Context

[Tree] (config>router>bgp>label-allocation>label-ipv6 disable-explicit-null)

Full Context

configure router bgp label-allocation label-ipv6 disable-explicit-null

Description

This command controls the allocation and use of explicit null MPLS labels with labeled-unicast ipv6 routes.

When this command is enabled (**no disable-explicit-null**), the following behaviors apply:

- during the times when the router is required to act as the BGP next-hop of a label-unicast IPv6 route that it is advertising, it sets the BGP label value to IPv6 explicit null (value 2), forcing a POP behavior for received packets.
- received label-unicast IPv6 routes never create tunnels in TTM that can be used to resolve other BGP routes (with an IPv6 next-hop).

- a received label-unicast IPv6 route can be resolved by a label-ipv4 BGP tunnel that is transported over a stacked tunnel (SR-TE LSP or LDPoRSVP LSP)

When this command is disabled (**disable-explicit-null**), the following behaviors apply:

- during those times when the router is required to act as the BGP next-hop of a label-unicast IPv6 route that it is advertising, it sets the BGP label value to a proper value in the dynamic label range and programs a POP or SWAP operation for that label, depending on the origin of the route and various import policy actions that could apply to the route
- received label-unicast IPv6 routes that have a prefix length of 128 bits are automatically installed in TTM so that they can be used to resolve other (non-labeled-unicast) BGP routes (with an IPv6 next-hop)
- a received label-unicast ipv6 route cannot be resolved by a label-ipv4 BGP tunnel that is transported over a stacked tunnel (SR-TE LSP or LDPoRSVP LSP)
- the label-ipv6 routes used for ECMP towards an IPv6 destination cannot be a mix of routes with regular label values and routes with special (IPv6 explicit null) label values

Changes in the operational status do not cause the BGP sessions of the base router to reset.

Platforms

7705 SAR Gen 2

8.94 disable-fast-external-failover

disable-fast-external-failover

Syntax

[**no**] **disable-fast-external-failover**

Context

[Tree] (config>service>vprn>bgp>group>neighbor disable-fast-external-failover)

[Tree] (config>service>vprn>bgp>group disable-fast-external-failover)

[Tree] (config>service>vprn>bgp disable-fast-external-failover)

Full Context

configure service vprn bgp group neighbor disable-fast-external-failover

configure service vprn bgp group disable-fast-external-failover

configure service vprn bgp disable-fast-external-failover

Description

This command configures BGP fast external failover.

Platforms
7705 SAR Gen 2

disable-fast-external-failover

Syntax
[**no**] **disable-fast-external-failover**

Context
[Tree] (config>router>bgp>group>neighbor disable-fast-external-failover)
[Tree] (config>router>bgp disable-fast-external-failover)
[Tree] (config>router>bgp>group disable-fast-external-failover)

Full Context
configure router bgp group neighbor disable-fast-external-failover
configure router bgp disable-fast-external-failover
configure router bgp group disable-fast-external-failover

Description
This command configures BGP fast external failover.

Default
no disable-fast-external-failover

Platforms
7705 SAR Gen 2

8.95 disable-graceful-shutdown

disable-graceful-shutdown

Syntax
[**no**] **disable-graceful-shutdown**

Context
[Tree] (config>system>login-control>ssh disable-graceful-shutdown)

Full Context
configure system login-control ssh disable-graceful-shutdown

Description

This command enables graceful shutdown of SSH sessions.

The **no** form of this command disables graceful shutdown of SSH sessions.

Platforms

7705 SAR Gen 2

8.96 disable-ldp-sync

disable-ldp-sync

Syntax

[**no**] **disable-ldp-sync**

Context

[Tree] (config>router>isis disable-ldp-sync)

Full Context

configure router isis disable-ldp-sync

Description

This command disables the IGP-LDP synchronization feature on all interfaces participating in the OSPF or IS-IS routing protocol. When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces which have the IGP-LDP synchronization enabled if the currently advertised cost is different. It will then disable IGP-LDP synchronization for all interfaces. This command does not delete the interface configuration. The **no** form of this command has to be entered to re-enable IGP-LDP synchronization for this routing protocol.

The **no** form of this command restores the default settings and re-enables IGP-LDP synchronization on all interfaces participating in the OSPF or IS-IS routing protocol and for which the ldp-sync-timer is configured.

Default

no disable-ldp-sync

Platforms

7705 SAR Gen 2

disable-ldp-sync

Syntax

[**no**] **disable-ldp-sync**

Context

[Tree] (config>router>ospf disable-ldp-sync)

[Tree] (config>router>ospf3 disable-ldp-sync)

Full Context

configure router ospf disable-ldp-sync

configure router ospf3 disable-ldp-sync

Description

This command disables the IGP-LDP synchronization feature on all interfaces participating in the OSPF routing protocol. When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces which have the IGP-LDP synchronization enabled if the currently advertised cost is different. It will then disable IGP-LDP synchronization for all interfaces. This command does not delete the interface configuration. The **no** form of this command has to be entered to re-enable IGP-LDP synchronization for this routing protocol.

The **no** form of this command restores the default settings and re-enables IGP-LDP synchronization on all interfaces participating in the OSPF or IS-IS routing protocol and for which the ldp-sync-timer is configured.

Default

no disable-ldp-sync

Platforms

7705 SAR Gen 2

8.97 disable-learning

disable-learning

Syntax

[**no**] **disable-learning**

Context

[Tree] (config>service>vpls>spoke-sdp disable-learning)

[Tree] (config>service>vpls disable-learning)

[Tree] (config>service>vpls>sap disable-learning)

[Tree] (config>service>template>vpls-template disable-learning)

Full Context

configure service vpls spoke-sdp disable-learning

configure service vpls disable-learning

configure service vpls sap disable-learning

configure service template vpls-template disable-learning

Description

This command disables learning of new MAC addresses in the VPLS forwarding database (FDB) for the service instance, SAP instance, spoke-SDP instance, or VXLAN instance.

When **disable-learning** is enabled, new source MAC addresses are not entered in the VPLS service forwarding database. This applies for both local and remote MAC addresses.

When **disable-learning** is disabled, new source MAC addresses are learned and entered into the VPLS forwarding database.

This parameter is mainly used in conjunction with the **discard-unknown** command.

The **no** form of this command enables learning of MAC addresses.

Default

no disable-learning

Normal MAC learning is enabled. The default mode for VXLAN instances is **disable-learning**.

Platforms

7705 SAR Gen 2

disable-learning

Syntax

[**no**] **disable-learning**

Context

[Tree] (config>service>pw-template disable-learning)

Full Context

configure service pw-template disable-learning

Description

This command enables learning of new MAC addresses.

This parameter is mainly used in conjunction with the **discard-unknown** command.

The **no** form of this command enables learning of MAC addresses.

Default

no disable-learning (Normal MAC learning is enabled)

Platforms

7705 SAR Gen 2

8.98 disable-route-table-install

disable-route-table-install

Syntax

[**no**] **disable-route-table-install**

Context

[Tree] (config>router>bgp disable-route-table-install)

Full Context

configure router bgp disable-route-table-install

Description

This command disables the installation of all IPv4, label-ipv4, IPv6 and label-ipv6 routes into the route table and tunnel table associated with the BGP instance.

Configuring this command prevents the advertisement of all IPv4, label-ipv4, IPv6 and label-ipv6 routes if there is a change of the BGP next-hop to one of the router's own addresses. Typically, this is only useful on a router that is a control-plane route reflector (not in the datapath).

The **no** form of the command enables the installation of all IPv4, label-ipv4, IPv6 and label-ipv6 routes into the route table and tunnel table associated with the BGP instance.

Default

no disable-route-table-install

Platforms

7705 SAR Gen 2

disable-route-table-install

Syntax

[**no**] **disable-route-table-install**

Context

[Tree] (config>router>policy-options>policy-statement>entry>action disable-route-table-install)

[Tree] (config>router>policy-options>policy-statement>default-action disable-route-table-install)

Full Context

configure router policy-options policy-statement entry action disable-route-table-install

configure router policy-options policy-statement default-action disable-route-table-install

Description

This command specifies that BGP routes (IPv4, IPv6, label-ipv4, label-ipv6) matching the policy entry (or, depending on the context, the policy's default-action) should not be installed in the route table, tunnel table (in the case of label-ipv4 routes), or IP FIB table.

This policy action has an effect only in BGP peer import policies. This policy action does not prevent the matched routes from contributing toward aggregate routes and does not prevent the matched routes from being advertised, even if next-hop-self has been applied. This means that incorrect use of this feature can blackhole traffic.

Marking label-ipv4 routes with this action does not prevent label-swap (ILM) entries from being programmed when such routes are re-advertised with a new BGP next-hop and label.

The **no** form of this command provides the default behavior of installing routes that are selected as the best path, ECMP path or backup path, depending on the circumstances.

Default

no disable-route-table-install

Platforms

7705 SAR Gen 2

8.99 disable-router-alert-check

disable-router-alert-check

Syntax

[**no**] **disable-router-alert-check**

Context

[Tree] (config>router>igmp>if disable-router-alert-check)

Full Context

configure router igmp interface disable-router-alert-check

Description

This command disables router alert checking for IGMP/MLD messages received on this interface.

The **no** form of this command enables router alert checking.

Platforms

7705 SAR Gen 2

disable-router-alert-check

Syntax

[**no**] **disable-router-alert-check**

Context

[Tree] (config>router>mld>if disable-router-alert-check)

Full Context

configure router mld interface disable-router-alert-check

Description

This command enables router alert checking for MLD messages received on this interface.

The **no** form of this command disables router alert checking.

Platforms

7705 SAR Gen 2

disable-router-alert-check

Syntax

[**no**] **disable-router-alert-check**

Context

[Tree] (config>service>vpls>spoke-sdp>mld-snooping disable-router-alert-check)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping disable-router-alert-check)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping disable-router-alert-check)

[Tree] (config>service>vpls>sap>mld-snooping disable-router-alert-check)

[Tree] (config>service>vpls>mesh-sdp>mld-snooping disable-router-alert-check)

[Tree] (config>service>vpls>sap>igmp-snooping disable-router-alert-check)

Full Context

configure service vpls spoke-sdp mld-snooping disable-router-alert-check

configure service vpls mesh-sdp igmp-snooping disable-router-alert-check

configure service vpls spoke-sdp igmp-snooping disable-router-alert-check

configure service vpls sap mld-snooping disable-router-alert-check

configure service vpls mesh-sdp mld-snooping disable-router-alert-check

configure service vpls sap igmp-snooping disable-router-alert-check

Description

This command disables the IGMP or MLD router alert check option.

The **no** form of this command enables the router alert check.

Default

no disable-router-alert-check

Platforms

7705 SAR Gen 2

disable-router-alert-check

Syntax

[**no**] **disable-router-alert-check**

Context

[Tree] (config>service>vprn>igmp>if disable-router-alert-check)

Full Context

configure service vprn igmp interface disable-router-alert-check

Description

This command disables the IGMP router alert check option.

The **no** form of this command enables the router alert check.

Platforms

7705 SAR Gen 2

disable-router-alert-check

Syntax

[**no**] **disable-router-alert-check**

Context

[Tree] (config>service>vprn>mld>if disable-router-alert-check)

Full Context

configure service vprn mld interface disable-router-alert-check

Description

This command disables router alert checking for MLD messages received on this interface.

The no form of this command enables the router alert checking.

Platforms

7705 SAR Gen 2

8.100 disable-stickiness

disable-stickiness

Syntax

[**no**] **disable-stickiness**

Context

[Tree] (config>aaa>radius-srv-plcy>servers disable-stickiness)

Full Context

configure aaa radius-server-policy servers disable-stickiness

Description

This command disables a subscriber RADIUS accounting session from sticking with a single server under normal working conditions. If a direct algorithm is used, all subscriber RADIUS sessions will go directly to the server with the lowest configured index. If a failure occurs, a new in-service server with the next lowest index is used. When the original server recovers, if stickiness is not disabled, all existing sessions will continue to use the new server. This command disables stickiness, and as a result, the recovered original RADIUS server will again service every subscriber. If a round-robin algorithm is used and stickiness is not disabled, an accounting session for a particular subscriber (or host, depending on the accounting mode) will stay with the same server. This command removes the stickiness and all subscriber accounting messages will go through the list of servers in a round-robin manner.

Platforms

7705 SAR Gen 2

8.101 disable-targeted-session

disable-targeted-session

Syntax

[**no**] **disable-targeted-session**

Context

[Tree] (config>router>ldp>targ-session disable-targeted-session)

Full Context

configure router ldp targeted-session disable-targeted-session

Description

This command disables support for SDP triggered automatic generated targeted sessions. Targeted sessions are LDP sessions between non-directly connected peers. The discovery messages for an indirect LDP session are addressed to the specified peer and not to the multicast address.

The **no** form of this command enables the set up of any targeted sessions.

Default

no disable-targeted-session

Platforms

7705 SAR Gen 2

8.102 disallow-igp

disallow-igp

Syntax

[**no**] **disallow-igp**

Context

[Tree] (config>router>static-route-entry>indirect>tunnel-next-hop disallow-igp)

Full Context

configure router static-route-entry indirect tunnel-next-hop disallow-igp

Description

This optional command determines if the associated static route can be resolved via an IGP next-hop in the RTM if no tunnel next-hops are found in TTM.

When configured, the associated static route will not be resolved to an available IGP route in the RTM.

The **no** form of this command returns the behavior to the default, which allows the static route to be resolved via an IGP route in the RTM if no tunnel next-hop can be found in the TTM.

Default

no disallow-igp

Platforms

7705 SAR Gen 2

8.103 disallow-sequence-keys

disallow-sequence-keys

Syntax

disallow-sequence-keys *number-of-characters*

no disallow-sequence-keys

Context

[Tree] (config>system>sec>passwd>compl disallow-sequence-keys)

Full Context

configure system security password complexity-rules disallow-sequence-keys

Description

This command configures the number of consecutive characters that are not allowed to be entered as part of the password on a U.S. English or Korean keyboard. These characters can be lowercase or uppercase letters, or numbers. Special characters are not taken into account. These consecutive characters can be horizontal (left to right) or (right to left) or diagonal (up to bottom or bottom to top). If the number of consecutive characters is equal to or larger than the configured value, the password is disallowed.

For example, if the user attempts to use the password "dsalkjhgfdsa", with this command configured to 8, the system rejects the password because the first consecutive sequence “dsa” is 3 lowercase letters, which passes the check, but the second consecutive sequence is “lkjhgfdsa”, which consists of 9 consecutive lowercase letters and this does not pass the check.

The **no** form of this command removes the restriction on the number of characters.

Default

no disallow-sequence-keys

Parameters

number-of-characters

Specifies the number of characters.

Values 2 to 8

Platforms

7705 SAR Gen 2

8.104 discard

discard

Syntax

discard [**now**]

Context

[Tree] (candidate discard)

Full Context

candidate discard

Description

This command deletes the entire contents of the candidate configuration and exits the edit-cfg mode. Undo cannot be used to recover a candidate that has been discarded with **candidate discard**.

Parameters

now

Avoids a confirmation prompt for the discard.

Platforms

7705 SAR Gen 2

8.105 discard-changes

discard-changes

Syntax

[**no**] **discard-changes**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization discard-changes)

Full Context

configure system security profile netconf base-op-authorization discard-changes

Description

This command enables the NETCONF <discard-changes> RPC.

The **no** form of this command disables the RPC.

Default

no discard-changes

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

8.106 discard-rx-pause-frames

discard-rx-pause-frames

Syntax

[**no**] **discard-rx-pause-frames**

Context

[Tree] (config>port>ethernet discard-rx-pause-frames)

Full Context

configure port ethernet discard-rx-pause-frames

Description

This command discards received pause frames. Pause frames are used for local link flow control.

The **no** form of this command processes pause frames upon reception and the transmit side of the receiving port pauses in its transmissions.

Default

no discard-rx-pause-frames

Platforms

7705 SAR Gen 2

8.107 discard-unknown

discard-unknown

Syntax

[**no**] **discard-unknown**

Context

[Tree] (config>service>template>vpls-template discard-unknown)

[Tree] (config>service>vpls discard-unknown)

Full Context

configure service template vpls-template discard-unknown

configure service vpls discard-unknown

Description

By default, packets with unknown destination MAC addresses are flooded. If **discard-unknown** is enabled at the VPLS level, packets with unknown destination MAC address will be dropped instead (even when configured FDB size limits for VPLS or SAP are not yet reached).

The **no** form of this command allows flooding of packets with unknown destination MAC addresses in the VPLS.

Default

no discard-unknown

Platforms

7705 SAR Gen 2

8.108 discard-unknown-source

discard-unknown-source

Syntax

[**no**] **discard-unknown-source**

Context

[Tree] (config>service>template>vpls-sap-template discard-unknown-source)

[Tree] (config>service>vpls>sap discard-unknown-source)

[Tree] (config>service>vpls>spoke-sdp discard-unknown-source)

Full Context

configure service template vpls-sap-template discard-unknown-source

configure service vpls sap discard-unknown-source

configure service vpls spoke-sdp discard-unknown-source

Description

When this command is enabled, packets received on a SAP, a spoke-SDP, or a static VXLAN instance with an unknown source MAC address will be dropped only if the maximum number of MAC addresses for that SAP or spoke-SDP (see **max-nbr-mac-addr** [**config>service>vpls>sap max-nbr-mac-addr**, **config>service>vpls>spoke-sdp max-nbr-mac-addr**]) has been reached. If **max-nbr-mac-addr** has not been set for the SAP or spoke-SDP, enabling **discard-unknown-source** has no effect.

When disabled, the packets are forwarded based on the destination MAC addresses.

The **no** form of this command causes packets with an unknown source MAC addresses to be forwarded by destination MAC addresses in VPLS.

Default

no discard-unknown-source

Platforms

7705 SAR Gen 2

discard-unknown-source

Syntax

[**no**] **discard-unknown-source**

Context

[Tree] (config>service>pw-template discard-unknown-source)

Full Context

configure service pw-template discard-unknown-source

Description

When this command is enabled, packets received with an unknown source MAC address will be dropped only if the maximum number of MAC addresses have been reached.

When disabled, the packets are forwarded based on the destination MAC addresses.

The **no** form of this command causes packets with an unknown source MAC addresses to be forwarded by destination MAC addresses.

Default

no discard-unknown-source

Platforms

7705 SAR Gen 2

8.109 disconnect

disconnect

Syntax

disconnect [**address** *ip-address* | **session-id** *session-id* | **username** *user-name* | {**console** | **bluetooth** | **telnet** | **ftp** | **ssh** | **netconf** | **grpc**}]

Context

[Tree] (admin disconnect)

Full Context

admin disconnect

Description

This command disconnects a user from a session.

Issuing the **disconnect** command without any parameters disconnects the session in which the command was executed.

If any of the session type options (for example, **console, bluetooth, telnet, FTP, SSH**) are specified, only the respective sessions are affected.

If no session type options are specified, all sessions from the IP address or from the specified user are disconnected.

Any task that the user is executing is terminated. FTP files accessed by the user are not removed.

A major severity security log event is created specifying what was terminated and by whom.

By default, no disconnect options are configured.

Parameters

ip-address

Specifies the IP address to disconnect, specified in dotted decimal notation.

Values

ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x* (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

session-id

The model-driven session ID. Can be obtained using the **show system management-interface datastore-locks** [**detail**] command.

user-name

Specifies the name of the user. The name can be up to 32 characters.

console

Disconnects the console session.

bluetooth

Disconnects the Bluetooth session.

telnet

Disconnects the Telnet session.

ftp

Disconnects the FTP session.

ssh

Disconnects the SSH session.

netconf

Disconnects the NETCONF session.

grpc

Disconnects the gRPC session.

Platforms

7705 SAR Gen 2

8.110 discovery-interval

discovery-interval

Syntax

discovery-interval *interval-secs* [**boot** *interval-secs*]

no discovery-interval

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ipsec discovery-interval)

Full Context

configure redundancy multi-chassis peer mc-ipsec discovery-interval

Description

This command specifies the time interval of tunnel-group stays in the Discovery state. Interval-1 is used as discovery-interval when a new tunnel-group is added to multi-chassis redundancy (mp-ipsec); interval-2 is used as discovery-interval when the system boots up, it is optional, when it is not specified, the interval-1 will be used.

Default

discovery-interval 300 boot 300

Parameters

interval-secs

Specifies the maximum duration, in seconds, of the discovery interval during which a newly activated multi- chassis IPsec tunnel-group will remain dormant while trying to contact its redundant peer. Groups held dormant in this manner will neither pass traffic nor negotiate security keys. This interval ends when either the redundant peer is contacted and a master election occurs, or when the maximum duration expires.

Values 1 to 1800

boot *interval-secs*

Specifies the maximum duration of an interval immediately following system startup. When the normal discovery interval for a group would expire while the post-boot discovery interval is still active, then the group's discovery interval is extended until the post-boot discovery interval expires. This allows an extension to the normal discovery stage of groups following a chassis reboot, to account for the larger variance in routing.

Values 1 to 1800

Platforms

7705 SAR Gen 2

8.111 disjointness-reference

disjointness-reference

Syntax

[no] disjointness-reference

Context

[Tree] (config>router>pcep>pcc>pce-assoc>div disjointness-reference)

Full Context

configure router pcep pcc pce-associations diversity disjointness-reference

Description

This command configures the value conveyed in the P-flag of the DISJOINTNESS-CONFIGURATION TLV. When enabled, it indicates that this LSP path is the reference path for the disjoint set of paths. The PCE must first compute the path of this LSP and then apply the requested disjointness type to compute the path of all other paths in the same diversity association ID.

The **no** form of this command sets the P-flag to false.

Default

P-flag to false

Platforms

7705 SAR Gen 2

8.112 disjointness-type

disjointness-type

Syntax

disjointness-type {**loose** | **strict**}

no disjointness-type

Context

[Tree] (config>router>pcep>pcc>pce-assoc>div disjointness-type)

Full Context

configure router pcep pcc pce-associations diversity disjointness-type

Description

This command configures the disjointness type for the association group.

The **no** form of this command reverts to the default value.

Default

disjointness-type loose

Parameters

loose

Keyword to specify the loose disjointness type.

strict

Keyword to specify the strict disjointness type.

Platforms
7705 SAR Gen 2

8.113 dispersion

dispersion

Syntax
dispersion *dispersion*

Context
[Tree] (config>port>dwdm>coherent dispersion)

Full Context
configure port dwdm coherent dispersion

Description
This command configures the residual chromatic dispersion to be compensated when the coherent receiver is operating in manual dispersion control mode.

Default
0

Parameters
dispersion
Specifies the dispersion compensation.

Values -50000 to 50000

Platforms
7705 SAR Gen 2

8.114 display

display

Syntax
display type {*type*} *url-string* **format** {*format*} [**password** [*32 chars max*]]

Context

[Tree] (admin>certificate display)

Full Context

admin certificate display

Description

This command displays the content of an input file in plain text.

The following list summarizes the formats supported by this command:

- System
  - system format
  - PKCS #12
  - PKCS #7 PEM encoded
  - PKCS #7 DER encoded
  - RFC 4945
- Certificate Request
  - PKCS #10
- Key
  - system format
  - PKCS #12
- CRL
  - system format
  - PKCS #7 PEM encoded
  - PKCS #7 DER encoded
  - RFC 4945

Parameters

url-string

Specifies the local CF card url of the input file.

Values

url-string	<local-url> [up to 99 characters]
local-url	<cflash-id>/<file-path>
cflash-id	cf1: | cf2: | cf3:

type

Specifies the type of input file.

Values cert, key, crl, cert-request

format

Specifies the format of input file.

Values pkcs10, pkcs12, pkcs7-der, pkcs7-pem, pem, der

password

Specifies the password to decrypt the input file in case that it is an encrypted PKCS#12 file, up to 99 characters.

Platforms

7705 SAR Gen 2

8.115 display-config

display-config

Syntax

display-config [**detail** | **index**]

Context

[Tree] (admin display-config)

Full Context

admin display-config

Description

This command displays the system's running configuration.

By default, only non-default settings are displayed.

Specifying the **detail** option displays all default and non-default configuration parameters.

Parameters

detail

Displays default and non-default configuration parameters.

index

Displays only persistent-indices.

Platforms

7705 SAR Gen 2

8.116 display-key

display-key

Syntax

display-key type {**ike** | **esp**} **gateway name** *name* **dynamic-tunnel** *ip-address*: *port*

display-key type {**ike** | **esp**} **tunnel** *ipsec-tunnel-name*

Context

[Tree] (admin>ipsec display-key)

Full Context

admin ipsec display-key

Description

This command displays existing IKE-SA or CHILD-SA keys..

Note:
This command does not work if **config>ipsec>no show-ipsec-keys** or **no max-history-{esp|ike}-key-records** is configured under corresponding **ipsec-gw** or **ipsec-tunnel**.

Parameters

name

The name, up to 32 characters.

ip-address

The IP address of the remote client.

Values	<ip-address>	ip-address	a.b.c.d
		ipv6-address	x:x:x:x:x:x:x:x
			x:x:x:x:x:x:d.d.d.d
			x - [0 to FFFF]H
			d - [0 to 255]D

port

The port of the remote client.

Values 0 to 65535

ipsec-tunnel-name

The IPsec tunnel name, up to 32 characters.

Platforms

7705 SAR Gen 2

Output

The following outputs are examples of the **admin ipsec display-key** command.

Output Example

```
admin ipsec display-key type ike gateway name "rw" dynamic-tunnel 11.1.1.100:500
===============================================================================
IKE-SA history: max-num-records 3 current-num-saved-records 1
                local: 172.16.100.1 remote: 11.1.1.100
record [0]: established time: 01/25/2018 20:51:55
  Initiator-SPI: d67ac71d73656496 Responder-SPI: d67ac71d73656496  Ike Version: 2
  SK_er: aes128, len: 16, val: a5da1c57f09a7eb7dbe9526cd52e2189
  SK_ar: sha1, len: 20, val: c11797bb8ebe5a1fadf46363bf5e763552bb45d0
  SK_ei: aes128, len: 16, val: 467124009cc577a8b23882a81ab9df70
  SK_ai: sha1, len: 20, val: 7dfef89bad31cb72d1ca8da2c04a9521993c7f9
```

Output Example

```
admin ipsec display-key type esp gateway name "rw" dynamic-tunnel 11.1.1.100:500

ESP-SA history: max-num-records 48 current-num-saved-records 2ynamic-tunnel 11.1.1.100:500
                local: 172.16.100.1 remote: 11.1.1.100
record [0]: established time: 01/25/2018 20:54:56
  InSpi:  154532(0x00025ba4)
          encr-alg: aes128 len: 16 val: 0xd26aa32d8bd328b1e8332fa5c7b5eeaa
          auth-alg: sha1 len: 20 val: 0x0b37ddb824a43921d3b0ee81a6910eed065a9845
  OutSpi: 3286259439(0xc3e056ef)
          encr-alg: aes128 len: 16 val: 0x3acd95376ce04fcded2e0c80cc4289cf
          alg: sha1 len: 20 val: 0x9f5a46b5cdc572972b44cdbd36b5f824ab060634
record [1]: established time: 01/25/2018 20:51:55
  InSpi:  261186(0x0003fc42)
          encr-alg: aes128 len: 16 val: 0x8bf97675d37de3e3f6e634e3e11fc3aa
          auth-alg: sha1 len: 20 val: 0xf10c0f0821488cc14f8715cc323441fc967a79dd
  OutSpi: 3246917342(0xc18806de)
          encr-alg: aes128 len: 16 val: 0xf36aaaa7a3a09734fe4fc6cd0ac9043e
          alg: sha1 len: 20 val: 0x40c13a444e4fb1d42a13812f70b17041ed0f56ee
```

8.117 dist-cpu-protection

dist-cpu-protection

Syntax

dist-cpu-protection *policy-name*

no dist-cpu-protection

Context

[Tree] (config>service>epipe>sap dist-cpu-protection)

Full Context

configure service epipe sap dist-cpu-protection

Description

This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid existing DCP policy can be assigned to a SAP or a network interface (this rule does not apply to templates, such as an **msap-policy** template).

If no dist-cpu-protection policy is assigned to a SAP, then the default access DCP policy (_default-access-policy) is used.

If no DCP functionality is required on the SAP, then an empty DCP policy can be created and explicitly assigned to the SAP.

Parameters

policy-name

Specifies the name of the DCP policy, up to 32 characters in length

Platforms

7705 SAR Gen 2

dist-cpu-protection

Syntax

dist-cpu-protection *policy-name*

no dist-cpu-protection

Context

[Tree] (config>service>vpls>sap dist-cpu-protection)

Full Context

configure service vpls sap dist-cpu-protection

Description

This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid existing DCP policy can be assigned to a SAP or a network interface (this rule does not apply to templates, such as an msap-policy template).

Default

If no dist-cpu-protection policy is assigned to a SAP, then the default access DCP policy (_default-access-policy) is used. If no DCP functionality is required on the SAP, then an empty DCP policy can be created and explicitly assigned to the SAP.

Parameters

policy-name

Specifies the name of the DCP policy, up to 32 characters in length

Platforms

7705 SAR Gen 2

dist-cpu-protection

Syntax

dist-cpu-protection *policy-name*

no dist-cpu-protection

Context

[Tree] (config>service>ies>if>sap dist-cpu-protection)

Full Context

configure service ies interface sap dist-cpu-protection

Description

This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid DCP policy can be assigned to a SAP or a network interface. This rule does not apply to templates such as an msap-policy.

Default

If no dist-cpu-protection policy is assigned to an SAP, then the default access DCP policy (default-access-policy) is used. If no DCP functionality is required on the SAP, then an empty DCP policy can be created and explicitly assigned to the SAP policy.

Parameters

policy-name

Specifies the name of the DCP policy, up to 32 characters in length

Platforms

7705 SAR Gen 2

dist-cpu-protection

Syntax

dist-cpu-protection *policy-name*

no dist-cpu-protection

Context

[Tree] (config>service>vprn>nw-if dist-cpu-protection)

Full Context

configure service vprn network-interface dist-cpu-protection

Description

This command assigns a Distributed CPU Protection (DCP) policy to the network interface. Only a valid created DCP policy can be assigned to a network interface (this rule does not apply to templates such as an msap-policy).

Default

If no dist-cpu-protection policy is assigned to the VPRN network interface, then the default network DCP policy (_default-network-policy) is used.

If no DCP functionality is required on the VPRN network interface then an empty DCP policy can be created and explicitly assigned to the VPRN network interface.

Parameters

policy-name

Specifies the name of the DCP policy, up to 32 characters in length

Platforms

7705 SAR Gen 2

dist-cpu-protection

Syntax

dist-cpu-protection *policy-name*

no dist-cpu-protection

Context

[Tree] (config>service>vprn>if>sap dist-cpu-protection)

Full Context

configure service vprn interface sap dist-cpu-protection

Description

This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid created DCP policy can be assigned to a SAP or a network interface (This rule does not apply to templates such as an msap-policy).

Default

If no dist-cpu-protection policy is assigned to an SAP policy, then the default access DCP policy (default-access-policy) is used. If no DCP functionality is required on the SAP policy, then an empty DCP policy can be created and explicitly assigned to the SAP policy.

Parameters

policy-name

Specifies the name of the DCP policy, up to 32 characters in length

Platforms

7705 SAR Gen 2

dist-cpu-protection

Syntax

dist-cpu-protection *policy-name*

no dist-cpu-protection

Context

[Tree] (config>router>if dist-cpu-protection)

Full Context

configure router interface dist-cpu-protection

Description

This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid created DCP policy can be assigned to a SAP or a network interface (note that this rule does not apply to templates such as an msap-policy).

If the user does not assign a DCP policy to a router interface, the system uses the default network DCP policy.

Default

no dist-cpu-protection

Parameters

policy-name

Specifies the name of the DCP policy, up to 32 characters in length

Platforms

7705 SAR Gen 2

dist-cpu-protection

Syntax

dist-cpu-protection

Context

[Tree] (config>system>security dist-cpu-protection)

Full Context

configure system security dist-cpu-protection

Description

Commands in this context configure the Distributed CPU Protection (DCP) feature.

Platforms

7705 SAR Gen 2

8.118 distinguisher

distinguisher

Syntax

distinguisher *id*

no distinguisher

Context

[Tree] (conf>router>segment-routing>sr-policies>policy distinguisher)

Full Context

configure router segment-routing sr-policies static-policy distinguisher

Description

This command associates a distinguisher value with a statically defined segment routing policy. This is a mandatory parameter and configuration command for non-local segment routing policies (for which the **head-end** parameter is set to a value other than "local”). Every non-local segment routing policy must have a unique distinguisher value. When a non-local static segment routing policy is imported into BGP and originated as a BGP route, the configured distinguisher value is copied into the NLRI of the route.

The **no** form of this command removes the distinguisher association.

Default

no distinguisher

Parameters

id

Specifies the distinguisher ID.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

distinguisher

Syntax

distinguisher *distinguisher-id*

no distinguisher

Context

[Tree] (config>router>policy-options>policy-statement>entry>from distinguisher)

Full Context

configure router policy-options policy-statement entry from distinguisher

Description

This command configures an SR Policy distinguisher as a route policy match criterion. This match criterion is only used in import policies.

The **no** form of this command removes the distinguisher ID match criterion from the configuration.

Parameters

distinguisher-id

Specifies the SR policy distinguisher ID.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

8.119 diversity

diversity

Syntax

[no] diversity *association-name*

Context

[Tree] (config>router>pcep>pcc>pce-assoc diversity)

Full Context

configure router pcep pcc pce-associations diversity

Description

This command creates a named diversity association from which the parameters for the specified diversity association are configured.

The **no** form of the command deletes the specified diversity association.

Parameters

association-name

Specifies the name of the diversity association, up to 32 characters.

Platforms

7705 SAR Gen 2

diversity

Syntax

[no] diversity *diversity-assoc-name*

Context

[Tree] (config>router>mpls>lsp>pce-assoc diversity)

[Tree] (config>router>mpls>lsp-template>pce-assoc diversity)

Full Context

configure router mpls lsp pce-associations diversity

configure router mpls lsp-template pce-associations diversity

Description

This command binds the LSP to a named diversity association. The diversity association must exist under the PCC. Up to five diversity associations can be configured per LSP.

The **no** form of the command removes the LSP binding from the specified diversity association.

Parameters

diversity-assoc-name

Specifies the name of an existing diversity association, up to 32 characters.

Platforms

7705 SAR Gen 2

8.120 diversity-type

diversity-type

Syntax

diversity-type {**link** | **node** | **srlg-link** | **srlg-node**}

no diversity-type

Context

[Tree] (config>router>pcep>pcc>pce-assoc>div diversity-type)

Full Context

configure router pcep pcc pce-associations diversity diversity-type

Description

This command configures the diversity type for the association group. This command is mandatory. If the command is not configured, the system does not validate the association configuration.

The **no** form of the command reverts to the default value.

Default

no diversity-type

Parameters

link

Keyword to specify the link diversity type.

node

Keyword to specify the node diversity type.

srlg-link

Keyword to specify the SRLG-link diversity type.

srlg-node

Keyword to specify the SRLG-node diversity type.

Platforms

7705 SAR Gen 2

8.121 dns

dns

Syntax

[**no**] **dns**

Context

[Tree] (config>service>vprn dns)

Full Context

configure service vprn dns

Description

Commands in this context configure domain name servers.

The **no** form of this command disables DNS for this service.

Platforms

7705 SAR Gen 2

dns

Syntax

dns

Context

[Tree] (config>router dns)

Full Context

configure router dns

Description
This command configures the DNS.

Default
dns

Platforms
7705 SAR Gen 2

dns

Syntax
dns

Context
[Tree] (config>system dns)

Full Context
configure system dns

Description
This command configures DNS settings.

Platforms
7705 SAR Gen 2

8.122 dns-domain

dns-domain

Syntax
dns-domain *dns-name*

no dns-domain

Context
[Tree] (bof dns-domain)

Full Context
bof dns-domain

Description

This command configures the domain name used when performing DNS address resolution. This is a required parameter if DNS address resolution is required. Only a single domain name can be configured. If multiple domain statements are configured, the last one encountered is used.

The **no** form of this command removes the domain name from the configuration.

Default

no dns-domain

Parameters

dns-name

Specifies the DNS domain name, up to 178 characters.

Platforms

7705 SAR Gen 2

8.123 dns-options

dns-options

Syntax

[**no**] **dns-options**

Context

[Tree] (config>service>vprn>router-advert>if dns-options)

[Tree] (config>service>vprn>router-advert dns-options)

Full Context

configure service vprn router-advertisement interface dns-options

configure service vprn router-advertisement dns-options

Description

Commands in this context configure DNS information for Stateless Address Auto-Configuration (SLAAC) hosts.

When specified at the router-advertisement level in the routing context, this command allows configuration of service-wide parameters. These can then be inherited at the interface level by specifying the **config>service>vprn>router-advert>if>dns-options>include-dns** command.

The **no** form of this command disables configuration of DNS information for Stateless Address Auto-Configuration (SLAAC) hosts.

Platforms

7705 SAR Gen 2

dns-options

Syntax

[**no**] **dns-options**

Context

[Tree] (config>router>router-advert>if dns-options)

[Tree] (config>router>router-advert dns-options)

Full Context

configure router router-advertisement interface dns-options

configure router router-advertisement dns-options

Description

Commands in this context configure DNS information for Stateless Address Auto-Configuration (SLAAC) hosts. When specified at the router-advertisement level in the routing context, this command allows configuration of service-wide parameters. These can then be inherited at the interface level by specifying the **config>router>router-advert>if>dns-options>include-dns** command.

The **no** form of this command disables configuration of DNS information for Stateless Address Auto-Configuration (SLAAC) hosts.

Platforms

7705 SAR Gen 2

8.124 dns-server

dns-server

Syntax

dns-server *ip-address* [*ip-address*]

no dns-server

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options dns-server)

[Tree] (config>service>vprn>dhcp>server>pool>options dns-server)

[Tree] (config>router>dhcp>server>pool>options dns-server)

Full Context

configure subscriber-mgmt local-user-db ipoe host options dns-server

configure service vprn dhcp local-dhcp-server pool options dns-server

configure router dhcp local-dhcp-server pool options dns-server

Description

This command configures the IPv4 address of the DNS server.

The **no** form of this command removes the IPv4 address of the DNS server from the configuration.

Parameters

ip-address

Specifies up to four DNS server IP addresses.

Platforms

7705 SAR Gen 2

dns-server

Syntax

dns-server *ipv6-address* [*ipv6-address*]

no dns-server

Context

[Tree] (config>router>dhcp6>server>pool>options dns-server)

[Tree] (config>service>vprn>dhcp6>server>pool>options dns-server)

[Tree] (config>service>vprn>dhcp6>server>pool>prefix>options dns-server)

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options6 dns-server)

[Tree] (config>router>dhcp6>server>pool>prefix>options dns-server)

Full Context

configure router dhcp6 local-dhcp-server pool options dns-server

configure service vprn dhcp6 local-dhcp-server pool options dns-server

configure service vprn dhcp6 local-dhcp-server pool prefix options dns-server

configure subscriber-mgmt local-user-db ipoe host options6 dns-server

configure router dhcp6 local-dhcp-server pool prefix options dns-server

Description

This command configures IPv6 DNS server addresses that can be used for name resolution.

The **no** form of this command removes the IPv6 address of the DNS server from the configuration.

Parameters

ipv6-address

Specifies up to four IPv6 DNS server addresses.

Values ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x - [0 to FFFF]H
d - [0 to 255]D

Platforms

7705 SAR Gen 2

8.125 dnssec

dnssec

Syntax

dnssec

Context

[Tree] (config>system>dns dnssec)

Full Context

configure system dns dnssec

Description

This command configures system Domain Name System Security Extensions (DNSSEC) settings.

Platforms

7705 SAR Gen 2

8.126 do-not-fragment

do-not-fragment

Syntax

[**no**] **do-not-fragment**

Context

[Tree] (config>oam-pm>session>ip do-not-fragment)

Full Context

configure oam-pm session ip do-not-fragment

Description

This command configures the Do Not Fragment (DF) bit in the IPv4 header of the TWAMP Light test packet in order to prevent packet fragmentation. This is only applicable to IPv4. IPv6 does not include the bit as part of the specification. This parameter is ignored but not blocked when the address is IPv6.

The **no** form of this command allows packet fragmentation.

Platforms

7705 SAR Gen 2

8.127 dod-label-distribution

dod-label-distribution

Syntax

[**no**] **dod-label-distribution**

Context

[Tree] (config>router>ldp>session-params>peer dod-label-distribution)

Full Context

configure router ldp session-parameters peer dod-label-distribution

Description

This command enables the use of the LDP Downstream-on-Demand (DoD) label distribution procedures.

When this option is enabled, LDP will set the A-bit in the Label Initialization message when the LDP session to the peer is established. When both peers set the A-bit, they will both use the DoD label distribution method over the LDP session (RFC 5036).

This feature can only be enabled on a link-level LDP session and therefore will apply to prefix labels only, not service labels.

As soon as the link LDP session comes up, the router will send a label request to its DoD peer for the FEC prefix corresponding to the peer's LSR-id. The DoD peer LSR-id is found in the basic Hello discovery messages the peer used to establish the Hello adjacency with the router.

Similarly if the router and the directly attached DoD peer entered into extended discovery and established a targeted LDP session, the router will immediately send a label request for the FEC prefix corresponding to the peer's LSR-id found in the extended discovery messages.

However, the router will not advertise any <FEC, label> bindings, including the FEC of its own LSR-id, unless the DoD peer requested it using a Label Request Message.

When the DoD peer sends a label request for any FEC prefix, the router will reply with a <FEC, label> binding for that prefix if the FEC was already activated on the router. If not, the router replies with a notification message containing the status code of "no route." The router will not attempt in the latter case to send a label request to the next-hop for the FEC prefix when the LDP session to this next-hop uses the DoD label distribution mode. Hence the reference to single-hop LDP DoD procedures.

As soon as the link LDP session comes up, the router will send a label request to its DoD peer for the FEC prefix corresponding to the peer's LSR-id. The DoD peer LSR-id is found in the basic Hello discovery messages the peer used to establish the Hello adjacency with the router.

Similarly if the router and the directly attached DoD peer entered into extended discovery and established a targeted LDP session, the router immediately sends a label request for the FEC prefix corresponding to the peer's LSR-id found in the extended discovery messages. Peer address has to be the peer LSR-ID address.

The **no** form of this command disables the DoD label distribution with an LDP neighbor.

Default

no dod-label-distribution

Platforms

7705 SAR Gen 2

8.128 domain

domain

Syntax

domain [*value*] [**create**]

no domain [*value*]

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ipsec domain)

Full Context

configure redundancy multi-chassis peer mc-ipsec domain

Description

This command configures domain information. This command is mutually exclusive to the **tunnel-group** command.

The **no** form of this command removes the multi-chassis IPsec domain value.

Parameters

value

Specifies the domain multi-chassis IPsec domain, up to 255 characters.

create

Keyword used to create the command instance.

Platforms

7705 SAR Gen 2

8.129 domain-id

domain-id

Syntax

domain-id *global-field*:*local-field*

no domain-id

Context

[Tree] (config>service>epipe>bgp-evpn>mpls domain-id)

[Tree] (config>service>vprn>bgp-ipvpn>mpls domain-id)

[Tree] (config>service>vprn>bgp-evpn>mpls domain-id)

Full Context

configure service epipe bgp-evpn mpls domain-id

configure service vprn bgp-ipvpn mpls domain-id

configure service vprn bgp-evpn mpls domain-id

Description

This command specifies the domain ID that identifies the network from which a BGP route was received before that route is exported to a different neighbor. The domain ID is part of a domain, represented as *domain-id*:*isf_safi_type* in the D-PATH attribute, as described in *draft-ietf-bess-evpn-ipvpn-interworking*. The D-PATH attribute is modified by gateway routers, where a gateway is defined as a PE where a VPRN is instantiated, and that VPRN advertises or receives routes from multiple BGP owners (for example, EVPN-IFL and BGP-IPVPN) or multiple instances of the same owner (for example, VPRN with two BGP-IPVPN instances).

In the following example, consider that a gateway receives prefix P in an EVPN-IFL instance with the following D-PATH from neighbor N:

Seg Len=1 / 65000:1:128

If the router imports the route in VPRN-1, BGP-EVPN SRv6 instance with domain 65000:2, it readvertises it to its BGP-IPVPN MPLS instance as follows:

Seg Len=2 / 65000:2:70 / 65000:1:128

That is, the gateway prepends the local domain ID and family to the D-PATH before readvertising the route into a different instance.

The D-PATH attribute is used on gateways to detect loops (for received routes where the D-PATH contains a local domain ID) and to make BGP best-path selection decisions based on the D-PATH length (shorter D-PATH is preferred).

The command is also supported in Epipe services with two instances. As in the case of multi-instance VPRN services, the configured domain ID in an Epipe instance is prepended to the AD per EVI route redistributed to the other instance.

The **no** form of this command removes the configured domain ID.

Default

no domain-id

Parameters

global-field:local-field

Specifies the domain ID.

Values *4byte-GlobalAdminValue:2byte-LocalAdminValue*

4byte-GlobalAdminValue: 0 to 4294967295

2byte-LocalAdminValue 0 to 65535

Platforms

7705 SAR Gen 2

8.130 domain-name

domain-name

Syntax

domain-name *domain-name*

no domain-name

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options domain-name)

[Tree] (config>router>dhcp>server>pool>options domain-name)

[Tree] (config>service>vprn>dhcp6>server>pool>options domain-name)

[Tree] (config>router>dhcp6>server>pool>options domain-name)

[Tree] (config>service>vprn>dhcp>server>pool>options domain-name)

[Tree] (config>router>dhcp6>server>pool>prefix>options domain-name)

[Tree] (config>service>vprn>dhcp6>server>pool>prefix>options domain-name)

Full Context

configure subscriber-mgmt local-user-db ipoe host options domain-name

configure router dhcp local-dhcp-server pool options domain-name

configure service vprn dhcp6 local-dhcp-server pool options domain-name

configure router dhcp6 local-dhcp-server pool options domain-name

configure service vprn dhcp local-dhcp-server pool options domain-name

configure router dhcp6 local-dhcp-server pool prefix options domain-name

configure service vprn dhcp6 local-dhcp-server pool prefix options domain-name

Description

This command configures the default domain for a DHCP client that the router uses to complete unqualified host names (without a dotted-decimal domain name).

The **no** form of this command removes the name from the configuration.

Parameters

domain-name

Specifies the domain name for the client, up to 127 characters.

Platforms

7705 SAR Gen 2

8.131 dot1p

dot1p

Syntax

dot1p *dot1p-priority* [**fc** *fc-name*] [**priority** {**low** | **high**}]

no dot1p *dot1p-priority*

Context

[Tree] (config>qos>sap-ingress dot1p)

Full Context

configure qos sap-ingress dot1p

Description

This command explicitly sets the forwarding class or subclass or enqueuing priority when a packet is marked with a *dot1p-priority* specified. Adding a dot1p rule on the policy forces packets that match the *dot1p-priority* specified to override the forwarding class and enqueuing priority based on the parameters included in the dot1p rule. When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueuing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueuing priority derived from earlier matches in the classification hierarchy.

The *dot1p-priority* is derived from the most significant three bits in the IEEE 802.1q or IEEE 802.1p header. The three dot1p bits define 8 Class-of-Service (CoS) values commonly used to map packets to per-hop QoS behavior.

The **no** form of this command removes the explicit dot1p classification rule from the SAP ingress policy. Removing the rule on the policy immediately removes the rule on all ingress SAPs using the policy.

Parameters

dot1p-priority

This value is a required parameter that specifies the unique IEEE 802.1p value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-value*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight dot1p rules are allowed on a single policy.

Values 0 to 7

fc *fc-name*

Specifies the value given for the *fc-name* parameter must be one of the predefined forwarding classes in the system. Specifying the fc-name is optional. When a packet matches the rule, the forwarding class is only overridden when the **fc** *fc-name* parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

priority

This parameter overrides the default enqueuing priority for all packets received on an ingress SAP using this policy that match this rule. Specifying the priority is optional. When a packet matches the rule, the enqueuing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and priority is not explicitly defined in the rule, the enqueuing priority is inherited based on previous rule matches.

Default Inherits the priority defined by the default-priority statement.

high

This parameter is used in conjunction with the **priority** parameter. Setting the enqueuing parameter to **high** for a packet increases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

low

This parameter is used in conjunction with the **priority** parameter. Setting the enqueuing parameter to **low** for a packet decreases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

Platforms

7705 SAR Gen 2

dot1p

Syntax

dot1p *dot1p-value* [*dot1p-mask*]

no dot1p

Context

[Tree] (config>qos>sap-ingress>mac-criteria>entry>match dot1p)

Full Context

configure qos sap-ingress mac-criteria entry match dot1p

Description

The IEEE 802.1p value to be used as the match criterion.

Use the **no** form of this command to remove the dot1p value as the match criterion.

Default

no dot1p

Parameters

dot1p-value

Specifies the IEEE 802.1p value in decimal.

Values 0 to 7

dot1pmask

This 3-bit mask can be configured using the following formats.

Table 23: Format Styles to Configure Mask

Format Style	Format Syntax	Example
Decimal	D	4
Hexadecimal	0xH	0x4

Format Style	Format Syntax	Example
Binary	0bBBB	0b100

To select a range from 4 up to 7, specify *p-value* of 4 and a *mask* of 0b100 for value and mask.

Values 0 to 7 (decimal hex or binary)

Default 7

Platforms

7705 SAR Gen 2

dot1p

Syntax

dot1p *dot1p-value* [**fc** *fc-name*] [**profile** {**in** | **out** | **use-de** | **exceed** | **inplus**}]

no dot1p *dot1p-value*

Context

[Tree] (config>qos>sap-egress dot1p)

Full Context

configure qos sap-egress dot1p

Description

This command defines a specific dot1p value that must be matched to perform the associated reclassification actions. If an egress packet on the SAP matches the specified dot1p value, the forwarding class or profile may be overridden. By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions.

The dot1p priority is derived from the most significant three bits in the IEEE 802.1q or IEEE 802.1p header. The three dot1p bits define 8 Class-of-Service (CoS) values commonly used to map packets to per-hop QoS behavior.

The reclassification actions from a dot1p reclassification rule may be overridden by a DSCP, IP precedence, or IP flow matching event.

The **fc** keyword is optional. When specified, the egress classification rule will overwrite the forwarding class derived from ingress. The new forwarding class is used for egress remarking and queue mapping decisions. If a DSCP, IP precedence, IPv6 criteria, or IP criteria match occurs after the dot1p match, the new forwarding class may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new FC, the FC from the dot1p match will be used.

The **profile** keyword is optional. When specified, the egress classification rule will overwrite the profile of the packet derived from ingress. The new profile value is used for egress remarking and queue congestion behavior. If a DSCP, IP precedence, IPv6 criteria, or IP criteria match occurs after the dot1p match, the

new profile may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new profile, the profile from the dot1p match will be used.

The **no** form of this command removes the reclassification rule from the SAP egress QoS policy.

Parameters

dot1p-value

This value is a required parameter that specifies the unique IEEE 802.1p value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-value*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight dot1p rules are allowed on a single policy.

Values 0 to 7

fc *fc-name*

Specifies the value given for the *fc-name* parameter must be one of the predefined forwarding classes in the system. Specifying the FC name is optional. When a packet matches the rule, the forwarding class is only overridden when the **fc** *fc-name* parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

profile {in | out | use-de | exceed | inplus}

Specifies the profile reclassification action is optional. When specified, packets matching the dot1p value will be explicitly reclassified to the profile specified regardless of the ingress profiling decision. The explicit profile reclassification may be overwritten by a DSCP, IP precedence, IPv6 criteria, or IP criteria reclassification match. To remove the profile reclassification action for the specified dotp1 value, the **dot1p** command must be re-executed without the profile reclassification action defined.

Values **in** — Specifies that any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.

out — Specifies that any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.

use-de — Specifies that the DE bit is used to determine the profile of the packet (in-profile is used when DE = 0 and out-of-profile is used when DE = 1).

exceed — Specifies that any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.

inplus — Specifies that any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

Platforms

7705 SAR Gen 2

dot1p

Syntax

dot1p {*dot1p-value* | **in-profile** *dot1p-value* **out-profile** *dot1p-value* [**exceed-profile** *dot1p-value*]}

no dot1p

Context

[Tree] (config>qos>sap-egress>fc dot1p)

Full Context

configure qos sap-egress fc dot1p

Description

This command explicitly defines the egress IEEE 802.1p (dot1p) bits marking for *fc-name*. When the marking is set, all packets of *fc-name* that have either an IEEE 802.1q or IEEE 802.1p encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1q or IEEE 802.1p encapsulated, the dot1p command has no effect.

The optional **in-profile** *dot1p-value* **out-profile** *dot1p-value* [**exceed-profile** *dot1p-value*] parameters added to the existing **dot1p** command adds the capability to mark on an egress SAP the in, out, and exceed-profile status via a certain dot1p combination, similarly with the DE options. All inplus-profile traffic is marked with the same value as in-profile traffic.

When the **in-profile** keyword is added, the **out-profile** keyword must be specified; however, **exceed-profile** is optional. If the optional **exceed-profile** *dot1p-value* is not included, any exceed-profile traffic will be marked with the same dot1p value as configured for the out-of-profile traffic.

The command with the additional structure may be used on the SAP when the internal in, out, and exceed-profile status needs to be communicated to an access network or customer device that does not support the DE bit.

When these commands are used, the DE bit or the equivalent field is left unchanged by the egress processing if a tag exists. If a new tag is added, the related DE bit is set to 0.

When the previous command (**dot1p** *dot1p-value*) is used without the new structure, it means that the dot1p value is used for the entire forwarding class, as it did before. The two versions of the command are mutually exclusive.

The in-profile or out-of-profile/exceed-profile status may be indicated via the DE bit setting if the **de-mark** command is used. The DE value used for exceed-profile is the same as that used for out-of-profile.

In the PBB case, for a Backbone SAP (B-SAP) and for packets originated from a local I-VPLS/PBB-Epipe, the command dictates the marking of the dot1p bits for both the BVID and ITAG.

The commands **dot1p-inner** and **dot1p-outer** take precedence over the **dot1p** command if both are specified in the same policy.

The **no** form of this command sets the IEEE 802.1p or IEEE 802.1q priority bits to 0.

Default

no dot1p

Parameters

in-profile *dot1p-value*

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

out-profile *dot1p-value*

Specifies the 802.1p value to set for out-profile frames in this forwarding class.

Values 0 to 7

exceed-profile *dot1p-value*

Specifies the 802.1p value to set for exceed-profile frames in this forwarding class.

Values 0 to 7

Platforms

7705 SAR Gen 2

dot1p

Syntax

dot1p *dot1p-priority* **fc** *fc-name* **profile** {**in** | **out** | **use-de**}

no dot1p

Context

[Tree] (config>qos>network>ingress dot1p)

Full Context

configure qos network ingress dot1p

Description

This command explicitly sets the forwarding class or enqueuing priority and profile of the packet when a packet is marked with a *dot1p-priority* specified. Adding a dot1p rule on the policy forces packets that match the *dot1p-priority* specified to override and be assigned to the forwarding class and enqueuing priority and profile of the packet, based on the parameters included in the dot1p rule. When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueuing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueuing priority derived from earlier matches in the classification hierarchy.

The *dot1p-priority* is derived from the most significant three bits in the IEEE 802.1q or IEEE 802.1p header. The three dot1p bits define 8 Class-of-Service (CoS) values commonly used to map packets to per-hop Quality of Service (QoS) behavior.

The **no** form of this command removes the explicit dot1p classification rule from the policy. Removing the rule on the policy immediately removes the rule on all ingress SAPs using the policy.

Parameters

dot1p-priority

This value is a required parameter that specifies the unique IEEE 802.1p value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-value*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight dot1p rules are allowed on a single policy.

Values 0 to 7

fc-name

Specifies the value given for the *fc-name* parameter must be one of the predefined forwarding classes in the system. Specifying the *fc-name* is optional. When a packet matches the rule, the forwarding class is only overridden when the **fc** *fc-name* parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

Values be, l2, af, l1, h2, ef, h1, nc

profile {in | out | use-de}

All packets that are assigned to this forwarding class will be considered in-profile or out-of-profile based on this command or will use the DE bit to determine the profile of the packets (in-profile is used when DE = 0 and out-of-profile is used when DE = 1). In case of congestion, the in-profile packets are preferentially queued over the out-of-profile packets.

Platforms

7705 SAR Gen 2

dot1p

Syntax

dot1p *dot1p-priority*

no dot1p

Context

[Tree] (config>qos>network>egress>fc dot1p)

Full Context

configure qos network egress fc dot1p

Description

This command is used whenever the dot1p bits are set to a common value regardless of the internal profile of the packets. Although it is not mandatory, this command should be used in combination with the **de-mark** command to enable the marking of the DE bit according to the internal profile of the packet.

This command acts as a shortcut for configuring the two existing commands with the same dot1p priority.

The **dot1p** *dot1p-priority* command is saved in the configuration as **dot1p-in-profile** *dot1p-priority* and **dot1p-out-profile** *dot1p-priority*. The inplus-profile traffic is marked with the same value as in-profile traffic. The exceed-profile traffic is marked with the same value as out-of-profile traffic.

Platforms

7705 SAR Gen 2

dot1p

Syntax

dot1p *dot1p-value* [*dot1p-mask*]

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry>match dot1p)

Full Context

configure system security management-access-filter mac-filter entry match dot1p

Description

This command configures Dot1p match conditions.

Table 24: Management Access Filter dot1p Mask Format

Format Style	Format Syntax	Example
Decimal	D	4
Hexadecimal	0xH	0x4
Binary	0bBBB	0b100

Parameters

dot1p-value

Specifies the IEEE 802.1p value in decimal.

Values 0 to 7

mask

Specifies the 3-bit mask can be configured using the following formats.

Platforms

7705 SAR Gen 2

8.132 dot1p-in-profile

dot1p-in-profile

Syntax

dot1p-in-profile *dot1p-priority*

no dot1p-in-profile

Context

[Tree] (config>qos>network>egress>fc dot1p-in-profile)

Full Context

configure qos network egress fc dot1p-in-profile

Description

This command specifies dot1p in-profile mappings. The inplus-profile traffic is marked with the same value as in-profile traffic.

The **no** form of this command resets the configuration to the default in-profile *dot1p-priority* setting for *policy-id* 1.

Parameters

dot1p-priority

Specifies the unique IEEE 802.1p value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-value*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight dot1p rules are allowed on a single policy.

Values 0 to 7

Platforms

7705 SAR Gen 2

8.133 dot1p-inner

dot1p-inner

Syntax

dot1p-inner *dot1p-value*

dot1p-inner in-profile*dot1p-value* **out-profile** *dot1p-value* [**exceed-profile** *dot1p-value*]

no dot1p-inner

Context

[Tree] (config>qos>sap-egress>fc dot1p-inner)

Full Context

configure qos sap-egress fc dot1p-inner

Description

This command explicitly configures the egress inner VLAN tag IEEE 802.1p (dot1p) bits marking for the forwarding class name. When the marking is set, all packets of the forwarding class name that have either an inner IEEE 802.1q or IEEE 802.1p encapsulation on a QinQ SAP will use the explicitly defined *dot1p-value*. If the egress packets for the forwarding class are not IEEE 802.1q or IEEE 802.1p QinQ encapsulated, this command has no effect.

The optional **in-profile** *dot1p-value*, **out-profile** *dot1p-value*, and **exceed-profile** *dot1p-value* parameters on the **dot1p-inner** command add the capability to mark the in-profile and out-of-profile status on an egress QinQ SAP. The command with the additional parameters may be used on the SAP when the internal in-profile, out-of-profile, and exceed-profile status needs to be communicated to an access network or customer device that does not support the DE bit. When the in-profile keyword is added, the rest of the structure must be specified. All inplus-profile traffic is marked with the same value as in-profile traffic.

When these commands are used, the DE bit or the equivalent field is left unchanged by the egress processing if an inner tag exists. If a new inner tag is added, the related DE bit is set to 0. The inplus/in, out, or exceed-profile status may be indicated using the DE bit setting if the **de-mark** or **de-mark-inner** command is used.

The two versions of the command (with and without parameters) are mutually exclusive.

This command takes precedence over the **configure qos sap-ingress dot1p** command if both are specified in the same policy, and over the default action where the marking is taken from a packet received at ingress.

The configuration of **qinq-mark-top-only** under the SAP egress takes precedence over the use of the **dot1p-inner** command in the policy; that is, the inner VLAN tag is not remarked when **qinq-mark-top-only** is configured. The marking used for the inner VLAN tag is based on the current default, which is governed by the marking of the packet received at the ingress to the system.

The **no** form of this command sets the inner IEEE 802.1p or IEEE 802.1q priority bits to 0.

Default

no dot1p-inner

Parameters

dot1p-value

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

in-profile *dot1p-value*

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

out-profile *dot1p-value*

Specifies the 802.1p value to set for out-of-profile frames in this forwarding class.

Values 0 to 7

exceed-profile *dot1p-value*

Specifies the 802.1p value to set for exceed-profile frames in this forwarding class.

Values 0 to 7

Platforms

7705 SAR Gen 2

8.134 dot1p-out-profile

dot1p-out-profile

Syntax

dot1p-out-profile *dot1p-priority*

no dot1p-out-profile

Context

[Tree] (config>qos>network>egress>fc dot1p-out-profile)

Full Context

configure qos network egress fc dot1p-out-profile

Description

This command specifies dot1p out-of-profile mappings.

The exceed-profile traffic is marked with the same value as out-of-profile traffic.

The **no** form of this command resets the configuration to the default out-profile *dot1p-priority* setting for *policy-id* 1.

Parameters

dot1p-priority

Specifies the unique IEEE 802.1p value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-value*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight dot1p rules are allowed on a single policy.

Values 0 to 7

Platforms

7705 SAR Gen 2

8.135 dot1p-outer

dot1p-outer

Syntax

dot1p-outer {*dot1p-value* | **in-profile** *dot1p-value* **out-profile** *dot1p-value* [**exceed-profile** *dot1p-value*]}

no dot1p-outer

Context

[Tree] (config>qos>sap-egress>fc dot1p-outer)

Full Context

configure qos sap-egress fc dot1p-outer

Description

This command explicitly defines the egress outer or single VLAN tag IEEE 802.1p (dot1p) bits marking for *fc-name*. When the marking is set, all packets of fc-name that have either an outer or single IEEE 802.1q or IEEE 802.1p encapsulation on a qinq or a dot1p SAP, respectively, will use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1q or IEEE 802.1p encapsulated, this command has no effect.

The optional **in-profile** *dot1p-value* **out-profile** *dot1p-value* [**exceed-profile** *dot1p-value*] parameters on the dot1p-outer command add the capability to mark the in, out, and exceed-profile status on an egress qinq or dot1p SAP. The command with the additional parameters may be used on the SAP when the internal in, out, and exceed-profile status needs to be communicated to an access network or customer device that does not support the DE bit.

When the **in-profile** keyword is added, the **out-profile** keyword must be specified; however, **exceed-profile** is optional. If the optional **exceed-profile** *dot1p-value* is not included, any exceed-profile traffic will be marked with the same dot1p value as configured for the out-of-profile traffic. All inplus-profile traffic is marked with the same value as in-profile traffic.

When these commands are used, the DE bit or the equivalent field is left unchanged by the egress processing if a single or outer tag exists. If a new tag is added, the related DE bit is set to 0. The in, out, or exceed-profile status may be indicated via the setting of the DE bit setting if the **de-mark(-outer)** command is used. The DE value used for inplus is the same as that used for in-profile and the one used for exceed-profile is the same as that used for out of profile.

In the PBB case, for a Backbone SAP (B-SAP) and for packets originated from a local I-VPLS/PBB-Epipe, the command dictates the marking of the dot1p bits for both the BVID and ITAG.

The two versions of the command (with and without parameters) are mutually exclusive.

This command takes precedence over the **dot1p** command if both are specified in the same policy, and over the default action where the marking is taken from a packet received at ingress.

The **no** form of the command sets the IEEE 802.1p or IEEE 802.1q priority bits to 0.

Default

no dot1p-outer

Parameters

dot1p-value

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

in-profile *dot1p-value*

Specifies the 802.1p value to set for in-profile frames in this forwarding class.

Values 0 to 7

out-profile *dot1p-value*

Specifies the 802.1p value to set for out-of-profile frames in this forwarding class.

Values 0 to 7

exceed-profile *dot1p-value*

Specifies the 802.1p value to set for exceed-profile frames in this forwarding class.

Values 0 to 7

Platforms

7705 SAR Gen 2

8.136 dot1q-etype

dot1q-etype

Syntax

dot1q-etype *value*

no dot1q-etype

Context

[Tree] (config>port>ethernet dot1q-etype)

Full Context

configure port ethernet dot1q-etype

Description

This command specifies the Ethertype expected when the port's encapsulation type is dot1q. Dot1q encapsulation is supported only on Ethernet interfaces.

The **no** form of this command reverts to the default value.

Parameters

value

Specifies the Ethertype to expect, in either decimal or hex.

Values 1536 to 65535 (0x0600 to 0xffff)

Default If the encap-type is dot1p, then the default is 0x8100. If the encap-type is qinq, then the default is 0x8100.

Platforms

7705 SAR Gen 2

8.137 dot1x

dot1x

Syntax

dot1x

Context

[Tree] (config>port>ethernet dot1x)

Full Context

configure port ethernet dot1x

Description

This command enables access to the context to configure port-specific 802.1x authentication attributes. This context can only be used when configuring a Fast Ethernet, Gigabit or 10-Gb Ethernet LAN ports on an appropriate MDA.

Platforms

7705 SAR Gen 2

8.138 down

down

Syntax

down ip *seconds* [**init-only**]

no down ip

down ipv6 *seconds* [**init-only**]

no down ipv6

Context

[Tree] (config>service>ies>if>hold-time down)

[Tree] (config>router>if>hold-time down)

[Tree] (config>service>vprn>nw-if>hold-time down)

[Tree] (config>service>vprn>if>hold-time down)

[Tree] (config>service>vpls>if>hold-time down)

Full Context

configure service ies interface hold-time down

configure router interface hold-time down

configure service vprn network-interface hold-time down

configure service vprn interface hold-time down

configure service vpls interface hold-time down

Description

This command causes a delay in the activation of the associated IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface up, unless the **init-only** option is configured. If the **init-only** option is configured, the delay is only applied when the IP interface is first configured or after a system reboot.

The **no** form of this command removes the command from the active configuration and removes the delay in activating the associated IP interface. If the configuration is removed during a delay period, the currently running delay will continue until it completes.

Default

no down ip

Parameters

seconds

The time delay, in seconds, to make the interface operational.

Values 1 to 1200

init-only

Specifies that the **down** delay is only applied when the interface is configured or after a reboot.

Values 1 to 1200

Platforms

7705 SAR Gen 2

8.139 down-on-internal-error

down-on-internal-error

Syntax

down-on-internal-error [**tx-disable**]

no down-on-internal-error

Context

[Tree] (config>port>ethernet down-on-internal-error)

Full Context

configure port ethernet down-on-internal-error

Description

This command configures the system to bring a port operationally down in the event the system has detected internal MAC transmit errors (Int MAC Tx Errs).

Default

no down-on-internal-error

Parameters

tx-disable

Specifies that the laser should be disabled if an internal MAC transmit error is encountered. When used, this option requires that the operator explicitly cycle the admin state of the port to clear the error and re-enable the laser.

Platforms

7705 SAR Gen 2

8.140 down-timeout

down-timeout

Syntax

[**no**] **down-timeout**

Context

[Tree] (config>aaa>radius-srv-plcy>servers>health-check down-timeout)

Full Context

configure aaa radius-server-policy servers health-check down-timeout

Description

This command determines the interval to wait for a RADIUS reply message from the RADIUS server before a RADIUS server is declared out-of-service. By default, the value of the down-timeout is the number of retries multiplied by the timeout interval. Each host will use the configured timeout and retry value under the AAA RADIUS server policy.

timeout refers to the waiting period before the next retry attempt.

retry refers the number of times the host will attempt to contact the RADIUS server.

If a RADIUS server is declared out-of-service, the host pending retry attempts will move on to the next RADIUS server.

The **no** form of this command reverts to the default.

Parameters

minutes

Specifies the timer to wait, in minutes, before declaring the RADIUS server that is down.

Values 1 to 5

seconds

Specifies the timer to wait, in seconds, before declaring the RADIUS server that is down.

Values 1 to 59

Platforms

7705 SAR Gen 2

8.141 downstream-ip-filter

downstream-ip-filter

Syntax

downstream-ip-filter *filter-id*

no downstream-ip-filter

Context

[Tree] (config>router>nat>outside downstream-ip-filter)

[Tree] (config>service>vprn>nat>outside downstream-ip-filter)

Full Context

configure router nat outside downstream-ip-filter

configure service vprn nat outside downstream-ip-filter

Description

This command specifies a filter to apply to the downstream traffic after routing in the outside virtual router instance and before the NAT function; it is useful for traffic that bypasses the egress filters applied in the inside virtual router instance, such as DS-Lite traffic.

The **no** form of the command removes the filter from the configuration.

Default

no downstream-ip-filter

Parameters

filter-id

Specifies a filter up to 64 characters.

Platforms

7705 SAR Gen 2

8.142 dpd

dpd

Syntax

dpd [**interval** *interval*] [**max-retries** *max-retries*] [**reply-only**]

no dpd

Context

[Tree] (config>ipsec>ike-policy dpd)

Full Context

configure ipsec ike-policy dpd

Description

This command controls the dead peer detection mechanism.

The **no** form of this command removes the parameters from the configuration.

Default

no dpd

Parameters

interval

Specifies the DPD interval, in seconds. Since more time is necessary to determine if there is incoming traffic, the actual time needed to bring down the tunnel is larger than the DPD interval multiplied by max-retries.

Values 10 to 300

Default 30

max-retries

Specifies the maximum number of retries before the tunnel is removed.

Values 2 to 5

Default 3

reply-only

Specifies whether to initiate a DPD request if there is an incoming ESP or IKE packet. Issuing the command without the reply-only keyword does not initiate a DPD request if there is an incoming ESP packet.

Platforms

7705 SAR Gen 2

8.143 drain

drain

Syntax

[**no**] **drain**

Context

[Tree] (config>service>vprn>dhcp>server>pool>subnet drain)

Full Context

configure service vprn dhcp local-dhcp-server pool subnet drain

Description

This command means no new leases can be assigned from this subnet and existing leases are cleaned up upon renew/rebind.

The **no** form of this command means the subnet is active and new leases can be assigned from it.

Platforms

7705 SAR Gen 2

drain

Syntax

[**no**] **drain**

Context

[Tree] (config>service>vprn>nat>outside>pool>address-range drain)

[Tree] (config>router>nat>outside>pool>address-range drain)

Full Context

configure service vprn nat outside pool address-range drain

configure router nat outside pool address-range drain

Description

This command starts or stops draining this NAT address range. When an address-range is being drained, it will not be used to serve new hosts. Existing hosts, however, will still be able to use the address that was assigned to them even if it is being drained. An address-range can only be deleted if the parent pool is shut down or if the range itself is effectively drained (hosts are no longer using the addresses).

Platforms

7705 SAR Gen 2

8.144 drop

drop

Syntax

drop

drop packet-length {**lt** | **gt** | **eq**} *packet-length-value*

drop packet-length range *packet-length-value packet-length-value*

drop pattern expression *expression* **mask** *mask* **offset-type** *offset-type* **offset-value** *offset-value*

drop ttl {**lt** | **gt** | **eq**} *ttl-value*

drop ttl range *ttl-value ttl-value*

Context

[Tree] (config>filter>ip-filter>entry>action drop)

Full Context

configure filter ip-filter entry action drop

Description

This command configures the drop action for the traffic that matches this filter entry.

Traffic can, also, be dropped based on *pkt-length*, *packet-length range*, *ttl*, *ttl range*, or a pattern of conditional match criteria.

Packets that match the filter entry match criteria, and not the conditional match criteria value, are implicitly forwarded with no further match in the following filter entries.

For pattern match:

- the expression is left-aligned for odd number bytes, for example, the expression 0xABC is programmed 0x0ABC in the line card
- the 'data' offset requires protocol UDP or TCP to be selected in the filter entry match criteria.

Parameters

packet-length

Specifies drop packets matching both the filter entry match criteria and the *packet-length value* defined in the **drop** action statement. Packets matching the filter entry match criteria and not matching the *packet-length* value, as defined in the **drop** action statement, are implicitly forwarded with no further match in the following filter entries.

Values lt — Specifies "less than". The **lt** parameter cannot be used with the lowest possible numerical value for the parameter.

gt — Specifies "greater than”. The **gt** parameter cannot be used with the highest possible numerical value for the parameter.

eq — Specifies "equal to”.

packet-length-value

Specifies the packet length value for the rate limit action.

Values 0 to 65535

range

Specifies an inclusive range. When **range** is used, the start of the range (the first value entered) must be smaller than the end of the range (the second value entered).

expression

Specifies the hexadecimal pattern to match; up to eight bytes.

Values 0x0000000000000001 to 0xffffffffffffffff

mask

Specifies the mask for the pattern expression, up to eight bytes.

Values 0x0000000000000001 to 0xffffffffffffffff

offset-type

Specifies the starting point reference for the offset-value of this pattern.

Values layer-3, layer-4, data, dns-qtype

offset-value

Specifies the offset value for the pattern expression. Dns-qtype supports offset value of 0.

Values 0 to 255

ttl-value

Specifies drop packets matching both the filter entry match criteria and the TTL value defined in the drop action statement. Packets matching the filter entry match criteria and not matching the TTL value, as defined in the drop action statement, are implicitly forwarded with no further match in the following filter entries.

Values 0 to 255

Platforms

7705 SAR Gen 2

drop

Syntax

drop

drop hop-limit {**lt** | **gt** | **eq**} *hop-limit-value*

drop hop-limit range hop-limit-value *hop-limit-value*

drop pattern expression *expression* **mask** *mask* **offset-type** *offset-type* **offset-value** *offset-value*

drop payload-length {**lt** | **gt** | **eq**} *payload-length-value*

drop payload-length range *payload-length-value payload-length-value*

Context

[Tree] (config>filter>ipv6-filter>entry>action drop)

Full Context

configure filter ipv6-filter entry action drop

Description

This command configures the drop action for the traffic that matches this filter entry.

Traffic can, also, be dropped based on *payload-length*, *payload-length range*, *hop-limit*, *hop-limit range*, or a pattern of conditional match criteria.

Packets that match the filter entry match criteria, but do not match the conditional match criteria value, are implicitly forwarded with no further match in the following filter entries.

For pattern match:

- the expression is left-aligned for the odd number bytes, for example, the expression 0xABC is programmed 0x0ABC in the line card
- the 'data' offset requires protocol UDP or TCP to be selected in the filter entry match criteria

Parameters

hop-limit

Specifies the hop-limit value for the drop action.

Values lt — Specifies "less than”. The **lt** parameter cannot be used with the lowest possible numerical value for the parameter.

eq — Specifies "equal to”.

gt — Specifies "greater than”. The **gt** parameter cannot be used with the highest possible numerical value for the parameter.

hop-limit-value

Specifies the hop-limit value for the drop action.

Values 0 to 255

range

Specifies an inclusive range. When the **range** parameter is used, the start of the range (the first value entered) must be smaller than the end of the range (the second value entered).

expression

Specifies the hexadecimal pattern to match; up to eight bytes.

Values 0x0000000000000001 to 0xffffffffffffffff

mask

Specifies the mask for the pattern expression, up to eight bytes.

Values 0x0000000000000001 to 0xffffffffffffffff

offset-type

Specifies the starting point reference for the offset-value of this pattern.

Values layer-3, layer-4, data, dns-qtype

offset-value

Specifies the offset value for the pattern expression. Dns-qtype supports offset value of 0.

Values 0 to 255

payload-length

Specifies drop packets matching both the filter entry match criteria and the payload-length-value defined in the drop action statement. Packets matching the filter entry match criteria and not matching the payload-length-value, as defined in the drop action statement, are implicitly forwarded with no further match in the following filter entries.

Values lt — Specifies "less than”. The **lt** parameter cannot be used with the lowest possible numerical value for the parameter.

gt — Specifies "greater than”. The **gt** parameter cannot be used with the highest possible numerical value for the parameter.

eq — Specifies "equal to”.

payload-length-value

Specifies the payload length value for the drop action.

Values 0 to 65535

Platforms

7705 SAR Gen 2

8.145 drop-count

drop-count

Syntax

drop-count *count*

no drop-count

Context

[Tree] (config>service>vprn>static-route-entry>indirect>cpe-check drop-count)

[Tree] (config>service>vprn>static-route-entry>next-hop>cpe-check drop-count)

Full Context

configure service vprn static-route-entry indirect cpe-check drop-count

configure service vprn static-route-entry next-hop cpe-check drop-count

Description

This optional parameter specifies the number of consecutive ping-replies that must be missed to declare the CPE down and to deactivate the associated static route.

Default

drop-count 3

Parameters

count

An integer count value.

Values 1 to 255

Platforms

7705 SAR Gen 2

drop-count

Syntax

drop-count *consecutive-failures* [**hold-down** *seconds*]

no drop-count

Context

[Tree] (config>filter>redirect-policy>dest>ping-test drop-count)

Full Context

configure filter redirect-policy destination ping-test drop-count

Description

This command specifies the number of consecutive requests that must fail for the destination to be declared unreachable and the time to hold destination unreachable before repeating tests.

Default

drop-count 3 hold-down 0

Parameters

consecutive-failures

Specifies the number of consecutive ping test failures before declaring the destination down.

Values 1 to 60

hold-down *seconds*

Specifies the amount of time, in seconds, that the system should be held down if any of the test has marked it unreachable.

Values 0 to 86400

Platforms

7705 SAR Gen 2

drop-count

Syntax

drop-count *count*

no drop-count

Context

[Tree] (config>router>static-route-entry>next-hop>cpe-check drop-count)

[Tree] (config>router>static-route-entry>indirect>cpe-check drop-count)

Full Context

configure router static-route-entry next-hop cpe-check drop-count

configure router static-route-entry indirect cpe-check drop-count

Description

This optional parameter specifies the number of consecutive ping-replies that must be missed to declare the CPE down and to deactivate the associated static route.

Default

drop-count 3

Parameters

count

Specifies the integer count value.

Values 1 to 255

Platforms

7705 SAR Gen 2

drop-count

Syntax

drop-count *count*

no drop-count

Context

[Tree] (config>vrrp>policy>priority-event>host-unreachable drop-count)

Full Context

configure vrrp policy priority-event host-unreachable drop-count

Description

This command configures the number of consecutively sent ICMP echo request messages that must fail before the host unreachable priority control event is set.

The **drop-count** command is used to define the number of consecutive message send attempts that must fail for the **host-unreachable** priority event to enter the set state. Each unsuccessful attempt increments the event's consecutive message drop counter. With each successful attempt, the event's consecutive message drop counter resets to zero.

If the event's consecutive message drop counter reaches the **drop-count** value, the **host-unreachable** priority event enters the set state.

The event's **hold-set** value defines how long the event must stay in the set state even when a successful message attempt clears the consecutive drop counter. The event is not cleared until the consecutive drop counter is less than the **drop-count** value and the **hold-set** timer has a value of zero (expired).

The **no** form of the command reverts to the default value.

Default

drop-count 3 — 3 consecutive ICMP echo request failures are required before the host unreachable priority control event is set.

Parameters

count

The number of ICMP echo request message attempts that must fail for the event to enter the set state. It also defines the threshold so a lower consecutive number of failures can clear the event state.

Values 1 to 60

Platforms

7705 SAR Gen 2

8.146 drop-extracted-traffic

drop-extracted-traffic

Syntax

drop-extracted-traffic

Context

[Tree] (config>filter>ipv6-filter>entry>action drop-extracted-traffic)

[Tree] (config>filter>ip-filter>entry>action drop-extracted-traffic)

Full Context

configure filter ipv6-filter entry action drop-extracted-traffic

configure filter ip-filter entry action drop-extracted-traffic

Description

This command specifies that a packet matching this filter entry is dropped if extracted to the CPM. Packets matching the filter entry match criteria and not extracted to the CPM are forwarded with no further match in the following filter entries.

Platforms

7705 SAR Gen 2

8.147 drop-tail

drop-tail

Syntax

drop-tail

Context

[Tree] (config>service>vpls>sap>egress>queue-override>queue drop-tail)

[Tree] (config>service>vpls>sap>ingress>queue-override>queue drop-tail)

[Tree] (config>service>ies>if>sap>egress>queue-override>queue drop-tail)

[Tree] (config>service>ies>if>sap>ingress>queue-override>queue drop-tail)

Full Context

configure service vpls sap egress queue-override queue drop-tail

configure service vpls sap ingress queue-override queue drop-tail

configure service ies interface sap egress queue-override queue drop-tail

configure service ies interface sap ingress queue-override queue drop-tail

Description

Commands in this context configure queue drop tail parameters.

Platforms

7705 SAR Gen 2

drop-tail

Syntax

drop-tail

Context

[Tree] (config>port>eth>access>ing>qgrp>qover>q drop-tail)

[Tree] (config>port>eth>access>egr>qgrp>qover>q drop-tail)

[Tree] (config>port>ethernet>network>egr>qgrp>qover>q drop-tail)

Full Context

configure port ethernet access ingress queue-group queue-overrides queue drop-tail

configure port ethernet access egress queue-group queue-overrides queue drop-tail

configure port ethernet network egress queue-group queue-overrides queue drop-tail

Description

Commands in this context configure queue drop tail parameters.

Platforms

7705 SAR Gen 2

drop-tail

Syntax

drop-tail

Context

[Tree] (config>service>epipe>sap>ingress>queue-override>queue drop-tail)

[Tree] (config>service>epipe>sap>egress>queue-override>queue drop-tail)

Full Context

configure service epipe sap ingress queue-override queue drop-tail

configure service epipe sap egress queue-override queue drop-tail

Description

Commands in this context configure queue drop tail parameters.

Platforms

7705 SAR Gen 2

drop-tail

Syntax

drop-tail

Context

[Tree] (config>service>vprn>if>sap>egress>queue-override>queue drop-tail)

[Tree] (config>service>vprn>if>sap>ingress>queue-override>queue drop-tail)

Full Context

configure service vprn interface sap egress queue-override queue drop-tail

configure service vprn interface sap ingress queue-override queue drop-tail

Description

Commands in this context configure queue drop tail parameters.

Platforms

7705 SAR Gen 2

drop-tail

Syntax

drop-tail

Context

[Tree] (config>qos>sap-ingress>queue drop-tail)

[Tree] (config>qos>sap-egress>queue drop-tail)

Full Context

configure qos sap-ingress queue drop-tail

configure qos sap-egress queue drop-tail

Description

Commands in this context configure queue drop tail parameters.

Platforms

7705 SAR Gen 2

drop-tail

Syntax

drop-tail

Context

[Tree] (config>qos>network-queue>queue drop-tail)

Full Context

configure qos network-queue queue drop-tail

Description

Commands in this context configure queue drop tail parameters.

Platforms

7705 SAR Gen 2

drop-tail

Syntax

drop-tail

Context

[Tree] (config>qos>qgrps>ing>qgrp>queue drop-tail)

[Tree] (config>qos>qgrps>egr>qgrp>queue drop-tail)

Full Context

configure qos queue-group-templates ingress queue-group queue drop-tail

configure qos queue-group-templates egress queue-group queue drop-tail

Description

Commands in this context configure queue drop-tail parameters.

Platforms

7705 SAR Gen 2

8.148 dsap

dsap

Syntax

dsap *dsap-value* [*dsap-mask*]

no dsap

Context

[Tree] (config>qos>sap-ingress>mac-criteria>entry>match dsap)

Full Context

configure qos sap-ingress mac-criteria entry match dsap

Description

Configures an Ethernet 802.2 LLC DSAP value or range for an ingress SAP QoS policy match criterion.

This is a 1-byte field that is part of the 802.2 LLC header of the IEEE 802.3 Ethernet Frame.

The snap-pid field, etype field, ssap, and dsap fields are mutually exclusive and cannot be part of the same match criteria.

Use the **no** form of this command to remove the dsap value as the match criterion.

Default

no dsap

Parameters

dsap-value

The 8-bit dsap match criteria value in hexadecimal.

Values 0x00 to 0xFF (hex)

dsap-mask

This is optional and can be used when specifying a range of dsap values to use as the match criteria.

This 8-bit mask can be configured using the following formats.

Table 25: Format Styles to Configure Mask

Format Style	Format Syntax	Example
Decimal	DDD	240
Hexadecimal	0xHH	0xF0

Format Style	Format Syntax	Example
Binary	0bBBBBBBBB	0b11110000

Values 0x00 to 0xFF (hex)

Default FF

Platforms

7705 SAR Gen 2

dsap

Syntax

dsap *dsap-value* [*dsap-mask*]

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry>match dsap)

Full Context

configure system security management-access-filter mac-filter entry match dsap

Description

This command configures DSAP match conditions.

Parameters

dsap-value

Specifies the 8-bit DSAP match criteria value in hexadecimal.

Values 0x00 to 0xFF (hex)

mask

Specifies a range of DSAP values to use as the match criteria.

This 8 bit mask can be configured using the formats described in Table 26: Format Styles:

Table 26: Format Styles

Format Style	Format Syntax	Example
Decimal	DDD	240
Hexadecimal	0xHH	0xF0
Binary	0bBBBBBBBB	0b11110000

Default FF (hex) (exact match)

Values 0x00 to 0xFF

Platforms

7705 SAR Gen 2

8.149 dscp

dscp

Syntax

dscp *dscp-name*

no dscp

Context

[Tree] (config>service>ies>if>sap>ip-tunnel dscp)

Full Context

configure service ies interface sap ip-tunnel dscp

Description

This command sets the DSCP code-point in the outer IP header of encapsulated packets associated with a particular tunnel.

The **no** form of this command copies the DSCP value from the inner IP header (after remarking by the private tunnel SAP egress qos policy) to the outer IP header.

Default

no dscp

Parameters

dscp

Specifies the DSCP code-point to be used.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

Platforms

7705 SAR Gen 2

dscp

Syntax

dscp *dscp-name* **fc** *fc-name*

no dscp *dscp-name*

Context

[Tree] (config>service>vprn>sgt-qos dscp)

[Tree] (config>router>sgt-qos dscp)

Full Context

configure service vprn sgt-qos dscp

configure router sgt-qos dscp

Description

This command creates a mapping between the DiffServ Code Point (DSCP) of the self-generated traffic and the forwarding class.

Self-generated traffic that matches the specified DSCP will be assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all 64 DSCPs to the forwarding class.

All DSCP names that define a DSCP value must be explicitly defined.

The **no** form of this command removes the DSCP-to-forwarding class association.

Parameters

dscp-name

Specifies the name of the DSCP to be associated with the forwarding class. DiffServ code point can only be specified by its name and only an existing DiffServ code point can be specified. The software provides names for the well-known code points.

Values be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

fc *fc-name*

Specifies the forwarding class name. All packets with a DSCP value or MPLS EXP bit that are not defined will be placed in this forwarding class.

Values be, l2, af, l1, h2, ef, h1, nc

Platforms

7705 SAR Gen 2

dscp

Syntax

dscp *dscp-name*

no dscp

Context

[Tree] (config>service>vprn>if>sap>ip-tunnel dscp)

Full Context

configure service vprn interface sap ip-tunnel dscp

Description

This command sets the DSCP code-point in the outer IP header of GRE encapsulated packets associated with a particular GRE tunnel. The default, set using the **no** form of this command, is to copy the DSCP value from the inner IP header (after remarking by the private tunnel SAP egress qos policy) to the outer IP header.

Default

no dscp

Parameters

dscp

Specifies the DSCP code-point to be used.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

Platforms

7705 SAR Gen 2

dscp

Syntax

dscp *dscp-name*

dscp resolve

Context

[Tree] (config>oam-pm>session>ip dscp)

Full Context

configure oam-pm session ip dscp

Description

This command can be used to explicitly configure the DSCP value to the specified *dscp-name*, or to use the configured **fc** and **profile** values to derive the DSCP value from the egress network QoS policy 1.

Default

dscp resolve

Parameters

dscp-name

Specifies the Diffserv code point name.

Values be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

resolve

Specifies to use the configured **fc** and **profile** values to derive the DSCP value from the egress network QoS policy 1.

Platforms

7705 SAR Gen 2

dscp

Syntax

dscp *dscp-name* [*dscp-name*] **fc** *fc-name* [**priority** {**low** | **high**}]

no dscp *dscp-name* [*dscp-name*]

Context

[Tree] (config>qos>sap-ingress dscp)

Full Context

configure qos sap-ingress dscp

Description

This command explicitly sets the forwarding class or subclass or enqueuing priority when a packet is marked with the DiffServ Code Point (DSCP) value contained in the *dscp-name*. A list of up to eight dscp-names can be entered on a single command. The lists of dscp-names within the configuration are managed by the system to ensure that each list does not exceed eight names. Entering more than eight dscp-names with the same parameters (**fc**, **priority**) will result in multiple lists being created. Conversely, multiple lists with the same parameters (fc, priority) are merged and the lists repacked to a maximum of eight per list if DSCP names are removed or the parameters changed so the multiple lists use the same parameters. Also, if a subset of a list is entered with different parameters, then a new list will be created for the subset. When the list is stored in the configuration, the DSCP names are sorted by their DSCP value in ascending numerical order; consequently, the order in the configuration may not be exactly what the user entered.

Adding a DSCP rule on the policy forces packets that match the DSCP value specified to override the forwarding class and enqueuing priority based on the parameters included in the DSCP rule. When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueuing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueuing priority derived from earlier matches in the classification hierarchy.

The DSCP value (referred to here by *dscp-name*) is derived from the most significant six bits in the IPv4 header ToS byte field (DSCP bits) or the Traffic Class field from the IPv6 header. If the packet does not have an IP header, DSCP-based matching is not performed. The six DSCP bits define 64 DSCP values used to map packets to per-hop Quality of Service (QoS) behavior. The most significant three bits in the IP header ToS byte field are also commonly used in a more traditional manner to specify an IP precedence value, causing an overlap between the precedence space and the DSCP space. Both IP precedence and DSCP classification rules are supported.

DSCP rules have a higher match priority than IP precedence rules and where a *dscp-name* DSCP value overlaps an *ip-prec-value*, the DSCP rule takes precedence.

The **no** form of this command removes the specified the *dscp-names* from the explicit DSCP classification rule in the SAP ingress policy. As *dscp-names* are removed, the system repacks the lists of *dscp-names* with the same parameters (up to eight per list). As the **no** command does not have any additional parameters, it is possible to remove multiple *dscp-names* from multiple DSCP statements having different parameters with one command. If a *dscp-name* specified in a **no** command does not exist in any DSCP statement, then the command is aborted at that point with an error message displayed; any DSCP names in the list before the failed entry will be processed as normal but the processing will stop at the failed entry so that the remainder of the list is not processed.

Removing the *dscp-name* from the policy immediately removes the DSCP name on all ingress SAPs using the policy.

Parameters

dscp-name

The DSCP name is a required parameter that specifies the unique IP header ToS byte DSCP bits value that will match the DSCP rule. If the command is executed multiple times with the same *dscp-name*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of 64 DSCP rules are allowed on a single policy and a maximum of eight *dscp-names* can be specified in a single statement.

The specified name must exist as a *dscp-name*. SR OS software provides names for the well-known code points; these can be shown using the **show qos dscp-table** command.

fc *fc-name*

The value given for *fc-name* must be one of the predefined forwarding classes in the system. Specifying the *fc-name* is optional. When a packet matches the rule, the forwarding class is only overridden when the **fc** *fc-name* parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

The *subclass-name* parameter is optional and used with the *fc-name* parameter to define a preexisting subclass. The *fc-name* and *subclass-name* parameters must be separated by a period (.). If *subclass-name* does not exist in the context of *fc-name*, an error will occur. If *subclass-name* is removed using the **no fc** *fc-name.subclass-name* **force** command, the **default-fc** command will automatically drop the *subclass-name* and only use *fc-name* (the parent forwarding class for the subclass) as the forwarding class.

Values

fc: *class*[.*subclass*]

class: be, l2, af, l1, h2, ef, h1, nc

subclass: 29 characters max

Default Inherit (when **fc** *fc-name* is not defined, the rule preserves the previous forwarding class of the packet).

priority

This parameter overrides the default enqueuing priority for all packets received on an ingress SAP using this policy that match this rule. Specifying the priority is optional. When a packet matches the rule, the enqueuing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and priority is not explicitly defined in the rule, the enqueuing priority is inherited based on previous rule matches.

Default Inherits the priority defined by the default-priority statement.

high

This parameter is used in conjunction with the **priority** parameter. Setting the enqueuing parameter to **high** for a packet increases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

low

This parameter is used in conjunction with the **priority** parameter. Setting the enqueuing parameter to **low** for a packet decreases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

Platforms

7705 SAR Gen 2

dscp

Syntax

dscp *dscp-name*

no dscp

Context

[Tree] (config>qos>sap-ingress>ipv6-criteria>entry>match dscp)

[Tree] (config>qos>sap-egress>ip-criteria>entry>match dscp)

[Tree] (config>qos>sap-ingress>ip-criteria>entry>match dscp)

[Tree] (config>qos>sap-egress>ipv6-criteria>entry>match dscp)

Full Context

configure qos sap-ingress ipv6-criteria entry match dscp

configure qos sap-egress ip-criteria entry match dscp

configure qos sap-ingress ip-criteria entry match dscp

configure qos sap-egress ipv6-criteria entry match dscp

Description

This command configures a DSCP code point to be used as a SAP QoS policy match criterion.

The **no** form of this command removes the DSCP match criterion.

Default

no dscp

Parameters

dscp-name

Specifies a dscp name that has been previously mapped to a value using the **dscp-name** command. The DiffServ code point can only be specified by its name.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

Platforms

7705 SAR Gen 2

dscp

Syntax

dscp *dscp-name* [*dscp-name*] [**fc** *fc-name*] [**profile** {**in** | **out** | **exceed** | **inplus**}]

no dscp *dscp-name* [*dscp-name*]

Context

[Tree] (config>qos>sap-egress dscp)

Full Context

configure qos sap-egress dscp

Description

This command defines IP Differentiated Services Code Point (DSCP) names that must be matched to perform the associated reclassification actions. The specified name must exist as a DSCP name. SR OS software provides names for the well-known code points. A list of up to eight DSCP names can be entered on a single command. The lists of DSCP names within the configuration are managed by the system to ensure that each list does not exceed eight names. Entering more than eight DSCP names with the same parameters (**fc** and **profile**) results in multiple lists being created. Conversely, multiple lists with the same parameters (**fc** and **profile**) are merged and the lists repacked to a maximum of eight per list if DSCP names are removed or the parameters changed so the multiple lists use the same parameters. Also, if a subset of a list is entered with different parameters, a new list is created for the subset. When the list is stored in the configuration, the DSCP names are sorted by their DSCP value in ascending numerical order; consequently, the order in the configuration may not be exactly what the user entered.

If an egress packet on the SAP matches an IP DSCP value corresponding to a specified DSCP name, the forwarding class, profile egress queue accounting behavior may be overridden. By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions. Matching a DSCP-based reclassification rule will override all IP precedence-based reclassification rule actions.

The IP DSCP bits used to match against DSCP reclassification rules come from the Type of Service (ToS) field within the IPv4 header or the traffic class field from the IPv6 header. If the packet does not have an IP header, DSCP-based matching is not performed.

The reclassification actions from a DSCP reclassification rule may be overridden by an IP flow match event.

The **fc** keyword is optional. When specified, the egress classification rule will overwrite the forwarding class derived from ingress. The new forwarding class is used for egress remarking and queue mapping decisions. If an IP criteria match occurs after the DSCP match, the new forwarding class may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new fc, the fc from the dscp match will be used.

The **profile** keyword is optional. When specified, the egress classification rule will overwrite the profile of the packet derived from ingress. The new profile value is used for egress remarking and queue congestion behavior. If an IP criteria match occurs after the DSCP match, the new profile may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new profile, the profile from the DSCP match will be used.

The **no** form of this command removes the specified the *dscp-names* from the reclassification rule in the SAP egress QoS policy. As *dscp-names* are removed, the system repacks the lists of *dscp-names* with

the same parameters (up to 8 per list). As the **no** command does not have any additional parameters, it is possible to remove multiple *dscp-names* from multiple DSCP statements having different parameters with one command. If a *dscp-name* specified in a **no** command does not exist in any DSCP statement, the command is aborted at that point with an error message displayed. Any *dscp-names* in the list before the failed entry will be processed as normal but the processing will stop at the failed entry so that the remainder of the list is not processed.

Parameters

dscp-name

The *dscp-name* parameter is required when defining a DSCP reclassification rule. The specified name must exist as a DSCP name. A maximum of eight DSCP names can be specified in a single statement. SR OS software provides names for the well-known code points, which can be shown using the **show qos dscp-table** command.

***fc-name*:**

The **fc** reclassification action is optional. When specified, packets matching the IP DSCP value corresponding to a specified *dscp-name* will be explicitly reclassified to the forwarding class specified as *fc-name* regardless of the ingress classification decision. The explicit forwarding class reclassification may be overwritten by an IP criteria reclassification match. The **fc** name defined must be one of the eight forwarding classes supported by the system. To remove the forwarding class reclassification action for the specified DSCP value, the **dscp** command must be re-executed without the **fc** reclassification action defined.

Values be, l1, af, l2, h1, ef, h2 or nc

counter-id

Specifies the counter ID.

profile

The profile reclassification action is optional. When specified, packets matching the IP DSCP value corresponding to a specified *dscp-name* will be explicitly reclassified to the specified profile regardless of the ingress profiling decision. The explicit profile reclassification may be overwritten by an IPv6 criteria or IP criteria reclassification match. To remove the profile reclassification action for the specified *dscp-name*, the **dscp** command must be re-executed without the profile reclassification action defined.

in

Specifies that any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.

out

Specifies that any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.

exceed

Specifies that when **exceed** is specified, any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.

inplus

Specifies that any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

Platforms

7705 SAR Gen 2

dscp

Syntax

dscp {*dscp-name* | **in-profile** *dscp-name* **out-profile** *dscp-name* [**exceed-profile** *dscp-name*]}

no dscp

Context

[Tree] (config>qos>sap-egress>fc dscp)

Full Context

configure qos sap-egress fc dscp

Description

This command configures a DSCP to be used for remarking packets from the specified FC. If the optional **exceed-profile**, **in-profile**, or **out-profile** keyword is specified, the command will remark different DSCP depending on whether the packet was classified to be exceed, in-profile, or out-of-profile ingress to the node. All inplus-profile traffic is marked with the same value as in-profile traffic.

Default

no dscp

Parameters

dscp-name

Specifies a DSCP name that has been previously mapped to a value using the **dscp-name** command. The DSCP can only be specified by its name.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

exceed-profile *dscp-name*

This optional parameter specifies the DSCP name to be used to remark the traffic that is exceed-profile. If not specified, this defaults to the same value configured for **out-profile** parameter.

in-profile *dscp-name*

Specifies the DSCP name to be used to remark the traffic that is in-profile.

out-profile *dscp-name*

Specifies the DSCP name to be used to remark the traffic that is out-of-profile.

Platforms

7705 SAR Gen 2

dscp

Syntax

dscp *dscp-name* **fc** *fc-name* **profile** {**in** | **out**}

no dscp

Context

[Tree] (config>qos>network>ingress dscp)

Full Context

configure qos network ingress dscp

Description

This command creates a mapping between the DiffServ Code Point (DSCP) of the network ingress traffic and the forwarding class.

Ingress traffic that matches the specified DSCP will be assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all 64 DiffServ code points to the forwarding class. For undefined code points, packets are assigned to the forwarding class specified under the **default-action** command.

The **no** form of this command removes the DiffServ code point-to-forwarding class association. The **default-action** then applies to that code point value.

Parameters

dscp-name

The name of the DiffServ code point to be associated with the forwarding class. DiffServ code point can only be specified by its name and only an existing DiffServ code point can be specified. The software provides names for the well-known code points.

The system-defined names available are as follows. The system-defined names must be referenced as all lowercase, exactly as shown in the first column in Table 27: Default DSCP Names to DSCP Value Mapping and Table 28: Default Class Selector Code Points to DSCP Value Mapping.

Additional names-to-code point value associations can be added using the '**dscp-name** *dscp-name dscp-value*' command.

The actual mapping is being done on the *dscp-value*, not the *dscp-name* that references the *dscp-value*. If a second *dscp-name* that references the same *dscp-value* is mapped within the policy, an error will occur. The second name will not be accepted until the first name is removed.

Table 27: Default DSCP Names to DSCP Value Mapping

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary
nc1	48	0x30	0b110000
nc2	56	0x38	0b111000
ef	46	0x2e	0b101110
af41	34	0x22	0b100010
af42	36	0x24	0b100100
af43	38	0x26	0b100110
af31	26	0x1a	0b011010
af32	28	0x1c	0b011100
af33	30	0x1d	0b011110
af21	18	0x12	0b010010
af22	20	0x14	0b010100
af23	22	0x16	0b010110
af11	10	0x0a	0b001010
af12	12	0x0c	0b001100
af13	14	0x0e	0b001110
default	0	0x00	0b000000

Table 28: Default Class Selector Code Points to DSCP Value Mapping

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary
cs7	56	0x38	0b111000
cs6	48	0X30	0b110000
cs5	40	0x28	0b101000
cs4	32	0x20	0b100000
cs3	24	0x18	0b011000
cs2	16	0x10	0b010000

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary
cs1	08	0x8	0b001000

fc-name

Enter this required parameter to specify the *fc-name* with which the code point will be associated.

Values be, l2, af, l1, h2, ef, h1, nc

profile {in | out}

Enter this required parameter to indicate whether the DiffServ code point value is the in-profile or out-of-profile value. For every DSCP value defined, the profile must be indicated. If a DSCP value is not mapped, the default-action forwarding class and profile state will be used for that value.

DSCP values mapping to forwarding classes Expedited (ef), High-1 (h1) and Network-Control (nc) can only be set to in-profile.

DSCP values mapping to forwarding class "be" can only be set to out-of-profile.

Values in, out

Platforms

7705 SAR Gen 2

dscp

Syntax

dscp *dscp-name* **fc** *fc-name* **profile** {**in** | **out** | **exceed** | **inplus**}

no dscp *dscp-name*

Context

[Tree] (config>qos>network>egress dscp)

Full Context

configure qos network egress dscp

Description

This command configures an IP Differentiated Services Code Point (DSCP) value that must be matched to perform the associated reclassification actions. If an egress packet on an IES/VPRN interface spoke SDP, on a CSC network interface in a VPRN, or on a network interface that the network QoS policy is applied to, matches the specified IP DSCP value, the forwarding class and profile may be overridden.

By default, the forwarding class and profile of the packet are derived from ingress classification and profiling functions. Matching a DHCP-based reclassification rule will override all IP precedence-based reclassification rule actions.

The IP DSCP bits used to match against DSCP reclassification rules come from the Type of Service (ToS) field within the IPv4 header or the Traffic Class field from the IPv6 header. If the packet does not have an IP header, DSCP-based matching is not performed.

The configuration of egress DSCP classification and the configuration of an egress IP criteria or IPv6 criteria entry statement within a network QoS policy are mutually exclusive.

The IP precedence- and DSCP-based reclassification are supported on a network interface, on a CSC network interface in a VPRN, and on a PW used in an IES or VPRN spoke interface. The CLI will block the application of a network QoS policy with the egress reclassification commands to the spoke SDP part of a Layer 2 service.

Conversely, the CLI will not allow the user to add the egress reclassification commands to a network QoS policy if the policy is being used by a Layer 2 spoke SDP.

The egress reclassification commands will only take effect if the redirection of the spoke SDP or CSC interface to use an egress port queue group succeeds. For example, the following CLI command would be successful:

config>**service**>**vprn**>**if**>**spoke-sdp**>**egress**> **qos** *network-policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*

config>**service**>**ies**>**if**>**spoke-sdp>egress**> **qos** *network-policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*

config>service>vprn>nw-if>qos *network-policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*

If the redirection command fails, the PW will use the network QoS policy assigned to the network IP interface, however any reclassification in the network QoS policy applied to the network interface will be ignored.

The **no** form of this command removes the egress reclassification rule.

Parameters

dscp-name

be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

fc fc-name

be, l2, af, l1, h2, ef, h1, nc

***profile* {*in* | *out* | *exceed* | *inplus*}**

The profile reclassification action is mandatory. When specified, packets matching the DSCP value will be explicitly reclassified to the profile specified regardless of the ingress profiling decision. To remove the profile reclassification action for the specified DSCP value, the **no dscp** command must be executed.

in - Specifies that any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.

out - Specifies that any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.

exceed - Specifies that any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.

inplus - Specifies that any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

Platforms

7705 SAR Gen 2

dscp

Syntax

dscp *dscp-name*

no dscp

Context

[Tree] (config>qos>network>ingress>ip-criteria>entry>match dscp)

[Tree] (config>qos>network>ingress>ipv6-criteria>entry>match dscp)

[Tree] (config>qos>network>egress>ip-criteria>entry>match dscp)

[Tree] (config>qos>network>egress>ipv6-criteria>entry>match dscp)

Full Context

configure qos network ingress ip-criteria entry match dscp

configure qos network ingress ipv6-criteria entry match dscp

configure qos network egress ip-criteria entry match dscp

configure qos network egress ipv6-criteria entry match dscp

Description

This command configures a DSCP to be used as a network QoS policy match criterion.

The **no** form of this command removes the DSCP match criterion.

Parameters

dscp-name

Specifies a DSCP name that has been previously mapped to a value using the **dscp-name** command. The DSCP can only be specified by its name.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

Platforms

7705 SAR Gen 2

dscp

Syntax

dscp *dscp-name*

no dscp

Context

[Tree] (config>filter>ipv6-filter>entry>match dscp)

[Tree] (config>filter>ip-filter>entry>match dscp)

Full Context

configure filter ipv6-filter entry match dscp

configure filter ip-filter entry match dscp

Description

This command configures a DiffServ Code Point (DSCP) name to be used as an IP filter match criterion.

The **no** form of the command removes the DSCP match criterion.

Default

no dscp

Parameters

dscp-name

Configures a DSCP name. The DiffServ code point may only be specified by its name.

Values be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

Platforms

7705 SAR Gen 2

8.150 dscp-in-profile

dscp-in-profile

Syntax

dscp-in-profile *dscp-name*

no dscp-in-profile

Context

[Tree] (config>qos>network>egress>fc dscp-in-profile)

Full Context

configure qos network egress fc dscp-in-profile

Description

This command specifies the in-profile DSCP name for the forwarding class. The corresponding DSCP value will be used for all IP packets that require marking at egress on this forwarding class queue, and that are in-profile. The inplus-profile traffic is marked with the same value as in-profile traffic.

When multiple DSCP names are associated with the forwarding class at network egress, the last name entered will overwrite the previous value.

The **no** form of this command resets the configuration to the factory default in-profile DSCP name setting for *policy-id* 1.

Parameters

dscp-name

Specifies the system- or user-defined, case-sensitive *dscp-name.*

Values Any defined system- or user-defined *dscp-name*

Platforms

7705 SAR Gen 2

8.151 dscp-out-profile

dscp-out-profile

Syntax

dscp-out-profile *dscp-name*

no dscp-out-profile

Context

[Tree] (config>qos>network>egress>fc dscp-out-profile)

Full Context

configure qos network egress fc dscp-out-profile

Description

This command specifies the out-of-profile DSCP name for the forwarding class. The corresponding DSCP value will be used for all IP packets requiring marking the egress on this forwarding class queue that are out-of-profile. The exceed-profile traffic is marked with the same value as out-of-profile traffic.

When multiple DSCP names are associated with the forwarding class at network egress, the last name entered will overwrite the previous value.

The **no** form of this command resets the configuration to the factory default out-of-profile DSCP name setting for *policy-id* 1.

Default

Policy-id 1:	Factory setting
Policy-id 2 to 65535:	Policy-id setting

Parameters

dscp-name

Specifies the system- or user-defined, case-sensitive *dscp-name.*

Values Any defined system- or user-defined *dscp-name*

Platforms

7705 SAR Gen 2

8.152 dst-ip

dst-ip

Syntax

dst-ip {*ip-address/mask* | *ip-address* [*ipv4-address-mask*] | **ip-prefix-list** *prefix-list-name*}

no dst-ip

Context

[Tree] (config>qos>sap-ingress>ip-criteria>entry>match dst-ip)

[Tree] (config>qos>sap-egress>ip-criteria>entry>match dst-ip)

Full Context

configure qos sap-ingress ip-criteria entry match dst-ip

configure qos sap-egress ip-criteria entry match dst-ip

Description

This command configures a destination address range to be used as a SAP QoS policy match criterion.

To match on the IPv4 destination address, specify the address and its associated mask, e.g., 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used for IPv4.

The **no** form of this command removes the destination IPv4 address match criterion.

Default

no dst-ip

Parameters

ip-address

Specifies the destination IPv4 address specified in dotted decimal notation.

Values ip-address: a.b.c.d

mask

Specifies the length in bits of the subnet mask.

Values 1 to 32

ipv4-address-mask

Specifies the subnet mask in dotted decimal notation.

Values a.b.c.d (dotted quad equivalent of mask length)

prefix-list-name

Specifies the IPv4 prefix list name, a string of up to 32 printable ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

dst-ip

Syntax

dst-ip {*ipv6-address/prefix-length* | *ipv6-address ipv6-address-mask* | **ipv6-prefix-list** *ipv6-prefix-list-name*}

no dst-ip

Context

[Tree] (config>qos>sap-ingress>ipv6-criteria>entry>match dst-ip)

[Tree] (config>qos>sap-egress>ipv6-criteria>entry>match dst-ip)

Full Context

configure qos sap-ingress ipv6-criteria entry match dst-ip

configure qos sap-egress ipv6-criteria entry match dst-ip

Description

This command configures a destination address range to be used as a SAP QoS policy match criterion.

To match on the IPv6 destination address, specify the address and its associated mask, for example, 2001:db8:1000::/64.

The **no** form of this command removes the destination IPv6 address match criterion.

Default

no dst-ip

Parameters

ipv6-address

Specifies the IPv6 address for the IP match criterion in hexadecimal digits.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x: [0 to FFFF]H
d: [0 to 255]D

prefix-length

Specifies the IPv6 prefix length for the IPv6 address expressed as a decimal integer.

Values 1 to 128

ipv6-address-mask

Specifies the IPv6 address mask.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x: [0 to FFFF]H
d: [0 to 255]D

ipv6-prefix-list-name

Specifies the IPv6 prefix list name, a string of up to 32 printable ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

dst-ip

Syntax

dst-ip {*ip-address/mask* | *ip-address ipv4-address-mask* | **ip-prefix-list** *ip-prefix-list-name*}

dst-ip {*ipv6-address/mask* | *ipv6-address ipv6-address-mask* | **ipv6-prefix-list** *ipv6-prefix-list-name*}

no dst-ip

Context

[Tree] (config>qos>network>egress>ip-criteria>entry>match dst-ip)

[Tree] (config>qos>network>ingress>ipv6-criteria>entry>match dst-ip)

[Tree] (config>qos>network>ingress>ip-criteria>entry>match dst-ip)

[Tree] (config>qos>network>egress>ipv6-criteria>entry>match dst-ip)

Full Context

configure qos network egress ip-criteria entry match dst-ip

configure qos network ingress ipv6-criteria entry match dst-ip

configure qos network ingress ip-criteria entry match dst-ip

configure qos network egress ipv6-criteria entry match dst-ip

Description

This command configures a destination address range to be used as a network QoS policy match criterion.

To match on the destination address, specify the address and its associated mask, for example, when specifying an IPv4 address, 10.1.0.0/16 or 10.1.0.0 255.255.0.0 can be used.

The **no** form of this command removes the destination IP address match criterion.

Parameters

ip-address

Specifies the source IPv4 address specified in dotted decimal notation.

Values ip-address: a.b.c.d

mask

Specifies the length in bits of the subnet mask.

Values 1 to 32

ipv4-address-mask

Specifies the subnet mask in dotted decimal notation.

Values a.b.c.d (dotted quad equivalent of mask length)

ip-prefix-list-name

Specifies an IPv4 prefix list which contains IPv4 address prefixes to be matched.

Values A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

ipv6-address

Specifies the IPv6 prefix for the IP match criterion in hex digits.

Values

ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

mask

Specifies the length of the IPv6 address expressed as a decimal integer.

Values 1 to 128

ipv6-address-mask

Specifies the eight 16-bit hexadecimal pieces representing bit match criteria.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)

ipv6-prefix-list-name

Specifies an IPv6 prefix list which contains IPv6 address prefixes to be matched.

Values A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

dst-ip

Syntax

IPv4:

dst-ip {*ip-address/mask* | *ip-address ipv4-address-mask* | **ip-prefix-list** *prefix-list-name*}

IPv6:

dst-ip {*ipv6-address/prefix-length* | *ipv6-address ipv6-address-mask* | **ipv6-prefix-list** *prefix-list-name*}

no dst-ip

Context

[Tree] (config>filter>ip-exception>entry>match dst-ip)

[Tree] (config>filter>ipv6-filter>entry>match dst-ip)

[Tree] (config>filter>ipv6-exception>entry>match dst-ip)

[Tree] (config>filter>ip-filter>entry>match dst-ip)

Full Context

configure filter ip-exception entry match dst-ip

configure filter ipv6-filter entry match dst-ip

configure filter ipv6-exception entry match dst-ip

configure filter ip-filter entry match dst-ip

Description

This command configures a destination address range to be used as a filter policy match criterion.

To match on the destination address, specify the address and its associated mask, e.g., 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used for IPv4.

The **no** form of this command removes the destination IPv4 or IPv6 address match criterion.

Default

no dst-ip

Parameters

ip-address

Specifies the destination IPv4 address in dotted decimal notation.

Values a.b.c.d

mask

Specifies the length in bits of the subnet mask.

Values 1 to 32

ipv4-address-mask

Specifies the subnet mask in dotted decimal notation.

Values a.b.c.d (dotted quad equivalent of mask length)

ip-prefix-listor ipv6-prefix-list prefix-list-name

Specifies to use a list of IP prefixes, which is a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

ipv6-address

Specifies the IPv6 prefix for the IP match criterion in hex digits.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0..FFFF]H

d: [0..255]D

prefix-length

Specifies the IPv6 prefix length for the *ipv6-address* as a decimal integer.

Values 1 to 128

ipv6-address-mask

Specifies the eight 16-bit hexadecimal pieces representing bit match criteria.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0..FFFF]H

d: [0..255]D

Platforms

7705 SAR Gen 2

8.153 dst-mac

dst-mac

Syntax

dst-mac *ieee-address* [*ieee-address-mask*]

no dst-mac

Context

[Tree] (config>qos>sap-ingress>mac-criteria>entry>match dst-mac)

Full Context

configure qos sap-ingress mac-criteria entry match dst-mac

Description

Configures a destination MAC address or range to be used as a Service Ingress QoS policy match criterion.

The **no** form of this command removes the destination MAC address as the match criterion.

Default

no dst-mac

Parameters

ieee-address

The MAC address to be used as a match criterion.

Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

ieee-address-mask

A 48-bit mask to match a range of MAC address values.

This 48-bit mask can be configured using the following formats.

Table 29: Format Styles to Configure Mask

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHH	0xFFFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

All packets with a source MAC OUI value of 00-03-FA, subject to a match condition, should be specified as: 0003FA000000 0x0FFFFF000000

Values 0x00000000000000 to 0xFFFFFFFFFFFF (hex)

Default 0xFFFFFFFFFFFF

Platforms

7705 SAR Gen 2

dst-mac

Syntax

dst-mac *ieee-address* [*ieee-address-mask*]

no dst-mac

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry>match dst-mac)

Full Context

configure system security management-access-filter mac-filter entry match dst-mac

Description

This command configures the destination MAC match condition.

Parameters

ieee-address

Specifies the MAC address to be used as a match criterion.

Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

mask

Specifies a 48-bit mask to match a range of MAC address values.

Platforms

7705 SAR Gen 2

8.154 dst-port

dst-port

Syntax

dst-port {**lt** | **gt** | **eq**} *dst-port-number*

dst-port range *start end*

no dst-port

Context

[Tree] (config>qos>sap-egress>ip-criteria>entry>match dst-port)

[Tree] (config>qos>sap-egress>ipv6-criteria>entry>match dst-port)

[Tree] (config>qos>sap-ingress>ipv6-criteria>entry>match dst-port)

[Tree] (config>qos>sap-ingress>ip-criteria>entry>match dst-port)

Full Context

configure qos sap-egress ip-criteria entry match dst-port

configure qos sap-egress ipv6-criteria entry match dst-port

configure qos sap-ingress ipv6-criteria entry match dst-port

configure qos sap-ingress ip-criteria entry match dst-port

Description

This command configures a destination TCP or UDP port number or port range for a SAP QoS policy match criterion.

The **no** form of this command removes the destination port match criterion.

Default

no dst-port

Parameters

{lt | gt | eq} *dst-port-number*

The TCP or UDP port numbers to match, specified as less than (**lt**), greater than (**gt**), or equal to (**eq**) to the destination port value, specified as a decimal integer.

Values 1 to 65535 (decimal)

range *startend*

The range of TCP or UDP port values to match, specified as between the *start* and *end* destination port values inclusive.

Values 1 to 65535 (decimal)

Platforms

7705 SAR Gen 2

dst-port

Syntax

dst-port {**lt** | **gt** | **eq**} *dst-port-number*

dst-port port-list *port-list-name*

dst-port range *start end*

no dst-port

Context

[Tree] (config>qos>network>egress>ipv6-criteria>entry>match dst-port)

[Tree] (config>qos>network>ingress>ipv6-criteria>entry>match dst-port)

[Tree] (config>qos>network>egress>ip-criteria>entry>match dst-port)

[Tree] (config>qos>network>ingress>ip-criteria>entry>match dst-port)

Full Context

configure qos network egress ipv6-criteria entry match dst-port

configure qos network ingress ipv6-criteria entry match dst-port

configure qos network egress ip-criteria entry match dst-port

configure qos network ingress ip-criteria entry match dst-port

Description

This command configures a destination TCP or UDP port number, port range, or a port list for a network QoS policy match criterion.

The **no** form of this command removes the destination port match criterion.

Parameters

lt

Keyword used to specify TCP or UDP port numbers to match that are less than the destination port value.

gt

Keyword used to specify TCP or UDP port numbers to match that are greater than the destination port value.

eq

Keyword used to specify TCP or UDP port numbers to match that are equal to the destination port value.

dst-port-number

Specifies the TCP or UDP port numbers to match, specified as less than (lt), greater than (gt), or equal to (eq) the destination port value, specified as a decimal integer.

Values 1 to 65535

port-list-name

Specifies a port list name, up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

start

Specifies the starting range of TCP or UDP port values to match.

Values 1 to 65535

end

Specifies the end range of TCP or UDP port values to match.

Values 1 to 65535

Platforms

7705 SAR Gen 2

dst-port

Syntax

dst-port {**lt** | **gt** | **eq**} *dst-port-number*

dst-port port-list *port-list-name*

dst-port range *dst-port-number dst-port-number*

no dst-port

Context

[Tree] (config>filter>ip-filter>entry>match dst-port)

[Tree] (config>filter>ipv6-filter>entry>match dst-port)

[Tree] (config>filter>ipv6-exception>entry>match dst-port)

[Tree] (config>filter>ip-exception>entry>match dst-port)

Full Context

configure filter ip-filter entry match dst-port

configure filter ipv6-filter entry match dst-port

configure filter ipv6-exception entry match dst-port

configure filter ip-exception entry match dst-port

Description

This command configures a destination TCP, UDP, or SCTP port number or port range for an IP filter or IP exception match criterion. An entry containing Layer 4 non-zero match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the Layer 4 information. Similarly an entry containing the "**dst-port eq** 0" match criterion, may match non-initial fragments when the destination port value is not present in a packet fragment and other match criteria are also met.

The **no** form of the command removes the destination port match criterion.

Default

no dst-port

Parameters

lt

Specifies that all port numbers less than the *dst-port-number* match.

gt

Specifies that all port numbers greater than the *dst-port-number* match.

eq

Specifies that the *dst-port-number* must be an exact match.

dst-port-number

Specifies the destination port number to be used as a match criteria expressed as a decimal integer, as well as in hexadecimal or binary format. The following value is for decimal integer format only.

Values 0 to 65535

port-list-name

Specifies to use a list of ports referred to by *port-list-name*, which is a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

dst-port-number dst-port-number

Specifies inclusive port range between two *dst-port-number* values.

Platforms

7705 SAR Gen 2

dst-port

Syntax

dst-port *value* [*mask*]

no dst-port

Context

[Tree] (config>system>security>mgmt-access-filter>ip-filter>entry dst-port)

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter>entry dst-port)

Full Context

configure system security management-access-filter ip-filter entry dst-port

configure system security management-access-filter ipv6-filter entry dst-port

Description

This command configures a destination TCP or UDP port number or port range for a management access filter match criterion.

The **no** form of this command removes the destination port match criterion.

Parameters

value

Specifies the destination TCP or UDP port number as match criteria.

Values 1 **to** 65535 (decimal)

mask

Specifies the mask used to specify a range of destination port numbers as the match criterion.

This 16 bit mask can be configured using the formats described in Table 30: Format Styles to Configure Mask:

Table 30: Format Styles to Configure Mask

Format Style	Format Syntax	Example
Decimal	DDDDD	63488
Hexadecimal	0xHHHH	0xF800
Binary	0bBBBBBBBBBBBBBBB	0b1111100000000000

To select a range from 1024 up to 2047, specify 1024 0xFC00 for value and mask.

Default **65535** (exact match)

Values 1 to 65535 (decimal)

Platforms

7705 SAR Gen 2

8.155 dst-zone

dst-zone

Syntax

[**no**] **dst-zone** {*std-zone-name* | *non-std-zone-name*}

Context

[Tree] (config>system>time dst-zone)

Full Context

configure system time dst-zone

Description

This command configures the start and end dates and offset for summer time or daylight savings time to override system defaults or for user defined time zones.

When configured, the time is adjusted by adding the configured offset when summer time starts and subtracting the configured offset when summer time ends.

If the time zone configured is listed in the Time Zones section, then the starting and ending parameters and offset do not need to be configured with this command unless it is necessary to override the system defaults. The command returns an error if the start and ending dates and times are not available either the Time Zones section on or entered as optional parameters in this command.

Up to five summer time zones may be configured, for example, for five successive years or for five different time zones. Configuring a sixth entry will return an error message. If no summer (daylight savings) time is supplied, it is assumed no summer time adjustment is required.

The **no** form of the command removes a configured summer (daylight savings) time entry.

Parameters

std-zone-name

Specifies the standard time zone name. The standard name must be a system-defined zone in the Time Zones section. For zone names in the table that have an implicit summer time setting, for example MDT for Mountain Daylight Saving Time, the remaining **start-date**, **end-date** and **offset** parameters need to be provided unless it is necessary to override the system defaults for the time zone.

Values ADT, NDT, AKDT, CDT, CEST, EDT, EEST, MDT, NZDT, PDT, WEST

non-std-zone-name

Specifies the non-standard time zone name. Create a user-defined name created using the zone. The name can be a maximum of 5 characters in length.

Platforms

7705 SAR Gen 2

8.156 duid-en

duid-en

Syntax

duid-en *hex-string*

no duid-en

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>host-ident duid-en)

Full Context

configure subscriber-mgmt local-user-db ipoe host host-identification duid-en

Description

This command configures the hexadecimal value for use in matching against the concatenation of enterprise number and identifier fields of DHCPv6 option CLIENTID (1) with DUID type = 2 (assigned by the vendor based on the enterprise number) in the DHCPv6 client message.

The **no** form of this command removes the client ID type duid-en from the configuration.

Default

no duid-en

Parameters

hex-string

Specifies the string in hexadecimal format, up to 254 hex nibbles.

Values 0x0 to 0xFFFFFFFF

Platforms

7705 SAR Gen 2

8.157 duid-ll-llt

duid-ll-llt

Syntax

duid-ll-llt *ieee-address*

no duid-ll-llt

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>host-ident duid-ll-llt)

Full Context

configure subscriber-mgmt local-user-db ipoe host host-identification duid-ll-llt

Description

This command configures the value for use in matching against the link-layer address field of DHCPv6 option CLIENTID (1) with DUID type = 3 (based on link-layer address) or DUID type = 1 (based on link-layer address plus time) and hardware type = 1 (Ethernet) in the DHCPv6 client message. For DUID type = 1, the time field is ignored.

The **no** form of this command removes the client ID type duid-ll-llt from the configuration.

Default

no duid-ll-llt

Parameters

ieee-address

Specifies the unicast MAC address of the client ID. This value cannot be all zeros.

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

Platforms

7705 SAR Gen 2

8.158 dup-detect

dup-detect

Syntax

dup-detect [**anti-spoof-mac** *mac-address*] **window** *minutes* **num-moves** *count* **hold-down** [*minutes* | **max**]

dup-detect anti-spoof-mac *mac-address* **window** *minutes* **num-moves** *count* **hold-down** [*minutes* | **max**] [**static-black-hole**]

Context

[Tree] (config>service>vpls>proxy-nd dup-detect)

[Tree] (config>service>vpls>proxy-arp dup-detect)

Full Context

configure service vpls proxy-nd dup-detect

configure service vpls proxy-arp dup-detect

Description

This command enables a mechanism that detects duplicate IPs and ARP/ND spoofing attacks. Attempts (relevant to dynamic and EVPN entry types) to add the same IP (different MAC) are monitored for **window** <*minutes*>. When <*count*> is reached within that **window**, the proxy-ARP/ND entry for the suspected IP is marked as duplicate. An alarm is also triggered. This condition is cleared when **hold-down** time expires (max does not expire) or a **clear** command is issued.

If the **anti-spoof-mac** is configured, the proxy-ARP/ND offending entry's MAC is replaced with this <*mac-address*> and advertised in an unsolicited GARP/NA for local SAP/SDP-bindings, and in EVPN to remote PEs. This mechanism assumes that the same **anti-spoof-mac** is configured in all the PEs for the same service and that traffic with destination **anti-spoof-mac** received on SAPs/SDP-bindings will be dropped. An ingress **mac-filter** may be configured to drop traffic to the **anti-spoof-mac**.

The **anti-spoof-mac** can also be combined with the **static-black-hole** option. To use a black-hole MAC entry for the **anti-spoof-mac** function in a proxy-ARP/ND service, the following must be configured:

- **static-black-hole** option for the **anti-spoof-mac**
- a static black-hole MAC using the same MAC address used for the **anti-spoof-mac: static-mac mac** <*mac-address*> **create black-hole** command.

When both **anti-spoof-mac** and **static-black-hole** commands are configured, the MAC is advertised in EVPN as Static. Locally, the MAC will be shown in the FDB as CStatic and associated with a black-hole.

The combination of the **anti-spoof-mac** and the **static-black-hole** options ensures that any frame arriving in the system with MAC DA=**anti-spoof-mac** will be discarded, regardless of the ingress endpoint type (SAP/SDP-binding or EVPN) and without the need for a filter.

If the user wants to redirect the traffic with MAC DA=**anti-spoof-mac** instead of discarding it, redirect filters should be configured on SAPs/SDP-bindings instead of the **static-black-hole** option.

If the **static-black-hole** option is not configured for the **anti-spoof-mac**, the behavior is as follows:

- The **anti-spoof-mac** is not programmed in the FDB.
- Any attempt to add a Static MAC (or any other MAC) with the **anti-spoof-mac** value will be rejected by the system.
- A mac-filter is needed to discard traffic with MAC DA=**anti-spoof-mac**.

Any changes to the configuration of **anti-spoof-mac** require proxy-arp or proxy-nd to first be shut down. Refer to "ARP/ND Snooping and Proxy Support" in the *7705 SAR Gen 2 Layer 2 Services and EVPN Guide* for more information.

Default

dup-detect window 3 num-moves 5 hold-down 9

Parameters

window *minutes*

Specifies the window size in minutes.

Values 1 to 15

Default 3

count

Specifies the number of moves required so that an entry is declared duplicate.

Values 3 to 10

Default 5

hold-down *minutes*

Specifies the hold-down time for a duplicate entry.

Values 2 to 60

Default 9

hold-down max

Specifies permanent hold-down time for a duplicate entry.

mac-address

Specifies the optional anti-spoof-mac to use.

Platforms

7705 SAR Gen 2

8.159 duplex

duplex

Syntax

duplex {**full** | **half**}

Context

[Tree] (config>port>ethernet duplex)

Full Context

configure port ethernet duplex

Description

This command configures the duplex of a Fast Ethernet port when autonegotiation is disabled.

This configuration command allows for the configuration of the duplex mode of a Fast Ethernet port. If the port is configured to autonegotiate this parameter is ignored.

Default

duplex full

Parameters

full

Sets the link to full duplex mode.

half

Sets the link to half duplex mode.

Platforms

7705 SAR Gen 2

duplex

Syntax

duplex {**full** | **half**}

Context

[Tree] (bof duplex)

Full Context

bof duplex

Description

This command configures the duplex mode of the CPM management Ethernet port when autonegotiation is disabled in the running configuration and the Boot Option File (BOF). If the port is configured to autonegotiate this parameter will be ignored.

Parameters

full

Sets the link to full duplex mode.

half

Sets the link to half duplex mode.

Platforms

7705 SAR Gen 2

8.160 dwdm

dwdm

Syntax

dwdm

Context

[Tree] (config>port dwdm)

Full Context

configure port dwdm

Description

This command configures the Dense Wavelength Division Multiplexing (DWDM) parameters.

Platforms

7705 SAR Gen 2

8.161 dynamic

dynamic

Syntax

dynamic *ip-address* [**create**]

no dynamic *ip-address*

Context

[Tree] (config>service>vpls>proxy-nd dynamic)

[Tree] (config>service>vpls>proxy-arp dynamic)

Full Context

configure service vpls proxy-nd dynamic

configure service vpls proxy-arp dynamic

Description

This command creates a dynamic IP that can be associated to a MAC list. The configured dynamic IP is only converted to a dynamic entry when the resolve process for the IP has passed successfully.

A summary of the IP resolution process is as follows:

- A resolve message is sent for the configured IP as soon as the dynamic IP is configured. The message is sent with a configurable frequency of 1 to 60 minutes (using the **resolve** command); the default value is 5 minutes. The actual resolve interval is a "tittered" value of the configured interval.
- The resolve message is an ARP-request or NS message flooded to all the non-EVPN endpoints in the service, irrespective of the status of the **unknown-arp-request-flood-evpn** or **unknown-ns-flood-evpn** commands. The router sends resolve messages at the configured frequency until a dynamic entry for the IP is created in the proxy-ARP or proxy-ND table. The IP entry is created only if all of the following conditions are true.
  - An ARP, GARP, or NA message is received for the configured IP.
  - The associated MAC exists in the configured MAC list for the IP.

    If the MAC list is empty or not configured, the router does not create an entry for the IP.
- After a dynamic entry is created in the proxy-ARP or proxy-ND table, the IP->MAC entry is advertised in the EVPN.

The **no** form of the command deletes the dynamic IP and the associated proxy-ARP or proxy-ND entry, if it exists.

Parameters

ip-address

Specifies the IPv4 or IPv6 address.

Values ip-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

where:

x: [0 to FFFF]H

d: [0 to 255]D

Platforms

7705 SAR Gen 2

8.162 dynamic-arp-populate

dynamic-arp-populate

Syntax

[**no**] **dynamic-arp-populate**

Context

[Tree] (config>service>vpls>proxy-arp dynamic-arp-populate)

Full Context

configure service vpls proxy-arp dynamic-arp-populate

Description

This command enables the addition of dynamic entries to the proxy-ARP table (disabled by default). When executed, the system will populate proxy-ARP entries from snooped GARP/ARP messages on SAPs/SDP-bindings. These entries will be shown as dynamic.

When disabled, dynamic-arp entries will be flushed from the proxy-ARP table. Enabling dynamic-arp-populate is only recommended in networks with a consistent configuration of this command in all the PEs.

Default

no dynamic-arp-populate

Platforms

7705 SAR Gen 2

8.163 dynamic-bgp

dynamic-bgp

Syntax

[**no**] **dynamic-bgp**

Context

[Tree] (config>router>static-route-entry>black-hole dynamic-bgp)

Full Context

configure router static-route-entry black-hole dynamic-bgp

Description

This optional command controls the behavior of the associated static route so that if a matching BGP route to the same exact prefix is present in BGP, the static route's nexthop is set to the BGP's nexthop value. If there is no matching active BGP route, the static route's nexthop is set to be a black-hole nexthop.

Default

no dynamic-bgp

Platforms

7705 SAR Gen 2

8.164 dynamic-bypass

dynamic-bypass

Syntax

dynamic-bypass [**enable** | **disable**]

no dynamic-bypass

Context

[Tree] (config>router>mpls dynamic-bypass)

Full Context

configure router mpls dynamic-bypass

Description

This command disables the creation of dynamic bypass LSPs in FRR. One or more manual bypass LSPs must be configured to protect the primary LSP path at the PLR nodes.

Default

dynamic-bypass enable

Platforms

7705 SAR Gen 2

8.165 dynamic-cost

dynamic-cost

Syntax

[**no**] **dynamic-cost**

Context

[Tree] (config>lag dynamic-cost)

Full Context

configure lag dynamic-cost

Description

This command enables OSPF or ISIS costing of a Link Aggregation Group (LAG) based on the available aggregated, operational bandwidth.

The path cost is dynamically calculated based on the interface bandwidth. OSPF path cost can be changed through the interface metric or the reference bandwidth.

If dynamic cost is configured, then costing is applied based on the total number of links configured and the cost advertised is inversely proportional to the number of links available at the time. This is provided that the number of links that are up exceeds the configured LAG threshold value at which time the configured threshold action determines if, and at what cost, this LAG will be advertised.

For example: Assume a physical link in OSPF has a cost associated with it of 100, and the LAG consists of four physical links. The cost associated with the logical link is 25. If one link fails then the cost would automatically be adjusted to 33.

If dynamic cost is not configured and OSPF autocost is configured, then costing is applied based on the total number of links configured. This cost will remain static provided the number of links that are up exceeds the configured LAG threshold value at which time the configured threshold action determines if and at what cost this LAG will be advertised.

If dynamic-cost is configured and OSPF autocost is not configured, the cost is determined by the cost configured on the OSPF metric provided the number of links available exceeds the configured LAG threshold value at which time the configured threshold action determines if this LAG will be advertised.

If neither dynamic-cost nor OSPF autocost are configured, the cost advertised is determined by the cost configured on the OSPF metric provided the number of links available exceeds the configured LAG threshold value at which time the configured threshold action determines if this LAG will be advertised.

The **no** form of this command removes dynamic costing from the LAG.

Default

no dynamic-cost

Platforms

7705 SAR Gen 2

8.166 dynamic-egress-label-limit

dynamic-egress-label-limit

Syntax

[**no**] **dynamic-egress-label-limit**

Context

[Tree] (config>service>vpls>bgp-evpn>mpls dynamic-egress-label-limit)

[Tree] (config>service>epipe>bgp-evpn>mpls dynamic-egress-label-limit)

[Tree] (config>service>vprn>bgp-evpn>mpls dynamic-egress-label-limit)

[Tree] (config>service>vprn>bgp-ipvpn>mpls dynamic-egress-label-limit)

Full Context

configure service vpls bgp-evpn mpls dynamic-egress-label-limit

configure service epipe bgp-evpn mpls dynamic-egress-label-limit

configure service vprn bgp-evpn mpls dynamic-egress-label-limit

configure service vprn bgp-ipvpn mpls dynamic-egress-label-limit

Description

This command relaxes the egress MPLS label limit check when resolving BGP next hops in the tunnel table.

For VPRN services, the OAM label is never computed and, therefore, one more egress label is allowed.

For EVPN (Epipe and VPLS) services, the system only computes the control word and ESI label if they are used. For the control word, the system reduces the egress label limit by one label if the control word is configured in the service. When configured, the ESI label is not counted for Epipes or VPLS services without an ES.

The **no** form of this command, for EVPN, Epipe, and VPLS services, always accounts for the ESI label and control word.

Default

no dynamic-egress-label-limit

Platforms

7705 SAR Gen 2

8.167 dynamic-enforcement-policer-pool

dynamic-enforcement-policer-pool

Syntax

[**no**] **dynamic-enforcement-policer-pool** *number-of-policers*

Context

[Tree] (config>card>fp>ingress>dist-cpu-protection dynamic-enforcement-policer-pool)

Full Context

configure card fp ingress dist-cpu-protection dynamic-enforcement-policer-pool

Description

This command reserves a set of policers for use as dynamic enforcement policers for the Distributed CPU Protection (DCP) feature. Policers are allocated from this pool and instantiated as per-object-per-protocol dynamic enforcement policers after a local monitor is triggered for an object (such as a SAP or Network Interface). Any change to this configured value automatically clears the high water mark, timestamp, and failed allocation counts as seen under "show card x fp y dist-cpu-protection" and in the tmnxFpDcpDynEnfrcPlcrStatTable in the TIMETRA-CHASSIS-MIB. Decreasing this value to below the currently used/allocated number causes all dynamic policers to be returned to the free pool (and traffic returns to the local monitors).

Default

no dynamic-enforcement-policer-pool

Parameters

number-of-policers

specifies the number of policers to be reserved.

Values 1000 to 32000

Platforms

7705 SAR Gen 2

8.168 dynamic-keying

dynamic-keying

Syntax

[**no**] **dynamic-keying**

Context

[Tree] (config>router>if>ipsec>ipsec-tunnel dynamic-keying)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel dynamic-keying)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel dynamic-keying)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel dynamic-keying)

[Tree] (config>ipsec>trans-mode-prof dynamic-keying)

Full Context

configure router interface ipsec ipsec-tunnel dynamic-keying

configure service ies interface ipsec ipsec-tunnel dynamic-keying

configure service vprn interface ipsec ipsec-tunnel dynamic-keying

configure service vprn interface sap ipsec-tunnel dynamic-keying

configure ipsec ipsec-transport-mode-profile dynamic-keying

Description

This command enables dynamic keying for the IPsec tunnel.

The **no** form of this command disables dynamic keying.

Platforms

7705 SAR Gen 2

8.169 dynamic-mbs

dynamic-mbs

Syntax

[**no**] **dynamic-mbs**

Context

[Tree] (config>qos>qgrps>egr>qgrp>queue dynamic-mbs)

Full Context

configure qos queue-group-templates egress queue-group queue dynamic-mbs

Description

This command enables support for dynamically modifying the MBS size of a queue using HQoS in order to maintain the maximum latency for traffic in the queue based on the queue's configured MBS and the ratio of its operational PIR to its administrative PIR. As the HQoS algorithm updates the operational PIR, by reducing or increasing it, the MBS of the queue is adjusted accordingly.

The configuration of dynamic MBS and the configuration of queue depth monitoring (**monitor-queue-depth** command) are mutually exclusive. Queue depth monitoring is an override on the queue where the queue group is applied.

The **no** form of this command disables dynamic MBS resizing.

Default

no dynamic-mbs

Platforms

7705 SAR Gen 2

8.170 dynamic-nd-populate

dynamic-nd-populate

Syntax

[**no**] **dynamic-nd-populate**

Context

[Tree] (config>service>vpls>proxy-nd dynamic-nd-populate)

Full Context

configure service vpls proxy-nd dynamic-nd-populate

Description

This command enables the addition of dynamic entries to the proxy-ND table. The command is disabled by default. When executed, the system will populate proxy-ND entries from snooped Neighbor Advertisement (NA) messages on SAPs/SDP-bindings, in addition to the entries coming from EVPN (if the EVPN is enabled). These entries will be shown as dynamic, as opposed to EVPN entries or static entries.

When disabled, dynamic-ND entries will be flushed from the proxy-ND table. Enabling **dynamic-nd-populate** is only recommended in networks with a consistent configuration of this command in all the PEs.

Default

no dynamic-nd-populate

Platforms
7705 SAR Gen 2

8.171 dynamic-neighbor

dynamic-neighbor

Syntax
dynamic-neighbor

Context
[Tree] (config>service>vprn>bgp>group dynamic-neighbor)

Full Context
configure service vprn bgp group dynamic-neighbor

Description
Commands in this context configure dynamic BGP sessions for a peer group.

Platforms
7705 SAR Gen 2

dynamic-neighbor

Syntax
dynamic-neighbor

Context
[Tree] (config>router>bgp>group dynamic-neighbor)

Full Context
configure router bgp group dynamic-neighbor

Description
Commands in this context configure dynamic BGP sessions for a peer group.

Platforms
7705 SAR Gen 2

8.172 dynamic-neighbor-limit

dynamic-neighbor-limit

Syntax

dynamic-neighbor-limit *peers*

no dynamic-neighbor-limit

Context

[Tree] (config>service>vprn>bgp>group dynamic-neighbor-limit)

[Tree] (config>service>vprn>bgp dynamic-neighbor-limit)

Full Context

configure service vprn bgp group dynamic-neighbor-limit

configure service vprn bgp dynamic-neighbor-limit

Description

This command configures the maximum number of dynamic BGP sessions that are accepted from remote peers associated with the entire BGP instance or a specific peer group. If accepting a new dynamic session would cause either the group limit or the instance limit to be exceeded, then the new session attempt is rejected and a Notification message is sent back to the remote peer.

The **no** form of this command removes the limit on the number of dynamic sessions.

Default

no dynamic-neighbor-limit

Parameters

peers

Specifies the maximum number of dynamic BGP sessions.

Values 1 to 8192

Platforms

7705 SAR Gen 2

dynamic-neighbor-limit

Syntax

dynamic-neighbor-limit *peers*

no dynamic-neighbor-limit

Context

[Tree] (config>router>bgp dynamic-neighbor-limit)

[Tree] (config>router>bgp>group dynamic-neighbor-limit)

Full Context

configure router bgp dynamic-neighbor-limit

configure router bgp group dynamic-neighbor-limit

Description

This command configures the maximum number of dynamic BGP sessions that will be accepted from remote peers associated with the entire BGP instance or a specific peer group. If accepting a new dynamic session would cause either the group limit or the instance limit to be exceeded, then the new session attempt is rejected and a Notification message is sent back to the remote peer.

The **no** form of this command removes the limit on the number of dynamic sessions.

Default

no dynamic-neighbor-limit

Parameters

peers

Specifies the maximum number of dynamic BGP sessions.

Values 1 to 8192

Platforms

7705 SAR Gen 2

8.173 dynamic-parameters

dynamic-parameters

Syntax

dynamic-parameters

Context

[Tree] (config>sys>security>dist-cpu-protection>policy>protocol dynamic-parameters)

Full Context

configure system security dist-cpu-protection policy protocol dynamic-parameters

Description

The dynamic-parameters are used to instantiate a dynamic enforcement policer for the protocol when the associated **local-monitoring-policer** is considered as exceeding its rate parameters (at the end of a minimum monitoring time of 60 seconds).

Platforms

7705 SAR Gen 2

8.174 dynamic-tunnel-redundant-next-hop

dynamic-tunnel-redundant-next-hop

Syntax

dynamic-tunnel-redundant-next-hop *ip-address*

no dynamic-tunnel-redundant-next-hop

Context

[Tree] (config>service>vprn>if dynamic-tunnel-redundant-next-hop)

[Tree] (config>service>ies>if dynamic-tunnel-redundant-next-hop)

Full Context

configure service vprn interface dynamic-tunnel-redundant-next-hop

configure service ies interface dynamic-tunnel-redundant-next-hop

Description

This command specifies redundant next-hop address on a public or private IPsec interface (with public or private tunnel-sap) for dynamic IPsec tunnel. The specified next-hop address is used by a standby node to shunt traffic to master in case it receives the address.

The next-hop address is resolved in the routing table of a corresponding service.

Default

no dynamic-tunnel-redundant-next-hop

Parameters

ip-address

Specifies the dynamic ISA tunnel redundant next-hop address.

Platforms

7705 SAR Gen 2

8.175 dynmldp

dynmldp

Syntax

dynmldp [**detail**]

no dynmldp

Context

[Tree] (debug>router>pim dynmldp)

Full Context

debug router pim dynmldp

Description

This command enables debugging for dynamic MLDP.

The **no** form of this command disables dynamic MLDP debugging.

Parameters

detail

Debugs detailed dynamic MLDP information.

Platforms

7705 SAR Gen 2

9 e Commands

9.1 e-counters

e-counters

Syntax

e-counters [**all**]

no e-counters

Context

[Tree] (config>log>acct-policy>cr>ref-queue e-counters)

[Tree] (config>log>acct-policy>cr>queue e-counters)

[Tree] (config>log>acct-policy>cr>ref-policer e-counters)

[Tree] (config>log>acct-policy>cr>policer e-counters)

Full Context

configure log accounting-policy custom-record ref-queue e-counters

configure log accounting-policy custom-record queue e-counters

configure log accounting-policy custom-record ref-policer e-counters

configure log accounting-policy custom-record policer e-counters

Description

This command configures egress counter parameters for this custom record.

The **no** form of this command reverts all egress counters to their default value.

Default

e-counters

Parameters

all

Specifies that all egress counters should be included.

Platforms

7705 SAR Gen 2

9.2 eapol-destination-address

eapol-destination-address

Syntax

eapol-destination-address *mac*

no eapol-destination-address

Context

[Tree] (config>port>ethernet>dot1x>macsec>sub-port eapol-destination-address)

Full Context

configure port ethernet dot1x macsec sub-port eapol-destination-address

Description

The EAPoL destination MAC address uses a destination multicast MAC address of 01:80:C2:00:00:03. Some networks cannot tunnel this packet over the network and consume these packets, causing the MKA session to fail. This command can change the destination MAC of the EAPoL to the unicast address of the MACsec peer, and as such, the EAPoL and MKA signaling will be unicasted between two peers.

The **no** form of this command returns the value to the default.

Default

no eapol-destination-address

Parameters

mac

Specifies the desired destination MAC address to be used by the EAPOL MKA packets of this sub-port.

Values aa:bb:cc:dd:ee:ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers.

Platforms

7705 SAR Gen 2

9.3 ebgp-default-reject-policy

ebgp-default-reject-policy

Syntax

ebgp-default-reject-policy [**import**] [**export**]

no ebgp-default-reject-policy

Context

[Tree] (config>service>vprn>bgp ebgp-default-reject-policy)

[Tree] (config>service>vprn>bgp>group>neighbor ebgp-default-reject-policy)

[Tree] (config>service>vprn>bgp>group ebgp-default-reject-policy)

Full Context

configure service vprn bgp ebgp-default-reject-policy

configure service vprn bgp group neighbor ebgp-default-reject-policy

configure service vprn bgp group ebgp-default-reject-policy

Description

This command configures the default import and export policy behavior for EBGP neighbors.

The **no** form of this command removes the default import and export policy behavior.

Default

no ebgp-default-reject-policy

Parameters

import

Specifies the default reject import policy for EBGP neighbors.

export

Specifies the default reject export policy for EBGP neighbors.

Platforms

7705 SAR Gen 2

ebgp-default-reject-policy

Syntax

ebgp-default-reject-policy [**import**] [**export**]

no ebgp-default-reject-policy

Context

[Tree] (config>router>bgp ebgp-default-reject-policy)

[Tree] (config>router>bgp>group>neighbor ebgp-default-reject-policy)

[Tree] (config>router>bgp>group ebgp-default-reject-policy)

Full Context

configure router bgp ebgp-default-reject-policy

configure router bgp group neighbor ebgp-default-reject-policy

configure router bgp group ebgp-default-reject-policy

Description

This command configures the default import and export policy behavior for EBGP neighbors.

The **no** form of this command removes the default import and export policy behavior.

Default

no ebgp-default-reject-policy

Parameters

import

Specifies the default reject import policy for EBGP neighbors.

export

Specifies the default reject export policy for EBGP neighbors.

Platforms

7705 SAR Gen 2

9.4 ebgp-ibgp-equal

ebgp-ibgp-equal

Syntax

ebgp-ibgp-equal [**ipv4**] [**ipv6**] [**label-ipv4**] [**label-ipv6**]

no ebgp-ibgp-equal

Context

[Tree] (config>service>vprn>bgp>best-path-selection ebgp-ibgp-equal)

Full Context

configure service vprn bgp best-path-selection ebgp-ibgp-equal

Description

This command instructs the BGP decision process to ignore the difference between EBGP and IBGP routes in selecting the best path and eligible multipaths (if multipath and ECMP are enabled). The result is a form of EIBGP load-balancing in a multipath scenario.

The operator can apply the behavior selectively to only certain types of routes by specifying one or more address family names in the command.

The **no** form of this command configures the router in the BGP decision process to prefer an EBGP learned route over an IBGP learned route.

Default

no ebgp-ibgp-equal

Parameters

ipv4

Specifies that the command should be applied to unlabeled unicast IPv4 routes.

ipv6

Specifies that the command should be applied to unlabeled unicast IPv6 routes.

label-ipv4

Specifies that the command should be applied to labeled IPv4 routes.

label-ipv6

Specifies that the command should be applied to labeled IPv6 routes.

Platforms

7705 SAR Gen 2

ebgp-ibgp-equal

Syntax

ebgp-ibgp-equal [**ipv4**] [**ipv6**] [**label-ipv4**] [**label-ipv6**] [**vpn-ipv4**] [**vpn-ipv6**]

[**evpn**]

no ebgp-ibgp-equal

Context

[Tree] (config>router>bgp>best-path-selection ebgp-ibgp-equal)

Full Context

configure router bgp best-path-selection ebgp-ibgp-equal

Description

This command instructs the BGP decision process to ignore the difference between EBGP and IBGP routes in selecting the best path and eligible multipaths (if multipath and ECMP are enabled). The result is a form of EIBGP load balancing in a multipath scenario.

The behavior can be applied selectively to only certain types of routes by specifying one or more address family names in the command. If no families are specified, this command applies to IPv4, IPv6, label-IPv4, label-IPv6, VPN-IPv4, VPN-IPv6, and EVPN routes.

The **no** form of this command configures the router in the BGP decision process to prefer an EBGP learned route over an IBGP learned route.

Default

no ebgp-ibgp-equal

Parameters

ipv4

Specifies that the command should be applied to unlabeled unicast IPv4 routes.

ipv6

Specifies that the command should be applied to unlabeled unicast IPv6 routes.

label-ipv4

Specifies that the command should be applied to labeled unicast IPv4 routes.

label-ipv6

Specifies that the command should be applied to labeled unicast IPv6 routes.

vpn-ipv4

Specifies that the command should be applied to IPv4 VPN routes.

vpn-ipv6

Specifies that the command should be applied to IPv6 VPN routes.

evpn

Specifies that the command should be applied to EVPN routes.

Platforms

7705 SAR Gen 2

9.5 ecdsa

ecdsa

Syntax

ecdsa

Context

[Tree] (config>system>security>user>public-keys ecdsa)

Full Context

configure system security user public-keys ecdsa

Description

This command allows the user to enter the context to configure ECDSA public keys.

Platforms

7705 SAR Gen 2

9.6 ecdsa-key

ecdsa-key

Syntax

ecdsa-key *key-id* [**create**]

no ecdsa-key *key-id*

Context

[Tree] (config>system>security>user>public-keys>ecdsa ecdsa-key)

Full Context

configure system security user public-keys ecdsa ecdsa-key

Description

This command creates an ECDSA public key and associates it with the username. Multiple public keys can be associated with the user. The key ID is used to identify these keys for the user.

Parameters

create

Keyword used to create an ECDSA key. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

key-id

Specifies the key identifier.

Values 1 to 32

Platforms

7705 SAR Gen 2

9.7 echo

echo

Syntax

echo [*text-to-echo*] [*extra-text-to-echo*] [*more-text*]

Context

[Tree] (echo)

Full Context

echo

Description

This command echoes arguments on the command line. The primary use of this command is to allow messages to be displayed to the screen in files executed with the **exec** command.

Parameters

text-to-echo

Specifies a text string to be echoed, up to 256 characters.

extra-text-to-echo

Specifies more text to be echoed, up to 256 characters.

more-text

Specifies more text to be echoed, up to 256 characters.

Platforms

7705 SAR Gen 2

9.8 echo-receive

echo-receive

Syntax

echo-receive *echo-interval*

no echo-receive

Context

[Tree] (config>router>bfd>bfd-template echo-receive)

Full Context

configure router bfd bfd-template echo-receive

Description

This command sets the minimum echo receive interval, in milliseconds, for a session. This is not used by a BFD session for MPLS-TP.

The **no** form of this command reverts to the default value.

Default

echo-receive 100

Parameters

echo-interval

Specifies the echo receive interval.

Values 100 ms to 100,000 ms in 1 ms increments

Default 100

Platforms

7705 SAR Gen 2

9.9 ecmp

ecmp

Syntax

ecmp *max-ecmp-routes*

Context

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel ecmp)

[Tree] (config>service>epipe>bgp-evpn>mpls ecmp)

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel ecmp)

Full Context

configure service epipe bgp-evpn mpls auto-bind-tunnel ecmp

configure service epipe bgp-evpn mpls ecmp

configure service vpls bgp-evpn mpls auto-bind-tunnel ecmp

Description

When configured in a VPLS service, this command controls the number of paths that are allowed to reach a specified MAC address when that MAC in the FDB is associated to a remote all-active multi-homed ES.

The configuration of two or more ECMP paths to a specified MAC enables the aliasing function described in RFC 7432.

When used in an Epipe service, this command controls the number of paths that are allowed to reach a specified remote Ethernet tag that is associated to an ES destination.

Default

ecmp 1

Parameters

max-ecmp-routes

Specifies the number of paths allowed to the same multi-homed MAC address or Ethernet tag.

Values 1 to 32

Platforms

7705 SAR Gen 2

ecmp

Syntax

ecmp *max-ecmp-routes*

no ecmp

Context

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel ecmp)

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel ecmp)

Full Context

configure service vprn bgp-ipvpn mpls auto-bind-tunnel ecmp

configure service vprn bgp-evpn mpls auto-bind-tunnel ecmp

Description

This command configures the maximum number of tunnels that may be used as ECMP next-hops for the VPRN. This value overrides any values that are configured using the **config**>**service**>**vprn**>**ecmp** command.

The **no** form of this command removes the configured overriding value, and the value configured using the **config**>**service**>**vprn**>**ecmp** command is used.

Default

ecmp 1

Parameters

max-ecmp-routes

Specifies the maximum number of tunnels that may be used as ECMP next-hops for the VPRN.

Values 1 to 32

Default 1

Platforms

7705 SAR Gen 2

ecmp

Syntax

ecmp *max-ecmp-routes*

no ecmp

Context

[Tree] (config>router ecmp)

Full Context

configure router ecmp

Description

This command enables ECMP and configures the number of routes for path sharing; for example, the value 2 means two equal cost routes are used for cost sharing.

ECMP can be used only for routes with the same preference and same protocol.

If available ECMP routes at the best preference exceed the maximum ECMP routes allowed, the system selects using the following criteria:

1. The system selects the lowest next hop router ID.
2. If the next hop goes to the same neighbor, the system selects the next hop with the lowest interface index.

The **no** form of this command disables ECMP path sharing. If ECMP is disabled and multiple routes are available at the best preference and equal cost, the route with the lowest next-hop IP address is used.

Default

no ecmp

Parameters

max-ecmp-routes

Specifies the maximum number of equal cost routes allowed on this routing table instance, expressed as a decimal integer. Setting ECMP *max-ecmp-routes* to 1 yields the same result as entering **no ecmp**.

Values 1 to 64

Platforms

7705 SAR Gen 2

ecmp

Syntax

ecmp *max-ecmp-routes*

no ecmp

Context

[Tree] (config>service>vprn ecmp)

Full Context

configure service vprn ecmp

Description

This command enables equal-cost multipath (ECMP) and configures the number of routes for path sharing. For example, the value of 2 means that 2 equal cost routes are used for cost sharing.

ECMP groups form when the system routes to the same destination with equal cost values. Routing table entries can be entered manually (as static routes), or they can be formed when neighbors are discovered and routing table information is exchanged by routing protocols. The system can balance traffic across the groups with equal costs.

ECMP can only be used for routes learned with the same preference and same protocol.

If available ECMP routes at the best preference exceed the maximum ECMP routes allowed, the system selects using the following criteria:

1. The system selects the lowest next hop router ID.
2. If the next hop goes to the same neighbor, the system selects the next hop with the lowest interface index.

The **no** form of this command disables ECMP path sharing. If ECMP is disabled and multiple routes are available at the best preference and equal cost, the newly updated route is used.

Default

no ecmp

Parameters

max-ecmp-routes

Specifies the maximum number of routes for path sharing.

Values 1 to 64

Platforms

7705 SAR Gen 2

ecmp

Syntax

ecmp

Context

[Tree] (config>service>vprn>auto-bind-tunnel ecmp)

Full Context

configure service vprn auto-bind-tunnel ecmp

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

9.10 ecmp-unequal-cost

ecmp-unequal-cost

Syntax

[**no**] **ecmp-unequal-cost**

Context

[Tree] (config>service>vprn ecmp-unequal-cost)

Full Context

configure service vprn ecmp-unequal-cost

Description

This command relaxes the constraint that ECMP multipaths must have the same IGP cost to reach the BGP next-hop. When VPN routes for the same IP prefix are imported into a VPRN service, they are eligible to be used as multipaths. The resulting route is programmed as an ECMP IP route.

The BGP best path selection algorithm is the basis for choosing the set of imported VPN routes that can be combined to form an ECMP route. Normally (unless an **ignore-nh-metric** command is configured), the BGP decision process gives higher preference to VPN routes with a lower next-hop cost if other, more significant criteria, are tied. In these circumstances, a VPN route cannot be an eligible multipath if it does not have the same next-hop cost as the best VPN route. Configuring this command removes this restriction and allows the multipaths to have different (meaning lower) next-hop costs than the best route. This broadens the applicability of multipath and can result in better load balancing in the network.

This command applies only to the following types of routes imported by a VPRN.

- vpn-ipv4
- vpn-ipv6
- mcast-vpn-ipv4
- mcast-vpn-ipv6

The **no** form of this command restores the default behavior that requires next-hop costs of multipaths to be equal, unless the next-hop cost is completely removed from the BGP decision process.

Default

ecmp-unequal-cost

Platforms

7705 SAR Gen 2

9.11 edge-port

edge-port

Syntax

[**no**] **edge-port**

Context

[Tree] (config>service>template>vpls-sap-template>stp edge-port)

[Tree] (config>service>vpls>spoke-sdp>stp edge-port)

[Tree] (config>service>vpls>sap>stp edge-port)

Full Context

configure service template vpls-sap-template stp edge-port

configure service vpls spoke-sdp stp edge-port

configure service vpls sap stp edge-port

Description

This command configures the SAP or SDP as an edge or non-edge port. If **auto-edge** is enabled for the SAP, this value will be used only as the initial value.

Note:
The function of the **edge-port** command is similar to the **rapid-start** command. It tells RSTP that it is on the edge of the network (for example, there are no other bridges connected to that port) and, as a consequence, it can immediately transition to a forwarding state if the port becomes available.

RSTP, however, can detect that the actual situation is different from what **edge-port** may indicate.

Initially, the value of the SAP or spoke-SDP parameter is set to edge-port. This value will change if:

- A BPDU is received on that port. This means that after all there is another bridge connected to this port. Then the edge-port becomes disabled.
- If auto-edge is configured and no BPDU is received within a certain period of time, RSTP concludes that it is on an edge and enables the edge-port.

The **no** form of this command returns the edge port setting to the default value.

Default

no edge-port

Platforms

7705 SAR Gen 2

edge-port

Syntax

[**no**] **edge-port**

Context

[Tree] (config>service>pw-template>stp edge-port)

Full Context

configure service pw-template stp edge-port

Description

This command configures the SAP or SDP as an edge or non-edge port. If **auto-edge** is enabled for the SAP, this value will be used only as the initial value.

Note:
The **edge-port** command tells RSTP that it is on the edge of the network (for example, there are no other bridges connected to that port) and, as a consequence, it can immediately transition to a forwarding state if the port becomes available.

RSTP, however, can detect that the actual situation is different from what **edge-port** may indicate.

Initially, the value of the SAP or spoke SDP parameter is set to edge-port. This value will change if:

- A BPDU is received on that port. This means that after all there is another bridge connected to this port. Then the edge-port becomes disabled.
- If auto-edge is configured and no BPDU is received within a certain period of time, RSTP concludes that it is on an edge and enables the edge-port.

The **no** form of this command returns the edge port setting to the default value.

Default

no edge-port

Platforms

7705 SAR Gen 2

9.12 edit

edit

Syntax

edit [**exclusive**]

Context

[Tree] (candidate edit)

Full Context

candidate edit

Description

This command enables the edit-cfg mode where changes can be made to the candidate configuration and sets the edit-point to the end of the candidate. In edit-cfg mode the CLI prompt contains **edit-cfg** near the root of the prompt. Commands in the **candidate** CLI branch, except **candidate edit**, are available only when in edit-cfg mode.

Parameters

exclusive

Allows a user to exclusively create a candidate configuration by blocking other users (and other sessions of the same user) from entering edit-cfg mode. Exclusive edit-cfg mode can

only be entered if the candidate configuration is empty and no user is in edit-cfg mode. Once a user is in exclusive edit-cfg mode no other users/sessions are allowed in edit-cfg mode. The user must either commit or discard the exclusive candidate before leaving exclusive edit-cfg mode. If the CLI session times out while a user is in exclusive edit-cfg mode then the contents of the candidate are discarded. The **admin disconnect** command can be used to force a user to disconnect (and to clear the contents of the candidate) if they have the candidate locked.

Platforms

7705 SAR Gen 2

9.13 edit-config

edit-config

Syntax

[**no**] **edit-config**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization edit-config)

Full Context

configure system security profile netconf base-op-authorization edit-config

Description

This command enables the NETCONF <edit-config> RPC.

The **no** form of this command disables the RPC.

Default

no edit-config

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

9.14 ee-revocation

ee-revocation

Syntax

ee-revocation primary *primary* **secondary** *secondary*

Context

[Tree] (config>system>security>tls>server-tls-profile>status-verify ee-revocation)

[Tree] (config>system>security>tls>client-tls-profile>status-verify ee-revocation)

Full Context

configure system security tls server-tls-profile status-verify ee-revocation

configure system security tls client-tls-profile status-verify ee-revocation

Description

This command configures the method used to verify the revocation status of the TLS end-entity (EE) certificate.

Parameters

primary

Specifies the primary method.

Values ocsp, crl

Default crl

secondary

Specifies the secondary method.

Values ocsp, crl, none

Default none

Platforms

7705 SAR Gen 2

9.15 egr-ip-load-balancing

egr-ip-load-balancing

Syntax

egr-ip-load-balancing {**source** | **destination** | **inner-ip**}

no egr-ip-load-balancing

Context

[Tree] (config>service>ies>if>load-balancing egr-ip-load-balancing)

Full Context

configure service ies interface load-balancing egr-ip-load-balancing

Description

This command specifies whether to include the source address or destination address or both in the LAG/ECMP hash on IP interfaces. Additionally, when l4-load-balancing is enabled, the command also applies to the inclusion of source/destination port in the hash inputs.

The **no** form of this command includes both source and destination parameters.

Default

no egr-ip-load-balancing

Parameters

source

Specifies using the source address and, if l4-load balancing is enabled, the source port in the hash, ignore destination address/port.

destination

Specifies using the destination address and, if l4-load balancing is enabled, the destination port in the hash, ignore source address/port.

inner-ip

Specifies using the inner IP header parameters instead of the outer IP header parameters in the LAG/ECMP hash for IPv4 encapsulated traffic.

Platforms

7705 SAR Gen 2

egr-ip-load-balancing

Syntax

egr-ip-load-balancing {**source** | **destination** | **inner-ip**}

no egr-ip-load-balancing

Context

[Tree] (config>service>vprn>if>load-balancing egr-ip-load-balancing)

Full Context

configure service vprn interface load-balancing egr-ip-load-balancing

Description

This command specifies whether to include the source address or destination address or both in the LAG/ECMP hash on IP interfaces. Additionally, when l4-load-balancing is enabled, the command also applies to the inclusion of source/destination port in the hash inputs.

The **no** form of this command includes both source and destination parameters.

Default

no egr-ip-load-balancing

Parameters

source

Specifies using the source address and (if l4-load balancing is enabled) source port in the hash, ignore destination address/port.

destination

Specifies using the destination address and (if l4-load balancing is enabled) destination port in the hash, ignore source address/port.

inner-ip

Specifies use of the inner IP header parameters instead of outer IP header parameters in LAG/ECMP hash for IPv4 encapsulated traffic.

Platforms

7705 SAR Gen 2

egr-ip-load-balancing

Syntax

egr-ip-load-balancing {**source** | **destination** | **inner-ip**}

no egr-ip-load-balancing

Context

[Tree] (config>router>if>load-balancing egr-ip-load-balancing)

Full Context

configure router interface load-balancing egr-ip-load-balancing

Description

This command specifies whether to include source address or destination address or both in LAG/ECMP hash on IP interfaces. Additionally, when l4-load-balancing is enabled the command applies also to inclusion of source/destination port in the hash inputs.

The **no** form of this command includes both source and destination parameters.

Default

no egr-ip-load-balancing

Parameters

source

Specifies using source address and (if l4-load balancing is enabled) source port in the hash, ignore destination address/port

destination

Specifies using destination address and (if l4-load balancing is enabled) destination port in the hash, ignore source address/port.

inner-ip

Specifies use of the inner IP header parameters instead of outer IP header parameters in LAG/ECMP hash for IPv4 encapsulated traffic.

Platforms

7705 SAR Gen 2

9.16 egress

egress

Syntax

egress

Context

[Tree] (config>service>vprn>if>spoke-sdp egress)

Full Context

configure service vprn interface spoke-sdp egress

Description

This command configures egress SDP parameters.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>service>ies>if>sap egress)

[Tree] (config>service>vpls>sap egress)

Full Context

configure service ies interface sap egress

configure service vpls sap egress

Description

Commands in this context configure egress Quality of Service (QoS) policies and filter policies.

If no QoS policy is defined, the system default QoS policy is used for egress processing. If no egress filter is defined, no filtering is performed.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>service>vpls>sap egress)

Full Context

configure service vpls sap egress

Description

Commands in this context configure egress filter policies.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. If **no** egress filter is defined, no filtering is performed.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>service>vpls>mesh-sdp egress)

[Tree] (config>service>ies>if>spoke-sdp egress)

[Tree] (config>service>vpls>spoke-sdp egress)

Full Context

configure service vpls mesh-sdp egress

configure service ies interface spoke-sdp egress

configure service vpls spoke-sdp egress

Description

Commands in this context configure egress SDP parameters.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>port>ethernet>access egress)

[Tree] (config>port>ethernet>network egress)

Full Context

configure port ethernet access egress

configure port ethernet network egress

Description

This command configures Ethernet access egress port parameters.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>port>network egress)

[Tree] (config>card>mda>network egress)

[Tree] (config>port>access egress)

[Tree] (config>card>mda>access egress)

Full Context

configure port network egress

configure card mda network egress

configure port access egress

configure card mda access egress

Description

Commands in this context configure egress buffer pool parameters which define the percentage of the pool buffers that are used for CBS calculations and specify the slope policy that is configured in the **config>qos>slope-policy** context.

On the MDA level, network and access egress pools are only allocated on channelized MDAs.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>port>ethernet egress)

Full Context

configure port ethernet egress

Description

This command configures Ethernet egress port parameters.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>service>epipe>sap egress)

Full Context

configure service epipe sap egress

Description

Commands in this context configure egress SAP parameters.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>service>epipe>spoke-sdp egress)

Full Context

configure service epipe spoke-sdp egress

Description

This command configures the egress SDP context.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>service>ies>if>vpls egress)

Full Context

configure service ies interface vpls egress

Description

The egress node under the vpls binding is used to define the optional sap-egress QoS policy that will be used for reclassifying the egress forwarding class or profile for routed packets associated with the IP interface on the attached VPLS or I-VPLS service context.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>service>vprn>nw-if egress)

Full Context

configure service vprn network-interface egress

Description

Commands in this context configure egress network filter policies for the interface.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>service>vprn>if>sap egress)

Full Context

configure service vprn interface sap egress

Description

Commands in this context configure egress SAP Quality of Service (QoS) policies and filter policies.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. If no egress filter is defined, no filtering is performed.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>service>vprn>if>vpls egress)

Full Context

configure service vprn interface vpls egress

Description

The egress node under the vpls binding is used to define the optional sap-egress QoS policy that will be used for reclassifying the egress forwarding class or profile for routed packets associated with the IP interface on the attached VPLS service context.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>service>vprn>network-interface egress)

Full Context

configure service vprn network-interface egress

Description

Commands in this context configure egress network filter policies for the interface.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>mirror>mirror-dest>remote-src>spoke-sdp egress)

[Tree] (config>mirror>mirror-dest>spoke-sdp egress)

Full Context

configure mirror mirror-dest remote-source spoke-sdp egress

configure mirror mirror-dest spoke-sdp egress

Description

Commands in this context configure spoke SDP egress parameters.

Platforms

7705 SAR Gen 2

egress

Syntax

[**no**] **egress**

Context

[Tree] (config>mirror>mirror-dest>sap egress)

Full Context

configure mirror mirror-dest sap egress

Description

This command enables access to the context to associate an egress SAP Quality of Service (QoS) policy with a mirror destination SAP.

If no QoS policy is defined, the system default SAP egress QoS policy is used for egress processing.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>qos>network egress)

Full Context

configure qos network egress

Description

This command is used to enter the CLI node that creates or edits egress policy entries that specify the forwarding class queues to be instantiated when this policy is applied to the network port.

The forwarding class and profile state mapping to in- and out-of-profile DiffServ Code Points (DSCPs), dot1p, and MPLS EXP bits mapping for all labeled packets are also defined in this context.

All service packets are aggregated into DiffServ-based egress queues on the network interface. The service packets are transported either with IP GRE encapsulation or over a MPLS LSP. The exception is with the IES service. In this case, the actual customer IP header has the DSCP field mapped.

All out-of-profile service packets are marked with the corresponding out-of-profile DSCP, dot1p, or the EXP bit value at network egress. All the in-profile service ingress packets are marked with the corresponding in-profile DSCP, dot1p, or EXP bit value based on the forwarding class to which they belong. The exceed-profile traffic is marked with the same value as out-of-profile traffic and the inplus-profile traffic is marked with the same value as in-profile traffic.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>qos>queue-group-templates egress)

Full Context

configure qos queue-group-templates egress

Description

Commands in this context configure QoS egress queue groups. Egress queue group templates can be applied to egress Ethernet ports to create an egress queue group.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>router>if egress)

Full Context

configure router interface egress

Description

This command enables access to the context to configure egress network filter policies for the IP interface. If an egress filter is not defined, no filtering is performed.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>service>cust>multi-service-site egress)

Full Context

configure service customer multi-service-site egress

Description

Commands in this context configure the egress node associate an existing scheduler policy name with the customer site. The egress node is an entity to associate commands that complement the association.

Platforms

7705 SAR Gen 2

egress

Syntax

egress

Context

[Tree] (config>service>pw-template egress)

Full Context

configure service pw-template egress

Description

Commands in this context configure spoke SDP binding egress filter parameters.

Platforms

7705 SAR Gen 2

9.17 egress-amplifier-gain

egress-amplifier-gain

Syntax

egress-amplifier-gain *egress-amplifier-gain*

no egress-amplifier-gain

Context

[Tree] (configure>port>transceiver>optical-line-system egress-amplifier-gain)

Full Context

configure port transceiver optical-line-system egress-amplifier-gain

Description

This command configures the gain for the egress amplifier.

The **no** form of this command sets the gain for the egress amplifier to the default.

Default

no egress-amplifier-gain

Parameters

egress-amplifier-gain

Specifies the gain for the amplifier in decibels.

Values 0 to 25.00 dB

Default 25.00 dB

Platforms

7705 SAR Gen 2

9.18 egress-engineering

egress-engineering

Syntax

egress-engineering

no egress-engineering

Context

[Tree] (config>router>bgp>group>neighbor egress-engineering)

[Tree] (config>router>bgp>group egress-engineering)

Full Context

configure router bgp group neighbor egress-engineering

configure router bgp group egress-engineering

Description

Commands in this context configure egress engineering on a specific neighbor or all neighbors in a BGP group.

If egress engineering is not configured in the neighbor context, the configuration is inherited from the group context.

The **no** form of this command removes the egress engineering configuration.

Default

no egress-engineering

Platforms

7705 SAR Gen 2

9.19 egress-fc

egress-fc

Syntax

egress-fc *fc-name*

no egress-fc

Context

[Tree] (config>qos>sap-ingress>fc egress-fc)

Full Context

configure qos sap-ingress fc egress-fc

Description

This command configures the forwarding class to be used by the egress QoS processing. It overrides the forwarding class determined by ingress classification but not the QoS Policy Propagation via BGP.

The forwarding class or forwarding subclass can be overridden.

The new egress forwarding class is applicable to both SAP egress and network egress.

Default

no egress-fc

Parameters

fc-name

Specifies the forwarding class name to be used by the egress QoS processing.

Values be, l2, af, l1, h2, ef, h1, nc

Platforms

7705 SAR Gen 2

9.20 egress-peer-engineering

egress-peer-engineering

Syntax

egress-peer-engineering

no egress-peer-engineering

Context

[Tree] (config>router>bgp egress-peer-engineering)

Full Context

configure router bgp egress-peer-engineering

Description

Commands in this context configure EPE parameters in BGP.

The **no** form of this command removes the EPE parameters from the BGP context.

Default

no egress-peer-engineering

Platforms

7705 SAR Gen 2

9.21 egress-peer-engineering-label-unicast

egress-peer-engineering-label-unicast

Syntax

[**no**] **egress-peer-engineering-label-unicast**

Context

[Tree] (config>router>bgp>group>neighbor egress-peer-engineering-label-unicast)

[Tree] (config>router>bgp>group egress-peer-engineering-label-unicast)

Full Context

configure router bgp group neighbor egress-peer-engineering-label-unicast

configure router bgp group egress-peer-engineering-label-unicast

Description

This command enables the generation of a label-unicast route for each /32 or /128 prefix that corresponds to the BGP neighbor or group address in the scope of the command. These routes can be advertised to other routers to recursively resolve unlabeled BGP routes for AS external destinations. They support the Egress Peer Engineering (EPE) use case.

The **no** form of this command disables the generation of EPE label-unicast routes.

Default

no egress-peer-engineering-label-unicast

Platforms

7705 SAR Gen 2

9.22 egress-rate

egress-rate

Syntax

egress-rate *sub-rate*

no egress-rate

Context

[Tree] (config>port>ethernet egress-rate)

Full Context

configure port ethernet egress-rate

Description

This command configures the rate of traffic leaving the network. The configured *sub-rate* uses packet-based accounting. An event log is generated each time the egress rate is modified unless the port is part of a LAG.

The **no** form of this command returns the value to the default.

Default

no egress-rate

Parameters

sub-rate

Specifies the egress rate in kb/s.

Values 1 to 100000000

Platforms

7705 SAR Gen 2

9.23 egress-scheduler-override

egress-scheduler-override

Syntax

egress-scheduler-override [**create**]

no egress-scheduler-override

Context

[Tree] (config>port>ethernet egress-scheduler-override)

Full Context

configure port ethernet egress-scheduler-override

Description

This command applies egress scheduler overrides. When a port scheduler is associated with an egress port, it is possible to override the following parameters:

- The **max-rate** allowed for the scheduler.
- The maximum **rate** for each priority level 8 through 1.
- The CIR associated with each priority level 8 through 1.

See the *7705 SAR Gen 2 Quality of Service Guide* for command syntax and usage for the **port-scheduler-policy** command.

The **no** form of this command removes all override parameters from the egress port or channel scheduler context. Once removed, the port scheduler reverts all rate parameters back to the parameters defined on the port-scheduler-policy associated with the port.

Parameters

create

Mandatory while creating an entry.

Platforms

7705 SAR Gen 2

9.24 egress-scheduler-policy

egress-scheduler-policy

Syntax

egress-scheduler-policy *port-scheduler-policy-name*

no egress-scheduler-policy

Context

[Tree] (config>port>ethernet egress-scheduler-policy)

Full Context

configure port ethernet egress-scheduler-policy

Description

This command enables the provisioning of an existing port-scheduler-policy to a port or channel.

The egress-scheduler-override node allows for the definition of the scheduler overrides for a specific port or channel.

When a port scheduler is active on a port or channel, all queues and intermediate service schedulers on the port are subject to receiving bandwidth from the scheduler. Any policers, queues, or schedulers with port-parent associations are mapped to the appropriate port priority levels based on the port-parent command parameters. Any policers, queues, or schedulers that do not have a port-parent or valid intermediate scheduler parent defined are treated as orphaned and are handled based on the port scheduler policies default or explicit orphan behavior.

The port scheduler maximum rate and priority level rate parameters may be overridden to allow unique values separate from the port-scheduler-policy-name attached to the port or channel. Use the **egress-scheduler-override** command to specify the port or channel specific scheduling parameters.

The **no** form of this command removes a port scheduler policy from an egress port or channel. Once the scheduler policy is removed, all orphaned policers, queues, and schedulers revert to a free running state governed only by the local queue or scheduler parameters. This includes any queues or schedulers with a port-parent association.

Parameters

port-scheduler-policy-name

Specifies an existing port-scheduler-policy configured in the **config>qos** context. The name can be up to 32 characters.

Platforms

7705 SAR Gen 2

9.25 egress-statistics

egress-statistics

Syntax

[**no**] **egress-statistics**

Context

[Tree] (config>router>mpls>fwd-policies>fwd-policy egress-statistics)

Full Context

configure router mpls forwarding-policies forwarding-policy egress-statistics

Description

This command configures egress statistics in an MPLS forwarding policy.

The **no** form of this command removes any egress statistics in a forwarding policy.

Default

no egress-statistics

Platforms

7705 SAR Gen 2

9.26 eibgp-loadbalance

eibgp-loadbalance

Syntax

[**no**] **eibgp-loadbalance**

Context

[Tree] (config>service>vprn>bgp eibgp-loadbalance)

Full Context

configure service vprn bgp eibgp-loadbalance

Description

This command enables eiBGP load sharing so routes with both MP-BGP and IPv4 next-hops can be used simultaneously.

In order for this command to be effective, the **ecmp** and **multipath** commands for the associated VPRN instance must also be configured to allow for multiple routes to the same destination.

The **no** form of this command used at the global level reverts to default values.

Default

no eibgp-loadbalance

Platforms

7705 SAR Gen 2

9.27 embed-filter

embed-filter

Syntax

embed-filter *ip-filter-id* [**offset** *offset*] [{**active** | **inactive**}]

no embed-filter *ip-filter-id*

embed-filter *ipv6-filter-id* [**offset** *offset*] [{**active** | **inactive**}]

no embed-filter *ipv6-filter-id*

embed-filter flowspec [**group** *group-id*] [**router** {*router-instance* | **service-name** *vprn-service-name*}] [**offset** *offset*] [{**active** | **inactive**}]

no embed-filter flowspec [**group** *group-id*]

embed-filter open-flow *ofs-name* [{**system** | **service** {*service-id* | *service-name*} | **sap** *sap-id*}] [**offset** *offset*] [{**active** | **inactive**}]

no embed-filter open-flow *ofs-name* [{**system** | **service** {*service-id* | *service-name*} | **sap** *sap-id*}]

Context

[Tree] (config>filter>ipv6-filter embed-filter)

[Tree] (config>filter>ip-filter embed-filter)

Full Context

configure filter ipv6-filter embed-filter

configure filter ip-filter embed-filter

Description

This command embeds a previously defined IPv4, IPv6, or MAC embedded filter policy or Hybrid OpenFlow switch instance into this exclusive, template, or system filter policy at the specified offset value. Rules derived from the BGP FlowSpec can also be embedded into template filter policies only.

The **embed-filter open-flow** *ofs-name* form of this command enables OpenFlow (OF) in GRT either by embedding the specified OpenFlow switch (OFS) instance with **switch-defined-cookie** disabled, or by

embedding rules with sros-cookie:type "grt-cookie”, value 0, from the specified OFS instance with **switch-defined-cookie** enabled. The embedding filter can only be deployed in GRT context or be unassigned.

The **embed-filter open-flow** *ofs-name* **system** form of this command enables OF in system filters by embedding rules with sros-cookie:type "system-cookie”, value 0, from the specified OFS instance with **switch-defined-cookie** enabled. The embedding filter can only be of scope **system**.

The **embed-filter open-flow** *ofs-name* **service** {*service-id* | *service-name*} form of this command enables OF in VPRN/VPLS filters by embedding rules with sros-cookie:type "service-cookie”, value **service-id**, from the specified OFS instance with **switch-defined-cookie** enabled—per service rules. The embedding filter can only be deployed in the specified VPRN/VPLS service. A single VPLS service can only support OF rules per SAP or per service.

The **embed-filter open-flow** *ofs-name* **sap** *sap-id* form of this command enables OF in VPLS SAP filters by embedding rules with sros-cookie:type "service-cookie”, value *service-id* and flow match conditions specifying the sap-id from the specified OFS instance with **switch-defined-cookie** enabled—per SAP OF rules. The embedding filter must be of type exclusive and can only be deployed on the specified SAP in the context of the specified VPLS service. A single VPLS service can only support OF rules per SAP or per service.

The **no embed-filter open-flow** *ofs-name* form of this command removes the OF embedding for the GRT context.

The **embed-filter flowspec** form of this command enables the embedding of rules derived from BGP FlowSpec routes into the filter policy that is being configured. The optional **group** parameter specifies that only FlowSpec routes tagged with an interface-set extended community containing this group ID should be selected for embedding. The optional **router** parameter specifies the routing instance source of the BGP FlowSpec routes; if the parameter is not specified, the routing instance is derived automatically from the context in which the filter policy is applied.

The **no embed-filter flowspec** form of this command removes the FlowSpec filter embedding from this filter policy.

The **no embed-filter** *filter-id* form of this command removes the embedding from this filter policy.

See the description of embedded filter policies in this guide for further operational details.

Parameters

ip-filter-id

Specifies a previously defined IPv4 policy for embedding in this filter.

ipv6-filter-id

Specifies a previously defined IPv6 policy for embedding in this filter.

offset

Specifies that an embedded filter entry X will have an entry X + offset in the embedding filter.

Values 0 to 2097151

Default 0

active

Specifies that embedded filter entries are to be included in this embedding filter policy and activated on applicable line cards—default if no keyword is specified and omitted from **info** command output (but not **info detail**), or when saving the configuration.

inactive

Specifies that no embedded filter policy entries are to be included in this embedding filter policy. The embedding is configured but will not do anything.

flowspec

This keyword indicates that rules derived from BGP FlowSpec routes should be embedded into (or removed from, in case of the **no** form) the filter.

group-id

Specifies that only FlowSpec routes with an interface-set extended community with this value of *group-id* should be selected for embedding.

Values 0 to 16383

router-instance

Specifies a router instance.

vprn-service-name

Specifies the VPRN service name used for embedding FlowSpec rules.

open-flow

Indicates that rules derived from OpenFlow should be embedded into (or removed from, in case of the **no** form) the filter.

ofs-name

Specifies the name of the currently configured Hybrid OpenFlow Switch (OFS) instance.

Not including the **system**, **service** or **sap** parameters will specify OF in a GRT instance context by default. This allows embedding of OF rules into filters deployed in GRT instances from OFS with **switch-defined-cookie** disabled, or embedding rules from OFS with **switch-defined-cookie** enabled, when the FlowTable cookie encodes sros-cookie:type "grt-cookie”.

system

Used for OF control of system filters. Allows embedding of OF rules into system filters from OFS with **switch-defined-cookie** enabled. Only the rules with cookie value encoding "system-cookie” are embedded.

service-id

Specifies an existing VPRN or VPLS service ID that the embedding filter can be used for.

service-name — Specifies an existing VPRN or VPLS service name that the embedding filter can be used for.

Values 1 to 2147483647

service-name

Specifies an existing VPRN or VPLS service name up to 64 characters that the embedding filter can be used for.

sap-id

Used for OF control of VPLS services when a PortID and VLAN ID match is required. Allows embedding of OF rules with a PortID and VLAN ID match into exclusive VPLS SAP filters. Only the rules with cookie value encoding the VPLS service, and flow table match

encoding the specified SAP, are embedded into the filter. The embedding filter can only be deployed in the context of the specified SAP.

sap-id — Specifies an existing SAP that the embedding filter can be used for.

Platforms

7705 SAR Gen 2

9.28 embedded-rp

embedded-rp

Syntax

embedded-rp

Context

[Tree] (config>service>vprn>pim>rp>ipv6 embedded-rp)

Full Context

configure service vprn pim rp ipv6 embedded-rp

Description

This command enables context to configure IPv6 embedded RP parameters.

Platforms

7705 SAR Gen 2

embedded-rp

Syntax

[**no**] **embedded-rp**

Context

[Tree] (config>router>pim>rp>ipv6 embedded-rp)

Full Context

configure router pim rp ipv6 embedded-rp

Description

Commands in this context configure embedded RP parameters.

Embedded RP is required to support IPv6 inter-domain multicast because there is no MSDP equivalent in IPv6.

The detailed protocol specification is defined in RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*. This RFC describes a multicast address allocation policy in which the address of the RP is encoded in the IPv6 multicast group address, and specifies a PIM-SM group-to-RP mapping to use the encoding, leveraging, and extending unicast-prefix-based addressing. This mechanism not only provides a simple solution for IPv6 inter-domain ASM but can be used as a simple solution for IPv6 intra-domain ASM with scoped multicast addresses as well. It can also be used as an automatic RP discovery mechanism in those deployment scenarios that would have previously used the Bootstrap Router protocol (BSR).

The **no** form of this command disables embedded RP.

Platforms

7705 SAR Gen 2

9.29 emulated-server

emulated-server

Syntax

emulated-server *ip-address*

no emulated-server

Context

[Tree] (config>service>vprn>if>dhcp>proxy-server emulated-server)

[Tree] (config>service>ies>if>dhcp>proxy-server emulated-server)

[Tree] (config>service>vpls>sap>dhcp>proxy-server emulated-server)

Full Context

configure service vprn interface dhcp proxy-server emulated-server

configure service ies interface dhcp proxy-server emulated-server

configure service vpls sap dhcp proxy-server emulated-server

Description

This command configures the IP address which is used as the DHCP server address in the context of the SAP. Typically, the configured address should be in the context of the subnet represented by the service.

The **no** form of this command reverts to the default setting. The local proxy server will not become operational without the emulated-server address being specified.

The **no** form of this command reverts to the default.

Parameters

ip-address

Specifies the emulated server's IP address. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

Platforms

7705 SAR Gen 2

9.30 enable-admin

enable-admin

Syntax

enable-admin

Context

[Tree] (enable-admin)

Full Context

enable-admin

Description

See the description for the **admin-password** command. If the **admin-password** is configured in the **config>system>security>password** context, then any user can enter a special administrative mode by entering the **enable-admin** command.

enable-admin is in the default profile. By default, all users are given access to this command.

Once the **enable-admin** command is entered, the user is prompted for a password. If the password matches, the user is given unrestricted access to all the commands.

The minimum length of the password is determined by the **minimum-length** command. The complexity requirements for the password are determined by the **complexity** command.

The following shows a password configuration example:

```
A:ALA-1>config>system>security# info
----------------------------------------------
...
            password
                aging 365
                minimum-length 8
                attempts 5 time 5 lockout 20
                admin-password "rUYUz9XMo6I" hash
            exit
...
----------------------------------------------
A:ALA-1>config>system>security#
```

There are two ways to verify that a user is in the enable-admin mode:

- show users — administrator can know which users are in this mode
- Enter the **enable-admin** command again at the root prompt and an error message will be returned.

```
*A:node-1# show users
===============================================================================
User                            Type      Login time           Idle time
  Session ID   From
===============================================================================
                                Console         --             3d 10:16:12 --
  6            --
admin                           SSHv2     12OCT2018 20:44:15   0d 00:00:00 A-
 #83           192.168.0.10
admin                           SSHv2     12OCT2018 21:09:25   0d 00:05:10 --
  84           192.168.0.10
-------------------------------------------------------------------------------
Number of users: 2
'#' indicates the current active session
'A' indicates user is in admin mode
===============================================================================
*A:node-1# enable-admin
MINOR: CLI Already in admin mode.
*A:node-1#
```

Platforms

7705 SAR Gen 2

9.31 enable-admin-control

enable-admin-control

Syntax

enable-admin-control

Context

[Tree] (config>system>security>password enable-admin-control)

Full Context

configure system security password enable-admin-control

Description

Enable the user to become a system administrator.

Note:
This command applies to users on RADIUS, TACACS, and LDAP.

Platforms
7705 SAR Gen 2

9.32 enable-graceful-shutdown

enable-graceful-shutdown

Syntax
[**no**] **enable-graceful-shutdown**

Context
[Tree] (config>system>login-control>telnet enable-graceful-shutdown)

Full Context
configure system login-control telnet enable-graceful-shutdown

Description
This command enables graceful shutdown of telnet sessions.

The **no** form of this command disables graceful shutdown of telnet sessions.

Platforms
7705 SAR Gen 2

9.33 enable-grt

enable-grt

Syntax
[**no**] **enable-grt**

Context
[Tree] (config>service>vprn>grt-lookup enable-grt)

Full Context
configure service vprn grt-lookup enable-grt

Description
This command enables the functions required for looking up routes in the Global Route Table (GRT) when the lookup in the local VRF fails. If this command is enabled without the use of a **static-route** option (as

subcommand to this parent), a lookup in the local VRF is preferred over the GRT. When the local VRF returns no route table lookup matches, the result from the GRT is preferred.

The **no** form of this command disables the lookup in the GRT when the lookup in the local VRF fails.

Default

no enable-grt

Platforms

7705 SAR Gen 2

9.34 enable-icmp-vse

enable-icmp-vse

Syntax

[**no**] **enable-icmp-vse**

Context

[Tree] (config>system enable-icmp-vse)

Full Context

configure system enable-icmp-vse

Description

This command enables vendor specific extensions to ICMP.

Default

no enable-icmp-vse

Platforms

7705 SAR Gen 2

9.35 enable-inter-as-vpn

enable-inter-as-vpn

Syntax

[**no**] **enable-inter-as-vpn**

Context

[Tree] (config>router>bgp enable-inter-as-vpn)

Full Context

configure router bgp enable-inter-as-vpn

Description

This command specifies whether VPNs can exchange routes across autonomous system boundaries, providing model B connectivity.

The **no** form of this command disallows ASBRs to advertise VPRN routes to their peers in other autonomous systems.

Default

no enable-inter-as-vpn

Platforms

7705 SAR Gen 2

9.36 enable-mac-accounting

enable-mac-accounting

Syntax

[**no**] **enable-mac-accounting**

Context

[Tree] (config>service>ies>if enable-mac-accounting)

Full Context

configure service ies interface enable-mac-accounting

Description

This command enables MAC accounting functionality on this interface.

The **no** form of this command disables MAC accounting functionality on this interface.

Platforms

7705 SAR Gen 2

enable-mac-accounting

Syntax

[**no**] **enable-mac-accounting**

Context

[Tree] (config>service>vprn>if enable-mac-accounting)

Full Context

configure service vprn interface enable-mac-accounting

Description

This command enables MAC accounting functionality on this interface.

The **no** form of this command disables MAC accounting functionality on this interface.

Platforms

7705 SAR Gen 2

enable-mac-accounting

Syntax

[**no**] **enable-mac-accounting**

Context

[Tree] (config>router>if enable-mac-accounting)

Full Context

configure router interface enable-mac-accounting

Description

This command enables MAC Accounting functionality for the interface.

Default

no enable-mac-accounting

Platforms

7705 SAR Gen 2

9.37 enable-mdt-spt

enable-mdt-spt

Syntax

[**no**] **enable-mdt-spt**

Context

[Tree] (config>router>pim enable-mdt-spt)

Full Context

configure router pim enable-mdt-spt

Description

This command enables SPT switchover for default MDT. On enable, PIM instance resets all MDTs and re-initiate setup.

The **no** form of this command disables SPT switchover for default MDT. On disable, PIM instance resets all MDTs and re-initiate setup.

Default

no enable-mdt-spt

Platforms

7705 SAR Gen 2

9.38 enable-notification

enable-notification

Syntax

enable-notification

no enable-notification

Context

[Tree] (config>service>vprn>bgp>group>neighbor>graceful-restart enable-notification)

[Tree] (config>service>vprn>bgp>graceful-restart enable-notification)

[Tree] (config>service>vprn>bgp>group>graceful-restart enable-notification)

Full Context

configure service vprn bgp group neighbor graceful-restart enable-notification

configure service vprn bgp graceful-restart enable-notification

configure service vprn bgp group graceful-restart enable-notification

Description

When this command is present, the graceful restart capability sent by this router indicates support for NOTIFICATION messages. If the peer also supports this capability then the session can be restarted gracefully (while preserving forwarding) if either peer sends a NOTIFICATION message due to some type of event or error.

Default

no enable-notification

Platforms

7705 SAR Gen 2

enable-notification

Syntax

enable-notification

no enable-notification

Context

[Tree] (config>router>bgp>group>neighbor>graceful-restart enable-notification)

[Tree] (config>router>bgp>group>graceful-restart enable-notification)

[Tree] (config>router>bgp>graceful-restart enable-notification)

Full Context

configure router bgp group neighbor graceful-restart enable-notification

configure router bgp group graceful-restart enable-notification

configure router bgp graceful-restart enable-notification

Description

When this command is present, the graceful restart capability sent by this router indicates support for NOTIFICATION messages. If the peer also supports this capability, then the session can be restarted gracefully (while preserving forwarding) if either peer needs to send a NOTIFICATION message due to some type of event or error.

Default

no enable-notification

Platforms

7705 SAR Gen 2

9.39 enable-origin-validation

enable-origin-validation

Syntax

enable-origin-validation [**ipv4**] [**ipv6**] [**label-ipv4**]

no enable-origin-validation

Context

[Tree] (config>service>vprn>bgp>group enable-origin-validation)

[Tree] (config>service>vprn>bgp>group>neighbor enable-origin-validation)

Full Context

configure service vprn bgp group enable-origin-validation

configure service vprn bgp group neighbor enable-origin-validation

Description

When this command is added to the configuration of a group or neighbor, it causes every inbound IPv4, IPv6, and label-IPv4 route from that peer to be marked with one of the following origin validation states:

- Valid (0)
- Not-Found (1)
- Invalid (2)

By default (when no family parameter is present in the command) or when all the family options are specified, all unicast IPv4 (AFI1/SAFI1), label-IPv4 (AFI1/SAFI4), and unicast IPv6 (AFI2/SAFI1) routes are evaluated to determine their origin validation states. When only a subset of the family options are present, then only the corresponding address family routes are evaluated.

This command applies to all types of VPRN BGP peers, generally, it should only be applied to EBGP peers and groups that contain only EBGP peers.

The **no** form of this command disables the inspection of received routes from the peer to determine origin validation state.

Default

no enable-origin-validation

Parameters

ipv4

Enables origin validation processing for unlabeled unicast IPv4 routes.

ipv6

Enables origin validation processing for unlabeled unicast IPv6 routes.

label-ipv4

Enables origin validation processing for labeled IPv4 routes.

Platforms

7705 SAR Gen 2

enable-origin-validation

Syntax

enable-origin-validation [**ipv4**] [**ipv6**] [**label-ipv4**] [**label-ipv6**]

no enable-origin-validation

Context

[Tree] (config>router>bgp>group>neighbor enable-origin-validation)

[Tree] (config>router>bgp>group enable-origin-validation)

Full Context

configure router bgp group neighbor enable-origin-validation

configure router bgp group enable-origin-validation

Description

When the **enable-origin-validation** command is added to the configuration of a group or neighbor, it causes every inbound IPv4 or IPv6 route from that peer to be marked with one of the following origin validation states:

- Valid (0)
- Not-Found (1)
- Invalid (2)

By default (when neither the ipv4 or ipv6 option is present in the command) or when both the ipv4 and ipv6 options are specified, all unicast IPv4 (AFI1/SAFI1), label-IPv4 (AFI1/SAFI4), unicast IPv6 (AFI2/SAFI1), and label-IPv6 (AFI2/SAFI4) routes are evaluated to determine their origin validation states. When only the ipv4 or ipv6 option is present, only the corresponding address family routes (unlabeled and labeled) are evaluated.

The **enable-origin-validation** command applies to all types of BGP peers, but as a general rule, it should only be applied to EBGP peers and groups that contain only EBGP peers.

Default

no enable-origin-validation

Parameters

ipv4

Enables origin validation processing for unlabeled unicast IPv4 routes.

ipv6

Enables origin validation processing for unlabeled unicast IPv6 routes.

label-ipv4

Enables origin validation processing for labeled IPv4 routes.

label-ipv6

Enables origin validation processing for labeled IPv6 routes.

Platforms

7705 SAR Gen 2

9.40 enable-peer-tracking

enable-peer-tracking

Syntax

[**no**] **enable-peer-tracking**

Context

[Tree] (config>service>vprn>bgp enable-peer-tracking)

[Tree] (config>service>vprn>bgp>group>neighbor enable-peer-tracking)

[Tree] (config>service>vprn>bgp>group enable-peer-tracking)

Full Context

configure service vprn bgp enable-peer-tracking

configure service vprn bgp group neighbor enable-peer-tracking

configure service vprn bgp group enable-peer-tracking

Description

This command enables BGP peer tracking.

Default

no enable-peer-tracking

Platforms

7705 SAR Gen 2

enable-peer-tracking

Syntax

[**no**] **enable-peer-tracking**

Context

[Tree] (config>router>bgp>group enable-peer-tracking)

[Tree] (config>router>bgp>group>neighbor enable-peer-tracking)

[Tree] (config>router>bgp enable-peer-tracking)

Full Context

configure router bgp group enable-peer-tracking

configure router bgp group neighbor enable-peer-tracking

configure router bgp enable-peer-tracking

Description

This command enables BGP peer tracking. BGP peer tracking allows a BGP peer to be dropped immediately if the route used to resolve the BGP peer address is removed from the IP routing table and there is no alternative available. The BGP peer will not wait for the holdtimer to expire; therefore, the BGP re-convergence process is accelerated.

The **no** form of this command disables peer tracking.

Default

no enable-peer-tracking

Platforms

7705 SAR Gen 2

9.41 enable-rr-vpn-forwarding

enable-rr-vpn-forwarding

Syntax

[**no**] **enable-rr-vpn-forwarding**

Context

[Tree] (config>router>bgp enable-rr-vpn-forwarding)

Full Context

configure router bgp enable-rr-vpn-forwarding

Description

When this command is configured all received VPN-IP routes, regardless of route target, are imported into the dummy VRF, where the BGP next-hops are resolved. The **label-route-transport-tunnel** under **config>router>bgp>next-hop-resolution** determines what types of tunnels are eligible to resolve the next-hops. If a received VPN-IP route from IBGP peer X is resolved and selected as best so that it can be re-advertised to an IBGP peer Y, and the BGP next-hop is modified towards peer Y (by using the next-hop-self command in Y's group or neighbor context or by using a next-hop action in an export policy applied to Y) then BGP allocates a new VPRN service label value for the route, signals that new label value to Y and programs the IOM to do the corresponding label swap operation. The supported combinations of X and Y are outlined below:

- from X (client) to Y (client)
- from X (client) to Y (non-client)
- from X (non-client) to Y (client)

The **no** form of this command causes the re-advertisement of a VPN-IP route between one IBGP peer and another IBGP peer does not cause a new VPRN service label value to be signaled and programmed even if the BGP next-hop is changed through group/neighbor configuration or policy.

Nokia recommends leaving this command disabled for scaling and convergence reasons.

Default

no enable-rr-vpn-forwarding

Platforms

7705 SAR Gen 2

9.42 enable-subconfed-vpn-forwarding

enable-subconfed-vpn-forwarding

Syntax

[**no**] **enable-subconfed-vpn-forwarding**

Context

[Tree] (config>router>bgp enable-subconfed-vpn-forwarding)

Full Context

configure router bgp enable-subconfed-vpn-forwarding

Description

This command configures BGP to keep VPN-IPv4 and VPN-IPv6 routes within a subconfederation and allow a **next-hop-self** command to create label swap forwarding entries.

When this is enabled, the base router BGP instance retains all received VPN-IPv4 and VPN-IPv6 routes, even those with route targets not matching any VRF import policy of any locally configured VPRN. In addition, when this leaf is enabled and base router BGP is configured to apply a **next-hop-self** command to a peer of any type (EBGP, IBGP, or confed-EBGP), the VPN-IPv4 and VPN-IPv6 routes are advertised to the peer with a new BGP label and next-hop, and a label-swap forwarding entry is programmed.The preceding behaviors are applied when the **enable-inter-as-vpn** or the **enable-rr-vpn-forwarding** commands, both under the **configure router bgp** context, are also enabled in the same BGP instance and regardless of whether the base router has a confederation configuration.

The **no** form of this command disables subconfederation VPN forwarding.

Default

no enable-subconfed-vpn-forwarding

Platforms

7705 SAR Gen 2

9.43 enable-tech

enable-tech

Syntax

[**no**] **enable-tech**

Context

[Tree] (admin enable-tech)

Full Context

admin enable-tech

Description

This command enables the shell and kernel commands.

Note:
This command should only be used with authorized direction of Nokia support.

Platforms

7705 SAR Gen 2

9.44 encap-match

encap-match

Syntax

encap-match {**all-encap** | **double-tag** *encap-value* | **single-tag** *encap-value* | **untagged**}

no encap-match

Context

[Tree] (config>port>ethernet>dot1x>macsec>sub-port encap-match)

Full Context

configure port ethernet dot1x macsec sub-port encap-match

Description

This command defines the sub-set of traffic on this port affected by this MACsec sub-port.

In order to establish an end-to-end communication between the remote MACsec peers encrypting VLAN-tagged traffic, the MKA packets have to be able to travel over the network following the same path as the encrypted traffic. MKA packets are generated with specific tags depending on the traffic match criteria configured, as shown in Table 31: MKA Packet Generation .

The **no** form of this command removes all traffic sub-set definitions from the MACsec sub-port.

Table 31: MKA Packet Generation

Configuration	Config Example (<s-tag>.<c-tag>)	MKA Packet Generation	Traffic pattern match/behavior
PORT all-encap	Config>port>ethernet>dot1x>macsec Sub-port 10 encap-match all-encap ca-name 10	untagged MKA packet	Matches all traffic on the port, including untagged, single-tag, double-tag. This is the Release 15.0 default behavior.
Untagged	Config>port>ethernet>dot1x>macsec Sub-port 1 encap-match untagged ca-name 2	untagged MKA packet	Matches only untagged traffic on the port
802.1Q single S-TAG (specific S-TAG)	Config>port>ethernet>dot1x>macsec Sub-port 2 encap-match dot1q 1	MKA packet generated with S-TAG=1	Matches only single-tag traffic on port with tag ID of 1

Configuration	Config Example (<s-tag>.<c-tag>)	MKA Packet Generation	Traffic pattern match/behavior
	ca-name 3		
802.1Q single S-TAG (any S-TAG)	Config>port>ethernet>dot1x>macsec Sub-port 3 encap-match dot1q * ca-name 4	untagged MKA packet	Matches any single-tag traffic on port
802.1ad double tag (both tag have specific TAGs)	Config>port>ethernet>dot1x>macsec Sub-port 4 encap-match qinq 1.1 ca-name 5	MKA packet generated with S-tag=1 and C-TAG=1	Matches only double-tag traffic on port with service tag of 1 and customer tag of 1
802.1ad double tag (specific S-TAG, any C-TAG)	Config>port>ethernet>dot1x>macsec Sub-port 6 encap-match qinq 1.* ca-name 7	MKA packet generated with S-TAG=1	Matches only double-tag traffic on port with service tag of 1 and customer tag of any
802.1ad double tag (any S-TAG, any C-TAG	Config>port>ethernet>dot1x>macsec Sub-port 7 encap-match double-tag *.* ca-name 8	untagged MKA packet	Matches any double-tag traffic on port

Default

encap-match all-encap

Parameters

all-encap

Specifies that all traffic patterns are matched including untagged, single-tag or double-tag, and all will be encrypted.

untagged

Specifies that only untagged traffic are matched and encrypted.

single-tag

Specifies that only dot1q traffic are matched. Either all single tag traffic can be matched, by using *, or a specific dot1q tag can be matched.

double-tag

Specifies that only qinq traffic are matched. The service tag can be specifically matched or a wild card match (*.*) can be used.

encap-value

Specifies the type and value of the packet encapsulation to match for this MACsec sub-port.

Type	Parameter
all-encap	—
untagged	—
dot1q	[*| s] (s = 0..4094)
qinq	[*.*| s.*| s.c] (s and c = 0..4094)

where:

- S = service tag
- C = customer tag

Platforms

7705 SAR Gen 2

9.45 encap-type

encap-type

Syntax

encap-type {**dot1q** | **null** | **qinq**}

no encap-type

Context

[Tree] (config>port>ethernet encap-type)

Full Context

configure port ethernet encap-type

Description

This command configures the encapsulation method used to distinguish customer traffic on an Ethernet access port, or different VLANs on a network port.

The **no** form of this command restores the default.

Default

encap-type null

Parameters

dot1q

Ingress frames carry 802.1Q tags where each tag signifies a different service.

null

Ingress frames will not use any tags to delineate a service. As a result, only one service can be configured on a port with a null encapsulation type.

qinq

Specifies QinQ encapsulation.

Platforms

7705 SAR Gen 2

encap-type

Syntax

encap-type {**dot1q** | **null** | **qinq**}

no encap-type

Context

[Tree] (config>lag encap-type)

Full Context

configure lag encap-type

Description

This command configures the encapsulation method used to distinguish customer traffic on a LAG. The encapsulation type is configurable on a LAG port. The LAG port and the port member encapsulation types must match when adding a port member.

If the encapsulation type of the LAG port is changed, the encapsulation type on all the port members will also change. The encapsulation type can be changed on the LAG port only if there is no interface associated with it. If the MTU is set to a non-default value, it will be reset to the default value when the encap type is changed.

The **no** form of this command restores the default.

Default

encap-type null — All traffic on the port belongs to a single service or VLAN.

Parameters

dot1q

Ingress frames carry 802.1Q tags where each tag signifies a different service.

null

Ingress frames will not use any tags to delineate a service. As a result, only one service can be configured on a port with a null encapsulation type.

qinq

Specifies QinQ encapsulation.

Platforms

7705 SAR Gen 2

9.46 encapsulated-ip-mtu

encapsulated-ip-mtu

Syntax

encapsulated-ip-mtu *bytes*

no encapsulated-ip-mtu

Context

[Tree] (config>service>vprn>if>sap>ip-tunnel encapsulated-ip-mtu)

[Tree] (config>ipsec>tnl-temp encapsulated-ip-mtu)

[Tree] (config>service>ies>if>sap>ip-tunnel encapsulated-ip-mtu)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel encapsulated-ip-mtu)

[Tree] (config>service>vprn>if>sap>ipsec-tun encapsulated-ip-mtu)

[Tree] (config>router>if>ipsec>ipsec-tunnel encapsulated-ip-mtu)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel encapsulated-ip-mtu)

Full Context

configure service vprn interface sap ip-tunnel encapsulated-ip-mtu

configure ipsec tunnel-template encapsulated-ip-mtu

configure service ies interface sap ip-tunnel encapsulated-ip-mtu

configure service ies interface ipsec ipsec-tunnel encapsulated-ip-mtu

configure service vprn interface sap ipsec-tunnel encapsulated-ip-mtu

configure router interface ipsec ipsec-tunnel encapsulated-ip-mtu

configure service vprn interface ipsec ipsec-tunnel encapsulated-ip-mtu

Description

This command specifies the maximum size of encapsulated tunnel packet for the ipsec-tunnel, ip-tunnel, or the dynamic tunnels terminated on the ipsec-gw. If the encapsulated IPv4 or IPv6 tunnel packet exceeds the **encapsulated-ip-mtu**, then the system fragments the packet against the encapsulated-ip-mtu.

The **no** form of this command reverts to the default.

Default

no encapsulated-ip-mtu

Parameters

bytes

Specifies the maximum size in bytes.

Values 512 to 9000

Platforms

7705 SAR Gen 2

encapsulated-ip-mtu

Syntax

encapsulated-ip-mtu *octets*

no encapsulated-ip-mtu

Context

[Tree] (config>service>vprn>if>sap>ip-tunnel encapsulated-ip-mtu)

Full Context

configure service vprn interface sap ip-tunnel encapsulated-ip-mtu

Description

This command configures the tunnel encapsulated IP MTU.

The **no** form of this command reverts to the default.

Parameters

octets

Specifies the tunnel encapsulated IP MTU in octets.

Platforms

7705 SAR Gen 2

9.47 encoding

encoding

Syntax

encoding *encoding*

no encoding

Context

[Tree] (config>system>telemetry>persistent-subscriptions>subscription encoding)

Full Context

configure system telemetry persistent-subscriptions subscription encoding

Description

This command configures the encoding type that is used for telemetry notifications in accordance with the definitions in the gNMI OpenConfig standard.

Default

encoding json

Parameters

encoding

Specifies the encoding type.

Values json, bytes, proto

Platforms

7705 SAR Gen 2

9.48 encrypt

encrypt

Syntax

encrypt {**on** | **off**}

Context

[Tree] (bof encrypt)

Full Context

bof encrypt

Description

This command enables and disables encryption of the BOF using AES256 and SHA256.

When the BOF is encrypted on the compact flash, it is still reachable using the BOF interactive menu during node startup, and fields can be modified using the BOF interactive menu.

Default

encrypt off

Parameters

on

Enables BOF encryption

off

Disables BOF encryption

Platforms

7705 SAR Gen 2

9.49 encryption-key

encryption-key

Syntax

encryption-key *key* [**hash** | **hash2** | **custom**]

no encryption-key

Context

[Tree] (bof encryption-key)

Full Context

bof encryption-key

Description

This command creates a key to be used by AES256 and SHA256 for configuration file encryption and hashing. This key is used for all configuration files (primary, secondary, and tertiary).

After creating or deleting a key, use the **admin save** command to save the configuration file with the current encryption key state.

The **no** form of this command deletes the encryption key.

Default

no encryption-key

Parameters

key

Specifies the encryption key.

If the **hash**, **hash2**, or **custom** parameter is not configured, the key is entered in plaintext and the key length must be between 8 and 32 characters. A plaintext key cannot contain embedded nulls or end with " hash", " hash2", or " custom".

If the **hash**, **hash2**, or **custom** parameter is configured, the key is hashed and the key length must be between 1 and 64 characters.

hash

Keyword to specify that the key is entered in an encrypted form.

hash2

Keyword to specify that the key is entered in a more complex encrypted form. The **hash2** encryption scheme is node-specific and the key cannot be transferred between nodes.

custom

Keyword to specify that the key uses custom encryption.

Platforms

7705 SAR Gen 2

encryption-key

Syntax

encryption-key *key* [**hash** | **hash2** | **custom**]

no encryption-key

Context

[Tree] (config>log encryption-key)

Full Context

configure log encryption-key

Description

This command specifies the encryption key used by AES-256-CTR for log file encryption. The encryption key is used for all local log files on the system.

The **no** form of this command deletes the encryption key.

Default

no encryption-key

Parameters

key

Specifies the encryption key.

If the **hash**, **hash2**, or **custom** parameter is not configured, the key is entered in plaintext and the key length must be between 8 and 32 characters. A plaintext key cannot contain embedded nulls or end with " hash", " hash2", or " custom".

If the **hash**, **hash2**, or **custom** parameter is configured, the key is hashed and the key length must be between 1 and 64 characters.

hash

Keyword to specify that the key is entered in an encrypted form.

hash2

Keyword to specify that the key is entered in a more complex encrypted form. The **hash2** encryption scheme is node-specific and the key cannot be transferred between nodes.

custom

Keyword to specify that the key uses custom encryption.

Platforms

7705 SAR Gen 2

9.50 encryption-keygroup

encryption-keygroup

Syntax

encryption-keygroup *keygroup-id* **direction** {**inbound** | **outbound**}

no encryption-keygroup direction {**inbound** | **outbound**}

Context

[Tree] (config>router>if>group-encryption encryption-keygroup)

Full Context

configure router interface group-encryption encryption-keygroup

Description

This command is used to bind a key group to a router interface for inbound or outbound packet processing. When configured in the outbound direction, packets egressing the router use the **active-outbound-sa** associated with the configured key group. When configured in the inbound direction, received packets must be encrypted using one of the valid security associations configured for the key group.

The **no** form of this command removes the key group from the router interface in the specified direction.

Default

no encryption-keygroup direction inbound

no encryption-keygroup direction outbound

Parameters

keygroup-id

The ID number of the key group being configured.

Values 1 to 127, *keygroup-name* (64 characters maximum)

inbound

Binds the key group in the inbound direction.

outbound

Binds the key group in the outbound direction.

Platforms

7705 SAR Gen 2

encryption-keygroup

Syntax

encryption-keygroup *keygroup-id* [**create**]

no encryption-keygroup *keygroup-id*

Context

[Tree] (config>grp-encryp encryption-keygroup)

Full Context

configure group-encryption encryption-keygroup

Description

This command is used to create a key group. Once the key group is created, use the command to enter the key group context or delete a key group.

The **no** form of the command removes the key group. Before using the **no** form, the key group association must be deleted from all services that are using this key group.

Parameters

keygroup-id

The number or name of the key group being referenced.

Values 1 to 15, or *keygroup-name* (up to 64 characters)

create

Creates a key group.

Platforms

7705 SAR Gen 2

encryption-keygroup

Syntax

encryption-keygroup *keygroup-id* **direction** {**inbound** | **outbound**}

no encryption-keygroup direction {**inbound** | **outbound**}

Context

[Tree] (config>service>vprn encryption-keygroup)

[Tree] (config>service>sdp encryption-keygroup)

[Tree] (config>service>pw-template encryption-keygroup)

Full Context

configure service vprn encryption-keygroup

configure service sdp encryption-keygroup

configure service pw-template encryption-keygroup

Description

This command is used to bind a key group to an SDP, VPRN service, or PW template for inbound or outbound packet processing. When configured in the outbound direction, packets egressing the node use the **active-outbound-sa** associated with the key group configured. When configured in the inbound direction, received packets must be encrypted using one of the valid security associations configured for the key group. Services using the SDP will be encrypted.

The encryption (enabled or disabled) configured on an SDP used to terminate a Layer 3 spoke SDP of a VPRN always overrides any VPRN-level configuration for encryption.

Encryption is enabled after the outbound direction is configured.

For PW template changes, the following **tools** command must be executed after the configuration changes are made: **tools**>**perform**>**service**>**eval-pw-template**>**allow-service-impact**. This command applies the changes to services that use the PW template.

The **no** form of the command removes the key group from the SDP or service in the specified direction (inbound or outbound).

Parameters

keygroup-id

Specifies the number of the key group being configured.

Values 1 to 15 or *keygroup-name* (up to 64 characters)

direction {inbound | outbound}

Specifies the direction of the service that the key group will be bound to.

Platforms

7705 SAR Gen 2

9.51 encryption-offset

encryption-offset

Syntax

encryption-offset *encryption-offset*

no encryption-offset

Context

[Tree] (config>macsec>connectivity-association encryption-offset)

Full Context

configure macsec connectivity-association encryption-offset

Description

This command specifies the offset of the encryption in MACsec packet.

The encryption-offset is distributed by MKA (Key-server) to all parties.

It is signaled via MACsec capabilities. There are four basic settings for this. Table 32: MACsec Basic Settings breaks down the settings.

Table 32: MACsec Basic Settings

Setting	Description
0	MACsec is not implemented
1	Integrity without confidentiality
2	The following are supported: • Integrity without confidentiality • Integrity and confidentiality with a confidentiality offset of 0
3	The following are supported: • Integrity without confidentiality • Integrity and confidentiality with a confidentiality offset of 0, 30, or 50

Note:

1. SR OS supports setting (3) Integrity without confidentiality and Integrity and confidentiality with a confidentiality offset of 0, 30, or 50.

The **no** form of this command rejects all arriving traffic whether MACsec is secured or not.

Default

encryption-offset 0

Parameters

encryption-offset

Specifies the encryption.

Values 0 — encrypt the entire payload
30 — leave the IPv4 header in clear
50 — leave the IPv6 header in clear

Platforms

7705 SAR Gen 2

9.52 end

end

Syntax

end *end-week end-day end-month hours-minutes*

Context

[Tree] (config>system>time>dst-zone end)

Full Context

configure system time dst-zone end

Description

This command configures start of summer time settings.

Default

end first sunday january 00:00

Parameters

end-week

Specifies the starting week of the month when the summer time ends.

Values first, second, third, fourth, last

Default first

end-day

Specifies the starting day of the week when the summer time ends.

Values sunday, monday, tuesday, wednesday, thursday, friday, saturday

Default sunday

end-month

Specifies the starting month of the year when the summer time takes effect.

Values january, february, march, april, may, june, july, august, september, october, november, december

Default january

hours-minutes

Specifies the time at which the summer time ends, in hh:mm format.

Values hours: 00 to 23

minutes: 00 to 59

Default 00:00

Platforms

7705 SAR Gen 2

9.53 end-of-data

end-of-data

Syntax

[**no**] **end-of-data**

Context

[Tree] (debug>router>rpki-session>packet end-of-data)

Full Context

debug router rpki-session packet end-of-data

Description

This command enables debugging for end of data RPKI packets.

The **no** form of this command disables debugging for end of data RPKI packets.

Platforms

7705 SAR Gen 2

9.54 end-time

end-time

Syntax

end-time [*date* | *day-name*] *time*

no end-time

Context

[Tree] (config>system>cron>sched end-time)

Full Context

configure system cron schedule end-time

Description

This command is used concurrently with type **periodic** or **calendar**. Using the type of **periodic**, end-time determines at which interval the schedule will end. Using the type of **calendar**, end-time determines on which date the schedule will end.

When **no end-time** is specified, the schedule runs forever.

Default

no end-time

Parameters

date

Specifies the date to schedule a command.

Values YYYY:MM:DD in year:month:day number format

day-name

Specifies the day of the week to schedule a command.

Values sunday, monday, tuesday, wednesday, thursday, friday, saturday

time

Specifies the time of day to schedule a command.

Values hh:mm

Platforms

7705 SAR Gen 2

end-time

Syntax

end-time *date hours-minutes* [**UTC**]

end-time {**now** | **forever**}

no end-time

Context

[Tree] (config>system>security>keychain>direction>uni>receive>entry end-time)

Full Context

configure system security keychain direction uni receive entry end-time

Description

This command specifies the calendar date and time after which the key specified by the authentication key is no longer eligible to sign or authenticate the protocol stream.

Default

end-time forever

Parameters

date

Specifies the calendar date after which the key specified by the authentication key is no longer eligible to sign or authenticate the protocol stream in the YYYY/MM/DD format. When no year is specified the system assumes the current year.

hours-minutes

Specifies the time after which the key specified by the authentication key is no longer eligible to sign or authenticate the protocol stream in the hh:mm[:ss] format. Seconds are optional, and if not included, assumed to be 0.

UTC

Indicates that time is given with reference to Coordinated Universal Time in the input.

now

Specifies a time equal to the current system time.

forever

Specifies that the key is always active.

Platforms

7705 SAR Gen 2

9.55 endpoint

endpoint

Syntax

endpoint *endpoint-name* [**create**]

no endpoint *endpoint-name*

Context

[Tree] (config>service>epipe endpoint)

Full Context

configure service epipe endpoint

Description

This command configures a service endpoint.

Parameters

endpoint-name

Specifies an endpoint name.

Platforms

7705 SAR Gen 2

endpoint

Syntax

endpoint *endpoint-name* [**create**]

no endpoint

Context

[Tree] (config>service>vpls endpoint)

Full Context

configure service vpls endpoint

Description

This command configures a service endpoint.

Parameters

endpoint-name

Specifies an endpoint name up to 32 characters in length

create

This keyword is mandatory while creating a service endpoint

Platforms

7705 SAR Gen 2

endpoint

Syntax

endpoint *endpoint-name* [**create**]

no endpoint *endpoint-name*

Context

[Tree] (config>mirror>mirror-dest endpoint)

Full Context

configure mirror mirror-dest endpoint

Description

This command configures a service end point. A mirror service supports two implicit endpoints managed internally by the system. The following applies to endpoint configurations.

Up to two named endpoints may be created per service mirror or LI service. The endpoint name is locally significant to the service mirror or LI service.

- Objects (SAPs or SDPs) may be created on the service mirror or LI with the following limitations:
  - two implicit endpoint objects (without explicit endpoints defined)
  - one implicit and multiple explicit object with the same endpoint name
  - multiple explicit objects each with one of two explicit endpoint names
- All objects become associated implicitly or indirectly with the implicit endpoints 'x' and 'y'.
- Objects may be created without an explicit endpoint defined.
- Objects may be created with an explicit endpoint defined.
- Objects without an explicit endpoint may have an explicit endpoint defined without deleting the object.
- Objects with an explicit endpoint defined may be dynamically moved to another explicit endpoint or may have the explicit endpoint removed.

Creating an object without an explicit endpoint:

- If an object on a mirror or LI service has no explicit endpoint name associated, the system attempts to associate the object with implicit endpoint 'x' or 'y'.
- The implicit endpoint cannot have an existing object association.

- If both 'x' and 'y' are available, 'x' is selected.
- If an 'x' or 'y' association cannot be created, the object cannot be created.

Creating an object with an explicit endpoint name:

- The endpoint name must exist on the mirror or LI service.
- If this is the first object associated with the endpoint name:
  - the object is associated with either implicit endpoint 'x' or 'y'
  - the implicit endpoint cannot have an existing object associated
  - if both 'x' and 'y' are available, 'x' is selected
  - if 'x' or 'y' is not available, the object cannot be created
  - the implicit endpoint is now associated with the named endpoint
  - f this is not the first object associated with the endpoint name:
  - the object is associated with the named endpoint's implicit association

Changing an object’s implicit endpoint to an explicit endpoint name

- If the explicit endpoint name is associated with an implicit endpoint, the object is moved to that implicit endpoint
- If the object is the first to be associated with the explicit endpoint name:
  - the object is associated with either implicit endpoint 'x' or 'y'
  - the implicit endpoint cannot have an existing object associated (except this one)
  - if both 'x' and 'y' are available, 'x' is selected
  - if 'x' or 'y' is not available, the object cannot be moved to the explicit endpoint
  - if moved, the implicit endpoint is now associated with the named endpoint

Changing an object’s explicit endpoint to another explicit endpoint name

- If the new explicit endpoint name is associated with an implicit endpoint, the object is moved to that implicit endpoint
- If the object is the first to be associated with the new explicit endpoint name:
  - the object is associated with either implicit endpoint 'x' or 'y'
  - the implicit endpoint cannot have an existing object associated (except this one)
  - if both 'x' and 'y' are available, 'x' is selected
  - if 'x' or 'y' is not available, the object cannot be moved to the new endpoint
  - if moved, the implicit endpoint is now associated with the named endpoint

An explicitly named endpoint can have a maximum of one SAP and one ICB. Once a SAP is added to the endpoint, only one more object of type ICB sdp is allowed. The ICB sdp cannot be added to the endpoint if the SAP is not part of a MC-LAG instance. Conversely, a SAP which is not part of a MC-LAG instance cannot be added to an endpoint which already has an ICB sdp.

An explicitly named endpoint which does not have a SAP object can have a maximum of four SDPs which can include any of the following: a single primary SDP, one or many secondary SDPs with precedence, and a single ICB SDP.

The user can only add a SAP configured on a MC-LAG instance to this endpoint. Conversely, the user will not be able to change the mirror service type away from mirror service without first deleting the MC-LAG SAP.

The **no** form of this command removes the association of a SAP or an SDP with an explicit endpoint name. When removing an objects explicit endpoint association:

- The system attempts to associate the object with implicit endpoint 'x' or 'y'.
- The implicit endpoint cannot have an existing object association (except this one).
- If both 'x' and 'y' are available, 'x' is selected.
- If an 'x' or 'y' association cannot be created, the explicit endpoint cannot be removed.

Parameters

endpoint-name

Specifies the endpoint name.

create

Mandatory keyword to create this entry.

Platforms

7705 SAR Gen 2

endpoint

Syntax

endpoint *ip-address*

no endpoint

Context

[Tree] (conf>router>segment-routing>sr-policies>policy endpoint)

Full Context

configure router segment-routing sr-policies static-policy endpoint

Description

This command associates an IPv4 or IPv6 endpoint address with a statically-defined segment routing policy. This association is mandatory when enabling an SR segment-routing policy.

The endpoint address 0.0.0.0 is a special value that matches all BGP next-hops. To use it, the BGP route must have a color-extended community with the color-only bits set to '01' or '10'.

The **no** form of this command removes the endpoint association.

Default

no endpoint

Parameters

ip-address

Specifies the endpoint IP address to be associated with the statically-defined segment-routing policy.

Values

ipv4-address:	a.b.c.d	
ipv6-address:	x:x:x:x:x:x:x:x	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 to FFFF]H
	d:	[0 to 255]D

Platforms

7705 SAR Gen 2

endpoint

Syntax

endpoint *ip-address*

no endpoint

Context

[Tree] (config>router>policy-options>policy-statement>entry>from endpoint)

Full Context

configure router policy-options policy-statement entry from endpoint

Description

This command configures an SR Policy endpoint address as a route policy match criterion. This match criterion is only used in import policies.

The **no** form of this command removes the endpoint IP match criterion from the configuration.

Parameters

ip-address

Specifies the IPv4 or IPv6 address.

Values ipv4-address:

- a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x:x [-interface]
- x:x:x:x:x:x:d.d.d.d [-interface]

- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

endpoint

Syntax

endpoint *ip-address*

no endpoint

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls>sr-policy endpoint)

Full Context

configure oam-pm session ip tunnel mpls sr-policy endpoint

Description

This command configures the unicast IPv4 or globally routable IPv6 address endpoint of the tunnel.

The **no** form of this command removes IPv4 or IPv6 address.

Default

no endpoint

Parameters

ip-address

Specifies the IPv4 or IPv6 address.

Values		
	ipv4-address	- a.b.c.d
	ipv6-address	- x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x - [0 to FFFF]H
		d - [0 to 255]D

Platforms

7705 SAR Gen 2

9.56 enforce-first-as

enforce-first-as

Syntax

enforce-first-as

Context

[Tree] (config>service>vprn>bgp>group enforce-first-as)

[Tree] (config>service>vprn>bgp>group>neighbor enforce-first-as)

[Tree] (config>service>vprn>bgp enforce-first-as)

Full Context

configure service vprn bgp group enforce-first-as

configure service vprn bgp group neighbor enforce-first-as

configure service vprn bgp enforce-first-as

Description

When this command is configured so that it applies to an EBGP session, all routes (belonging to all address families) that are received from the EBGP peer are checked to ensure that the most recent autonomous system number (ASN) in the AS_PATH attribute of each route matches the configured **peer-as** of the session; if it does not match, then either the session is reset (if **update-fault-tolerance** is not enabled) or the session is left up but the route is treated as withdrawn (if **update-fault-tolerance** is enabled).

Enabling or disabling this command on a session that is already up does not flap the session. When **enforce-first-as** is enabled, previously received routes are not checked for compliance with the rule. Enforcement applies only to routes received after the command is enabled and stops when the command is disabled.

Platforms

7705 SAR Gen 2

enforce-first-as

Syntax

enforce-first-as

Context

[Tree] (config>router>bgp>group enforce-first-as)

[Tree] (config>router>bgp>group>neighbor enforce-first-as)

[Tree] (config>router>bgp enforce-first-as)

Full Context

configure router bgp group enforce-first-as

configure router bgp group neighbor enforce-first-as

configure router bgp enforce-first-as

Description

When this command is configured so that it applies to an EBGP session, all routes (belonging to all address families) that are received from the EBGP peer are checked to ensure that the most recent autonomous system number (ASN) in the AS_PATH attribute of each route matches the configured **peer-as** of the session; if it does not match, then either the session is reset (if **update-fault-tolerance** is not enabled) or the session is left up but the route is treated as withdrawn (if **update-fault-tolerance** is enabled).

Enabling or disabling this command on a session that is already up does not flap the session. When **enforce-first-as** is enabled, previously received routes are not checked for compliance with the rule. Enforcement applies only to routes received after the command is enabled and stops when the command is disabled.

Platforms

7705 SAR Gen 2

9.57 enforce-strict-tunnel-tagging

enforce-strict-tunnel-tagging

Syntax

[**no**] **enforce-strict-tunnel-tagging**

Context

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel enforce-strict-tunnel-tagging)

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel enforce-strict-tunnel-tagging)

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel enforce-strict-tunnel-tagging)

Full Context

configure service vpls bgp-evpn mpls auto-bind-tunnel enforce-strict-tunnel-tagging

configure service epipe bgp-evpn mpls auto-bind-tunnel enforce-strict-tunnel-tagging

configure service vprn bgp-evpn mpls auto-bind-tunnel enforce-strict-tunnel-tagging

Description

This command forces the system to only consider LSPs marked with an admin tag for next hop resolution. Untagged LSPs are not considered.

The **no** form of this command reverts to default value. While tagged RSVP and SR-TE LSPs are considered first, the system can fall back to using untagged LSPs of other types and does not exclude them depending on the **auto-bind-tunnel** configuration.

Default

no enforce-strict-tunnel-tagging

Platforms

7705 SAR Gen 2

enforce-strict-tunnel-tagging

Syntax

[**no**] **enforce-strict-tunnel-tagging**

Context

[Tree] (config>router>bgp>next-hop-resolution>shortcut-tunn>family enforce-strict-tunnel-tagging)

[Tree] (config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family enforce-strict-tunnel-tagging)

Full Context

configure router bgp next-hop-resolution shortcut-tunnel family enforce-strict-tunnel-tagging

configure router bgp next-hop-resolution labeled-routes transport-tunnel family enforce-strict-tunnel-tagging

Description

This command forces the system to only consider LSPs marked with an **admin-tag** for next-hop resolution. Untagged LSPs are not be considered.

The **no** form of this command reverts to the default behavior. While tagged RSVP and SR-TE LSPs will be considered first, the system can fall back to using tagged LSPs that are not explicitly excluded by a route admin tag policy and untagged LSPs of other types and not exclude them.

Default

no enforce-strict-tunnel-tagging

Platforms

7705 SAR Gen 2

enforce-strict-tunnel-tagging

Syntax

enforce-strict-tunnel-tagging

Context

[Tree] (config>service>vprn>auto-bind-tunnel enforce-strict-tunnel-tagging)

Full Context

configure service vprn auto-bind-tunnel enforce-strict-tunnel-tagging

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

9.58 enforce-test-session-start-time

enforce-test-session-start-time

Syntax

[**no**] **enforce-test-session-start-time**

Context

[Tree] (config>test-oam>twamp>server enforce-test-session-start-time)

Full Context

configure test-oam twamp server enforce-test-session-start-time

Description

This command configures the router to check the signalled test-session start time against the server time and discard TWAMP test packets that arrive before the negotiated test-session start time.

The **no** form of this command configures the router to process all TWAMP test packets without checking the test-session start time against the server time.

Default

enforce-test-session-start-time

Platforms

7705 SAR Gen 2

9.59 enforce-unique-if-index

enforce-unique-if-index

Syntax

[**no**] **enforce-unique-if-index**

Context

[Tree] (config>system>ip enforce-unique-if-index)

Full Context

configure system ip enforce-unique-if-index

Description

This command enables the options to force the creation of IP interface indexes so that they are globally unique across all routing contexts. In addition, the command ensures that any interface created using SNMP also has a system-wide unique IP interface index.

If this command is issued but the system has previously existing interface indexes that conflict, the command will be rejected until all the conflicts are removed. Pre-existing persistency tables should also be removed before enabling this system option.

The **no** form of the command disables this option and returns the system to the default behavior.

Default

no enforce-unique-if-index

Platforms

7705 SAR Gen 2

9.60 enforce-untagged-route

enforce-untagged-route

Syntax

enforce-untagged-route {**none** | **untagged-tunnel**}

Context

[Tree] (config>router>bgp>next-hop-res>lbl-routes>transport-tunn>family enforce-untagged-route)

Full Context

configure router bgp next-hop-resolution labeled-routes transport-tunnel family enforce-untagged-route

Description

This command configures the enforcement of BGP routes with no administrative tag policy applied by modifying the next-hop resolution behavior for autobind services.

Default

enforce-untagged-route none

Parameters

none

Keyword to specify that untagged routes can bind to tagged or untagged LSPs.

untagged-tunnel

Keyword to specify that untagged routes can only bind to LSPs with no administrative tags configured. If both tagged and untagged tunnels to the next hop exist, the system only considers untagged tunnels. If no untagged tunnels to the next hop exist, the resolution of untagged routes also fails. This keyword may be used in combination with the **enforce-strict-tunnel-tagging** command, in which case tagged routes resolve to tagged LSPs and untagged routes only resolve to untagged LSPs.

Platforms

7705 SAR Gen 2

enforce-untagged-route

Syntax

enforce-untagged-route {**none** | **untagged-tunnel**}

Context

[Tree] (config>router>bgp>next-hop-res>shortcut-tunn>family enforce-untagged-route)

Full Context

configure router bgp next-hop-resolution shortcut-tunnel family enforce-untagged-route

Description

This command configures the enforcement of BGP routes with no administrative tag policy applied by modifying the next-hop resolution behavior for autobind services.

Default

enforce-untagged-route none

Parameters

none

Keyword to specify that untagged routes can bind to tagged or untagged LSPs.

untagged-tunnel

Keyword to specify that untagged routes can only bind to LSPs with no administrative tags configured. If both tagged and untagged tunnels to the next hop exist, the system only considers untagged tunnels. If no untagged tunnels to the next hop exist, the resolution of untagged routes also fails. This keyword may be used in combination with the **enforce-strict-tunnel-tagging** command, in which case tagged routes resolve to tagged LSPs and untagged routes only resolve to untagged LSPs.

Platforms

7705 SAR Gen 2

enforce-untagged-route

Syntax

enforce-untagged-route {**none** | **untagged-tunnel**}

Context

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel enforce-untagged-route)

Full Context

configure service epipe bgp-evpn mpls auto-bind-tunnel enforce-untagged-route

Description

This command configures the enforcement of BGP routes with no administrative tag policy applied by modifying the next-hop resolution behavior for autobind services.

Default

enforce-untagged-route none

Parameters

none

Keyword to specify that untagged routes can bind to tagged or untagged LSPs.

untagged-tunnel

Keyword to specify that untagged routes can only bind to LSPs with no administrative tags configured. If both tagged and untagged tunnels to the next hop exist, the system only considers untagged tunnels. If no untagged tunnels to the next hop exist, the resolution of untagged routes also fails. This keyword may be used in combination with the **enforce-**

strict-tunnel-tagging command, in which case tagged routes resolve to tagged LSPs and untagged routes only resolve to untagged LSPs.

Platforms

7705 SAR Gen 2

enforce-untagged-route

Syntax

enforce-untagged-route {**none** | **untagged-tunnel**}

Context

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel enforce-untagged-route)

Full Context

configure service vpls bgp-evpn mpls auto-bind-tunnel enforce-untagged-route

Description

This command configures the enforcement of BGP routes with no administrative tag policy applied by modifying the next-hop resolution behavior for autobind services.

Default

enforce-untagged-route none

Parameters

none

Keyword to specify that untagged routes can bind to tagged or untagged LSPs.

untagged-tunnel

Keyword to specify that untagged routes can only bind to LSPs with no administrative tags configured. If both tagged and untagged tunnels to the next hop exist, the system only considers untagged tunnels. If no untagged tunnels to the next hop exist, the resolution of untagged routes also fails. This keyword may be used in combination with the **enforce-strict-tunnel-tagging** command, in which case tagged routes resolve to tagged LSPs and untagged routes only resolve to untagged LSPs.

Platforms

7705 SAR Gen 2

enforce-untagged-route

Syntax

enforce-untagged-route {**none** | **untagged-tunnel**}

Context

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel enforce-untagged-route)

Full Context

configure service vprn bgp-evpn mpls auto-bind-tunnel enforce-untagged-route

Description

This command configures the enforcement of BGP routes with no administrative tag policy applied by modifying the next-hop resolution behavior for autobind services.

Default

enforce-untagged-route none

Parameters

none

Keyword to specify that untagged routes can bind to tagged or untagged LSPs.

untagged-tunnel

Keyword to specify that untagged routes can only bind to LSPs with no administrative tags configured. If both tagged and untagged tunnels to the next hop exist, the system only considers untagged tunnels. If no untagged tunnels to the next hop exist, the resolution of untagged routes also fails. This keyword may be used in combination with the **enforce-strict-tunnel-tagging** command, in which case tagged routes resolve to tagged LSPs and untagged routes only resolve to untagged LSPs.

Platforms

7705 SAR Gen 2

enforce-untagged-route

Syntax

enforce-untagged-route {**none** | **untagged-tunnel**}

Context

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel enforce-untagged-route)

Full Context

configure service vprn bgp-ipvpn mpls auto-bind-tunnel enforce-untagged-route

Description

This command configures the enforcement of BGP routes with no administrative tag policy applied by modifying the next-hop resolution behavior for autobind services.

Default

enforce-untagged-route none

Parameters

none

Keyword to specify that untagged routes can bind to tagged or untagged LSPs.

untagged-tunnel

Keyword to specify that untagged routes can only bind to LSPs with no administrative tags configured. If both tagged and untagged tunnels to the next hop exist, the system only considers untagged tunnels. If no untagged tunnels to the next hop exist, the resolution of untagged routes also fails. This keyword may be used in combination with the **enforce-strict-tunnel-tagging** command, in which case tagged routes resolve to tagged LSPs and untagged routes only resolve to untagged LSPs.

Platforms

7705 SAR Gen 2

9.61 enforcement

enforcement

Syntax

enforcement {**static** *policer-name* | **dynamic** {*mon-policer-name* | **local-mon-bypass**}}

Context

[Tree] (config>sys>security>dist-cpu-protection>policy>protocol enforcement)

Full Context

configure system security dist-cpu-protection policy protocol enforcement

Description

This command configures the enforcement method for the protocol.

Default

enforcement dynamic local-mon-bypass

Parameters

static

Specifies that the protocol is always enforced using a **static-policer**. Multiple protocols can reference the same **static-policer**. Packets of protocols that are statically enforced bypass any local monitors.

policer name

Specifies which **static-policer** to use.

dynamic

Specifies that a specific enforcement policer for this protocol for this SAP/object is instantiated when the associated **local-monitoring-policer** is determined to be in a nonconforming state (at the end of a minimum monitoring time of 60 seconds to reduce thrashing).

mon-policer-name

Specifies which **local-monitoring-policer** to use.

local-mon-bypass

This parameter is used to not include packets from this protocol in the local monitoring function, and when the local-monitor "trips", do not instantiate a dynamic enforcement policer for this protocol.

Platforms

7705 SAR Gen 2

9.62 engineID

engineID

Syntax

[**no**] **engineID** *engine-id*

Context

[Tree] (config>system>snmp engineID)

Full Context

configure system snmp engineID

Description

This command sets the SNMP engine ID that uniquely identifies the SNMPv3 node.If unconfigured, the system uses an engine ID based on the information from the system backplane.If the SNMP engine ID is changed, the current configuration must be saved and a reboot must be executed. Otherwise, the previously configured SNMP communities and logger trap-target notify communities will not be valid for the new engine ID.

Note:
Changing the SNMP engine ID invalidates all SNMPv3 MD5 and SHA security digest keys, which may render the node unmanageable.

When replacing a chassis, configure the new router to use the same engine ID as the previous router. This preserves SNMPv3 security keys and allows management stations to use their existing authentication keys for the new router.

Ensure that the engine ID of each router is unique. A management domain can only maintain one instance of a specific engine ID.

The **no** form of the command configures the router to use the default value.

Parameters

engine-id

Specifies an identifier from 10 to 64 hexadecimal digits (5 to 32 octet number), uniquely identifying this SNMPv3 node. This string is used to access this node from a remote host with SNMPv3.

Platforms

7705 SAR Gen 2

9.63 enroll

enroll

Syntax

enroll est-profile *name* **key** *key-filename* **output** *output-cert-filename* [**hash-alg** *hash algorithm*] **subject-dn** *subject-dn* [**domain-name** *domain-names*] [**ip-addr** *ip-address* **|** *ipv6-address*] [**validate-cert-chain**] [**force**]

Context

[Tree] (admin>certificate>est enroll)

Full Context

admin certificate est enroll

Description

This command enrolls a new certificate with Certificate Authority (CA) by the EST protocol specified with the **est-profile** *name* parameter with a imported private key specified by the **key** *key-filename* parameter.

The **est-profile** *name* specifies the authentication between the system and EST server.

The **hash-alg** *hash-algorithm*, **subject-dn** *subject-dn*, **domain-name** *domain-names*, and **ip-addr** *ip-address* parameters are used to generate the Certificate Signing Request (CSR) in the EST request message. The **domain-name** *domain-names* and **ip-addr** *ip-address* parameters are used as subject alternative names.

If **validate-cert-chain** is specified, the system validates the certificate’s chain of result certificate before importing it. The "certificate chain” is the chain of all the certificates from the result certificate to the issuing CA. The "result certificate” is the new certificate returned by EST server.

The result certificate is imported and saved with the filename specified by the **output** *output-cert-filename*. If **force** is specified, the system overwrites the existing file with same name as the *output-cert-filename*.

Parameters

name

Specifies EST profile name, up to 32 characters

key-filename

Specifies the filename of a key, up to 95 characters

output-cert-filename

Specifies the output certificate filename, up to 200 characters

hash-algorithm

Specifies the hash algorithm used in a certificate request.

Values sha1, sha224, sha256, sha384, sha512

subject-dn

Specifies the distinguish name, up to 256 characters, used as the subject in a certificate request, including:

- C-Country
- ST-State
- O-Organization name
- OU-Organization Unit name
- CN-common name

This parameter is formatted as a text string including any of the preceding attributes. The attribute and its value is linked by using "=", and "," is used to separate different attributes.

For example: C=US,ST=CA,O=ALU,CN=SR12

Values attr1=val1,attr2=val2

where: attrN={C | ST | O | OU | CN}, up to 256 characters

domain-names

Specifies domain names, up to 512 characters, separated by commas

ip-address

Specifies an IPv4 or IPv6 address string, up to 64 characters

validate-cert-chain

Specifies that the system validates the certificate's chain of result certificate before importing it

force

Specifies that the system overwrites the existing file with same *output-cert-filename*

Platforms

7705 SAR Gen 2

9.64 enter

enter

Syntax

[**no**] **enter**

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>command-completion enter)

Full Context

configure system management-interface cli md-cli environment command-completion enter

Description

This command enables completion on the enter character.

The **no** form of this command reverts to the default value.

Default

enter

Platforms

7705 SAR Gen 2

9.65 entry

entry

Syntax

entry *entry-id* [**create**]

no entry *entry-id*

Context

[Tree] (config>filter>dhcp-filter entry)

[Tree] (config>filter>dhcp6-filter entry)

Full Context

configure filter dhcp-filter entry

configure filter dhcp6-filter entry

Description

This command configures DHCP filter entries.

The **no** form of this command removes the entry from the configuration.

Parameters

entry-id

Specifies the entry ID.

Values 1 to 65535

create

This keyword is required when first creating the DHCP filter entry. Once the context is created, it is possible to navigate into the context without the **create** keyword.

Platforms

7705 SAR Gen 2

entry

Syntax

entry *entry-id* [**name** *entry-name*]

no entry *entry-id*

Context

[Tree] (config>service>vprn>log>filter entry)

Full Context

configure service vprn log filter entry

Description

This command is used to create or edit an event filter entry. Multiple entries may be created using unique *entry-id* values. The SR OS implementation exits the filter on the first match found and executes the action in accordance with the action command.

Comparisons are performed in an ascending entry ID order. When entries are created, they should be arranged sequentially from the most explicit entry to the least explicit. Matching ceases when a packet matches an entry. The entry action is performed on the packet, either drop or forward. To be considered a match, the packet must meet all the conditions defined in the entry.

An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword action for it to be considered complete. Entries without the action keyword will be considered incomplete and are rendered inactive.

By default, no filter entries are defined. Entries must be explicitly configured.

The **no** form of this command removes the specified entry from the event filter. Entries removed from the event filter are immediately removed from all log-id's where the filter is applied.

Default

No event filter entries are defined. An entry must be explicitly configured.

Parameters

entry-id

The entry ID uniquely identifies a set of match criteria corresponding action within a filter. Entry ID values should be configured in staggered increments so you can insert a new entry in an existing policy without renumbering the existing entries.

Values 1 to 999

name entry-name

Configures an optional entry name for the event filter, up to 64 characters, that can be used to refer to the entry after it is created.

Platforms

7705 SAR Gen 2

entry

Syntax

entry *entry-id* [**create**]

no entry *entry-id*

Context

[Tree] (config>ipsec>cert-profile entry)

Full Context

configure ipsec cert-profile entry

Description

This command configures the certificate profile entry information

The **no** form of this command removes the *entry-id* value from the cert-profile configuration.

Parameters

entry-id

Specifies the entry ID.

Values 1 to 8

Platforms

7705 SAR Gen 2

entry

Syntax

entry *entry-id* [**create**]

no entry *entry-id*

Context

[Tree] (config>ipsec>ts-list>remote entry)

[Tree] (config>ipsec>ts-list>local entry)

Full Context

configure ipsec ts-list remote entry

configure ipsec ts-list local entry

Description

This command creates a new TS-list entry or enables the context to configure an existing TS-list entry.

The **no** form of this command removes the entry from the local or remote configuration.

Parameters

entry-id

Specifies the entry ID

Values 1 to 32

Platforms

7705 SAR Gen 2

entry

Syntax

entry *entry-id* [**create**]

no entry *entry-id*

Context

[Tree] (config>router>ipsec>sec-plcy entry)

[Tree] (config>service>vprn>ipsec>sec-plcy entry)

Full Context

configure router ipsec security-policy entry

configure service vprn ipsec security-policy entry

Description

This command configures an IPsec security policy entry.

Parameters

entry-id

Specifies the IPsec security policy entry.

Values 1 to 16

create

Keyword used to create the security policy entry instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

entry

Syntax

entry *entry-id* [**create**]

no entry *entry-id*

Context

[Tree] (config>qos>sap-ingress>ip-criteria entry)

[Tree] (config>qos>sap-ingress>mac-criteria entry)

[Tree] (config>qos>sap-egress>ip-criteria entry)

[Tree] (config>qos>sap-ingress>ipv6-criteria entry)

[Tree] (config>qos>sap-egress>ipv6-criteria entry)

Full Context

configure qos sap-ingress ip-criteria entry

configure qos sap-ingress mac-criteria entry

configure qos sap-egress ip-criteria entry

configure qos sap-ingress ipv6-criteria entry

configure qos sap-egress ipv6-criteria entry

Description

This command is used to create or edit an IP, IPv6, or MAC criteria entry for the policy. Multiple entries can be created using unique *entry-id* numbers.

The list of flow criteria is evaluated in a top-down manner with the lowest entry ID at the top and the highest entry ID at the bottom. If the defined match criteria for an entry within the list matches the information in the egress packet, the system stops matching the packet against the list and performs the

matching entries reclassification actions. If none of the entries match the packet, the IP flow reclassification list has no effect on the packet.

An entry is not populated in the list unless the action command is executed for the entry. An entry that is not populated in the list has no effect on egress packets. If the action command is executed without any explicit reclassification actions specified, the entry is populated in the list allowing packets matching the entry to exit the list, preventing them from matching entries lower in the list. Since this is the only flow reclassification entry that the packet matched and this entry explicitly states that no reclassification action is to be performed, the matching packet will not be reclassified.

The **no** form of this command removes the specified entry from the policy. Entries removed from the policy are immediately removed from all services where that policy is applied.

Parameters

entry-id

The *entry-id,* expressed as an integer, uniquely identifies a match criterion and the corresponding action. It is recommended that multiple entries be given *entry-ids* in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

An entry cannot have any match criteria defined (in which case, everything matches) but must have at least the keyword **action fc** *fc-name* for it to be considered complete. Entries without the action keyword will be considered incomplete and, therefore, will be rendered inactive.

Values 1 to 65535

create

Required parameter when creating a flow entry when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the flow entry already exists.

Platforms

7705 SAR Gen 2

entry

Syntax

entry *entry-id* [**create**]

no entry *entry-id*

Context

[Tree] (config>qos>network>ingress>ipv6-criteria entry)

[Tree] (config>qos>network>ingress>ip-criteria entry)

[Tree] (config>qos>network>egress>ipv6-criteria entry)

[Tree] (config>qos>network>egress>ip-criteria entry)

Full Context

configure qos network ingress ipv6-criteria entry

configure qos network ingress ip-criteria entry

configure qos network egress ipv6-criteria entry

configure qos network egress ip-criteria entry

Description

This command is used to create or edit an IP or IPv6 criteria entry for the policy. Multiple entries can be created using unique entry numbers.

The list of flow criteria is evaluated in a top-down manner with the lowest entry ID at the top and the highest entry ID at the bottom. If the defined match criteria for an entry within the list matches the information in the packet, the system stops matching the packet against the list and performs the matching entries reclassification actions. If none of the entries match the packet, the IP flow reclassification list has no effect on the packet.

An entry is not populated in the list unless the action command is executed for the entry. An entry that is not populated in the list has no effect on ingress packets. If the **action** command is executed without any explicit reclassification actions specified, the entry is populated in the list allowing packets matching the entry to exit the list, preventing them from matching entries lower in the list. Since this is the only flow reclassification entry that the packet matched, and this entry explicitly states that no reclassification action is to be performed, the matching packet will not be reclassified.

The configuration of egress prec/DSCP classification and the configuration of an egress IP criteria or IPv6 criteria entry statement within a network QoS policy are mutually exclusive.

Network QoS policies containing egress **ip-criteria** or **ipv6-criteria entry** statements are only applicable to network interfaces. Configuration of **ip-criteria** or **ipv6-criteria entry** statements in a network egress QoS policy and the application of the policy on any object other than a GRT network interface are mutually exclusive.

The **no** form of this command removes the specified entry from the policy. Entries removed from the policy are immediately removed from all services to which that policy is applied.

Parameters

entry-id

The entry identifier, expressed as an integer, uniquely identifies a match criterion and the corresponding action. It is recommended that multiple entries be given entry identifiers in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

An entry cannot have any match criteria defined (in which case, everything matches) but must have at least the keyword **action fc fc-name profile** *profile* for it to be considered complete. Entries without the action keyword will be considered incomplete and will be rendered inactive.

Values 1 to 65535

create

Required parameter when creating a flow entry when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled, and an object name is mistyped

when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the flow entry already exists.

Platforms

7705 SAR Gen 2

entry

Syntax

entry *entry-id* [**create**]

no entry *entry-id*

Context

[Tree] (config>filter>ip-filter entry)

[Tree] (config>filter>ipv6-exception entry)

[Tree] (config>filter>ipv6-filter entry)

[Tree] (config>filter>ip-exception entry)

Full Context

configure filter ip-filter entry

configure filter ipv6-exception entry

configure filter ipv6-filter entry

configure filter ip-exception entry

Description

This command creates or edits an IPv4, IPv6, MAC, IP exception filter, or IPv6 exception filter entry. Multiple entries can be created using unique *entry-id* numbers within the filter. Entries must be sequenced from most to least explicit.

An entry may not have any match criteria defined (in which case everything matches) but must have at least the keyword **action** for it to be considered complete. Entries without the **action** keyword will be considered incomplete and hence will be rendered inactive.

The **no** form of the command removes the specified entry from the filter. Entries removed from the filter are immediately removed from all services or network ports where that filter is applied.

Parameters

entry-id

Uniquely identifies a match criteria and the corresponding action. It is recommended that multiple entries be given *entry-id* in staggered increments. This allows users to insert a new entry in an existing policy without requiring to renumbering all the existing entries. The parameter is expressed as a decimal integer.

Values 1 to 2097151

create

This keyword is required to create the configuration context. Once the context is created, the user can enable the context with or without the **create** keyword.

Platforms

7705 SAR Gen 2

entry

Syntax

entry *entry-id* [**name** *entry-name*]

no entry *entry-id*

Context

[Tree] (config>log>filter entry)

Full Context

configure log filter entry

Description

This command creates or edits an event filter entry. Multiple entries can be created using unique *entry-id* values. The SR OS implementation exits the filter on the first match found and executes the action in accordance with the **action** command.

Comparisons are performed in an ascending entry ID order. When entries are created, they should be arranged sequentially from the most explicit entry to the least explicit. Matching ceases when a packet matches an entry. The entry action is performed on the packet, either drop or forward. To be considered a match, the packet must meet all the conditions defined in the entry.

An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword action for it to be considered complete. Entries without the action keyword will be considered incomplete and are rendered inactive.

By default, no filter entries are defined. Entries must be explicitly configured.

The **no** form of this command removes the specified entry from the event filter. Entries removed from the event filter are immediately removed from all log-id's where the filter is applied.

Parameters

entry-id

The entry ID uniquely identifies a set of match criteria corresponding action within a filter. Entry ID values should be configured in staggered increments so you can insert a new entry in an existing policy without renumbering the existing entries.

Values 1 to 999

name entry-name

Configures an optional entry name for the event filter, up to 64 characters, that can be used to refer to the entry after it is created.

Platforms

7705 SAR Gen 2

entry

Syntax

[**no**] **entry** *entry-id*

Context

[Tree] (config>log>event-handling>handler>action-list entry)

Full Context

configure log event-handling handler action-list entry

Description

This command configures an EHS handler action-list entry. A handler can have multiple actions where each action, for example, could request the execution of a different script. When the handler is triggered it will walk through the list of configured actions.

The **no** form of this command removes the specified EHS handler action-list entry.

Parameters

entry-id

Specifies the identifier of the EHS handler entry.

Values 1 to 1500

Platforms

7705 SAR Gen 2

entry

Syntax

[**no**] **entry** *entry-id*

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter entry)

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter entry)

[Tree] (config>system>security>mgmt-access-filter>ip-filter entry)

Full Context

configure system security management-access-filter mac-filter entry

configure system security management-access-filter ipv6-filter entry

configure system security management-access-filter ip-filter entry

Description

This command is used to create or edit a management access IP(v4), IPv6, or MAC filter entry. Multiple entries can be created with unique *entry-id* numbers. The OS exits the filter upon the first match found and executes the actions according to the respective action command. For this reason, entries must be sequenced correctly from most to least explicit.

An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword **action** defined to be considered complete. Entries without the **action** keyword are considered incomplete and inactive.

The **no** form of this command removes the specified entry from the management access filter.

Parameters

entry-id

Specifies an entry ID uniquely identifies a match criteria and the corresponding action. It is recommended that entries are numbered in staggered increments. This allows users to insert a new entry in an existing policy without having to renumber the existing entries.

Values 1 to 9999

Platforms

7705 SAR Gen 2

entry

Syntax

[**no**] **entry** *entry-id*

Context

[Tree] (config>system>security>profile entry)

Full Context

configure system security profile entry

Description

This command is used to create a user profile entry.

More than one entry can be created with unique *entry-id* numbers. Exits when the first match is found and executes the actions according to the accompanying **action** command. Entries should be sequenced from most explicit to least explicit.

An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword **action** for it to be considered complete.

The **no** form of this command removes the specified entry from the user profile.

Parameters

entry-id

Specifies an entry-id that uniquely identifies a user profile command match criteria and a corresponding action. If more than one entry is configured, the *entry-ids* should be numbered in staggered increments to allow users to insert a new entry without requiring renumbering of the existing entries.

Values 1 to 9999

Platforms

7705 SAR Gen 2

entry

Syntax

entry *entry-id* [**key** *authentication-key* | *hash-key* | *hash2-key* | *custom-key*] [**hash** | **hash2** | **custom**] **algorithm** *algorithm*]

no entry *entry-id*

Context

[Tree] (config>system>security>keychain>direction>bi entry)

[Tree] (config>system>security>keychain>direction>uni>send entry)

[Tree] (config>system>security>keychain>direction>uni>receive entry)

Full Context

configure system security keychain direction bi entry

configure system security keychain direction uni send entry

configure system security keychain direction uni receive entry

Description

This command defines a particular key in the keychain. Entries are defined by an entry ID. A keychain must have valid entries for the TCP Enhanced Authentication mechanism to work.

If the entry is the active entry for sending, then this causes a new active key to be selected (if one is available using the youngest key rule). If it is the only possible key to send, then the system rejects the command with an error indicating the configured key is the only available send key.

If the key is one of the eligible keys for receiving, it will be removed. If the key is the only possible eligible key, then the command is accepted, and an error indicating that this is the only eligible key will be generated.

The **no** form of this command removes the entry from the keychain.

Parameters

entry-id

Specifies an entry that represents a key configuration to be applied to a keychain.

Values 0 to 63, null-key

key

Specifies a key ID which is used along with *keychain-name* and **direction** to uniquely identify this particular key entry.

authentication-key

Specifies the *authentication-key* that is used by the encryption algorithm. The key is used to sign and authenticate a protocol packet.

The *authentication-key* can be any combination of letters or numbers.

Values A key must be 160 bits for algorithm hmac-sha-1-96 and must be 128 bits for algorithm aes-128-cmac-96. If the key given with the entry command amounts to less than this number of bits, then it is padded internally with zero bits up to the correct length.

algorithm

Specifies an enumerated integer that indicates the encryption algorithm to be used by the key defined in the keychain.

Values aes-128-cmac-96 — Specifies an algorithm based on the AES standard for TCP authentication as described in RFC 4494 for BGP and LDP.

aes-128-cmac-128 — Specifies an algorithm based on the AES standard as described in RFC 4493 for NTP.

aes-128-gcm-16 — Specifies an algorithm used for MCS.

hmac-sha-1-96 — Specifies an algorithm based on SHA-1 for RSVP-TE and TCP authentication.

message-digest — MD5 hash used for TCP authentication.

hmac-md5 — MD5 hash used for IS-IS and RSVP-TE.

password – Specifies a simple password authentication for OSPF, IS-IS, and RSVP-TE.

hmac-sha-1 — Specifies the sha-1 algorithm for OSPF, IS-IS, and RSVP-TE.

hmac-sha-256 — Specifies the sha-256 algorithm for OSPF and IS-IS.

hash-key | hash2-key | custom-key

Specifies the hash key. The key can be any combination of ASCII characters up to 33 for the *hash-key* and 96 characters for the *hash2-key* (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies a custom hash version is used while saving the configuration files.

Platforms

7705 SAR Gen 2

entry

Syntax

entry *entry-id* [**create**]

no entry *entry-id*

Context

[Tree] (config>system>security>tls>cert-profile entry)

Full Context

configure system security tls cert-profile entry

Description

This command configures an entry for the TLS certificate profile. A certificate profile may have up to eight entries. Currently, TLS uses the entry with the smallest ID number when responding to server requests.

The **no** form of the command deletes the specified entry.

Parameters

entry-id

Specifies the identification number of the TLS certificate profile entry.

Values 1 to 8

create

Keyword used to create the TLS certificate profile entry.

Platforms

7705 SAR Gen 2

entry

Syntax

entry *entry-id* **expression** *regular-expression*

no entry *entry-id*

Context

[Tree] (config>router>policy-options>as-path-group entry)

Full Context

configure router policy-options as-path-group entry

Description

This command creates the context to edit route policy entries within an autonomous system path group.

Multiple entries can be created using unique entries. The router exits the filter when the first match is found and executes the action specified. For this reason, entries must be sequenced correctly from most to least explicit.

An entry does not require matching criteria defined (in which case, everything matches) but must at least define an action in order to be considered complete. Entries without an action are considered incomplete and will be rendered inactive.

The **no** form of this command removes the specified entry from the autonomous system path group.

Parameters

entry-id

Specifies the entry ID expressed as a decimal integer. An *entry-id* uniquely identifies match criteria and the corresponding action. Nokia recommends that multiple entries be given *entry-ids* in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

Values 1 to 128

regular-expression

Specifies the AS path group regular expression. Allowed values are any string up to 255 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

An AS path in a BGP route matches an AS path group, if the pattern of the path matches the concatenation of all regular expressions in the group. A regular expression incorporates terms and operators that use the terms. An individual AS number is an elementary term in the AS path regular expression. More complex terms can be built from elementary terms. The following are key operators supported by SR OS:

- .

- *
- ?
- {n}
- {m,n}
- {m, }

To reverse the match criteria when specifying a list of ranges or single values using square brackets, use the non-match operator (^) before the elements within the square brackets.

Platforms

7705 SAR Gen 2

entry

Syntax

entry *entry-id*

no entry

Context

[Tree] (config>router>policy-options>policy-statement entry)

Full Context

configure router policy-options policy-statement entry

Description

This command creates the context to edit route policy entries within the route policy statement.

Multiple entries can be created using unique entries. The router exits the filter when the first match is found and executes the action specified. For this reason, entries must be sequenced correctly from most to least explicit.

An entry does not require matching criteria defined (in which case, everything matches) but must have at least define an action in order to be considered complete. Entries without an action are considered incomplete and will be rendered inactive.

The **no** form of this command removes the specified entry from the route policy statement.

Parameters

entry-id

Specifies the entry ID expressed as a decimal integer. An *entry-id* uniquely identifies match criteria and the corresponding action. It is recommended that multiple entries be given *entry-ids* in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

Values 1 to 128

Platforms
7705 SAR Gen 2

9.66 environment

environment

Syntax
environment

Context
[Tree] (environment)

Full Context
environment

Description
Commands in this context configure classic CLI session environment parameters.

Platforms
7705 SAR Gen 2

environment

Syntax
environment

Context
[Tree] (config>system>management-interface>cli>md-cli environment)

Full Context
configure system management-interface cli md-cli environment

Description
Commands in this context configure MD-CLI session environment parameters.

Platforms
7705 SAR Gen 2

9.67 epipe

epipe

Syntax

epipe *service-id* **customer** *customer-id* [**vpn** *vpn-id*] [**vc-switching**] [**create**] **name** [*name*] [**flexible-cross-connect**]

epipe *service-id* [**test**] [**create**] [**name** *name*] [**flexible-cross-connect**]

no epipe *service-id*

Context

[Tree] (config>service epipe)

Full Context

configure service epipe

Description

This command configures an Epipe service instance. This command is used to configure a point-to-point epipe service. An Epipe connects two endpoints defined as Service Access Points (SAPs). Both SAPs may be defined in one 7705 SAR Gen 2 or they may be defined in separate devices connected over the service provider network. When the endpoint SAPs are separated by the service provider network, the far end SAP is generalized into a service destination point (SDP). This SDP describes a destination and the encapsulation method used to reach it.

No MAC learning or filtering is provided on an Epipe.

When creating a service, you must enter the **customer** keyword and specify a *customer-id* to associate the service with a customer. The *customer-id* must already exist, having been created using the **customer** command in the **service** context. After a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and re-created with a new customer association.

After a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

By default, no epipe services exist until they are explicitly created with this command.

The **no** form of this command deletes the epipe service instance with the specified *service-id*. The service cannot be deleted until the service has been shut down.

Parameters

service-id

The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7705 SAR Gen 2 on which this service is defined.

Values *service-id*: 1 to 2147483647

svc-name: up to 64 characters

customer-id

Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn vpn-id

Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN ID. If this parameter is not specified, the VPN ID uses the same service ID number.

Values 1 to 2147483647

Default null (0)

vc-switching

Specifies if the pseudowire switching signaling is used for the spoke SDPs configured in this service.

test

Specifies a unique test service type for the service context which will contain only a SAP configuration. The test service can be used to test the throughput and performance of a path for MPLS-TP PWs.

create

Keyword used to create the service instance. The **create** keyword requirement can be enabled/disabled in the **environment**>**create** context.

name name

Configures an optional service name identifier, up to 64 characters, to a given service. This service name can then be used in configuration references, display, and show commands throughout the system. A defined service name can help the service provider or administrator to identify and manage services within the SR OS platforms.

To create a service, you must assign a service ID; however, after it is created, either the service ID or the service name can be used to identify and reference a service.

If a name is not specified at creation time, then SR OS assigns a string version of the *service-id* as the name.

Values *name*: up to 64 characters

flexible-cross-connect

Keyword to specify the Flexible Cross Connect (FXC) mode, which allows the configuration of two or more SAPs on the same Epipe.

Platforms

7705 SAR Gen 2

9.68 error

error

Syntax

[**no**] **error** [**neighbor** *ip-int-name* | *ip-address*]

Context

[Tree] (debug>router>rip error)

Full Context

debug router rip error

Description

This command enables debugging for RIP errors.

Parameters

ip-int-name* | *ip-address

Debugs the RIP errors sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

error

Syntax

[**no**] **error** [**neighbor** *ip-int-name* | *ipv6-address*]

Context

[Tree] (debug>router>ripng error)

Full Context

debug router ripng error

Description

This command enables debugging for RIPng errors.

Parameters

ip-int-name*| *ipv6-address

Debugs the RIPng errors sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

error

Syntax

[**no**] **error**

Context

[Tree] (debug>router>pcep>pcc>conn error)

[Tree] (debug>router>pcep>pcc error)

Full Context

debug router pcep pcc connection error

debug router pcep pcc error

Description

This command enables debugging for PCC or connection errors.

The **no** form of this command disables debugging.

Platforms

7705 SAR Gen 2

9.69 error-handling

error-handling

Syntax

error-handling

Context

[Tree] (config>service>vprn>bgp>group>neighbor error-handling)

[Tree] (config>service>vprn>bgp error-handling)

[Tree] (config>service>vprn>bgp>group error-handling)

Full Context

configure service vprn bgp group neighbor error-handling

configure service vprn bgp error-handling

configure service vprn bgp group error-handling

Description

This command specifies whether the error handling mechanism for optional transitive path attributes is enabled for this peer group.

Platforms

7705 SAR Gen 2

error-handling

Syntax

error-handling

Context

[Tree] (config>router>bgp>group>neighbor error-handling)

[Tree] (config>router>bgp>group error-handling)

[Tree] (config>router>bgp error-handling)

Full Context

configure router bgp group neighbor error-handling

configure router bgp group error-handling

configure router bgp error-handling

Description

This command specifies whether updated BGP error handling procedures should be applied.

Platforms

7705 SAR Gen 2

9.70 error-report

error-report

Syntax

[**no**] **error-report**

Context

[Tree] (debug>router>rpki-session>packet error-report)

Full Context

debug router rpki-session packet error-report

Description

This command enables debugging for error report RPKI packets.

The **no** form of this command disables debugging for error report RPKI packets.

Platforms

7705 SAR Gen 2

9.71 esp-auth-algorithm

esp-auth-algorithm

Syntax

esp-auth-algorithm {**null** | **md5** | **sha1** | **sha256** | **sha384** | **sha512** | **aes-xcbc** | **auth-encryption**}

no esp-auth-algorithm

Context

[Tree] (config>ipsec>transform esp-auth-algorithm)

Full Context

configure ipsec ipsec-transform esp-auth-algorithm

Description

This command specifies which hashing algorithm should be used for the authentication function Encapsulating Security Payload (ESP). Both ends of a manually configured tunnel must share the same configuration parameters for the IPsec tunnel to enter the operational state.

The **no** form of this command disables the authentication.

Default

esp-auth-algorithm sha1

Parameters

null

This is a very fast algorithm specified in RFC 2410, which provides no authentication.

md5

This parameter configures ESP to use the **hmac-md5** algorithm for authentication.

sha1

This parameter configures ESP to use the **hmac-sha1** algorithm for authentication.

sha256

This parameter configures ESP to use the **sha256** algorithm for authentication.

sha384

This parameter configures ESP to use the **sha384** algorithm for authentication.

sha512

This parameter configures ESP to use the **sha512** algorithm for authentication.

aes-xcbc

Specifies the **aes-xcbc** algorithm for authentication.

auth-encryption

This parameter must be configured when **esp-encryption-algorithm** is either **aes-gcm** or **aes-gmac**.

Platforms

7705 SAR Gen 2

esp-auth-algorithm

Syntax

esp-auth-algorithm {**sha256** | **sha512**}

no esp-auth-algorithm

Context

[Tree] (config>grp-encryp>encryp-keygrp esp-auth-algorithm)

Full Context

configure group-encryption encryption-keygroup esp-auth-algorithm

Description

This command specifies the hashing algorithm used to perform authentication on the Encapsulating Security Payload (ESP) within NGE packets for services configured using this key group. All SPI entries must be deleted before the **no** form of the command may be entered or the **esp-auth-algorithm** value changed from its current value.

The **no** form of the command reverts to the default value.

Default

esp-auth-algorithm sha256

Parameters

sha256

Configures the ESP to use the HMAC-SHA-256 algorithm for authentication.

sha512

Configures the ESP to use the HMAC-SHA-512 algorithm for authentication.

Platforms

7705 SAR Gen 2

9.72 esp-encryption-algorithm

esp-encryption-algorithm

Syntax

esp-encryption-algorithm {**null** | **des** | **3des** | **aes128** | **aes192** | **aes256**| **aes128-gcm8** | **aes128-gcm12** | **aes128-gcm16** | **aes192-gcm8** | **aes192-gcm12** | **aes192-gcm16** | **aes256-gcm8** | **aes256-gcm12** | **aes256-gcm16** | **null-aes128-gmac** | **null-aes192-gmac** | **null-aes256-gmac**}

no esp-encryption-algorithm

Context

[Tree] (config>ipsec>ipsec-transform esp-encryption-algorithm)

Full Context

configure ipsec ipsec-transform esp-encryption-algorithm

Description

This command specifies the encryption algorithm to use for the IPsec session. Encryption only applies to esp configurations. If encryption is not defined, esp will not be used.

For IPsec tunnels to come up, both ends need to be configured with the same encryption algorithm.

The **no** form of this command removes the specified encryption algorithm.

Note:
When **aes-gcm** or **aes-gmac** is configured:

- **esp-auth-algorithm** must be set to **auth-encryption**
- the system will not include the authentication algorithm in the ESP proposal of the SA payload
- **ipsec-transform** cannot be used for manual keying

Default

esp-encryption-algorithm aes128

Parameters

null

This parameter configures the high-speed null algorithm, which does nothing. This is the same as not having encryption turned on.

des

This parameter configures the 56-bit des algorithm for encryption. This is an older algorithm, with relatively weak security. Although slightly better than no encryption,

it should only be used where a strong algorithm is not available on both ends at an acceptable performance level.

3des

This parameter configures the 3-des algorithm for encryption. This is a modified application of the des algorithm which uses multiple des operations to make things more secure.

aes128

This parameter configures the aes algorithm with a block size of 128 bits. This is the mandatory implementation size for aes. As of today, this is a very strong algorithm choice.

aes192

This parameter configures the aes algorithm with a block size of 192 bits. This is a stronger version of aes.

aes256

This parameter configures the aes algorithm with a block size of 256 bits. This is the strongest available version of aes.

aes128-gcm8

Configures ESP to use aes-gcm with a 128-bit key size and an 8-byte ICV for encryption and authentication.

aes128-gcm12

Configures ESP to use aes-gcm with a 128-bit key size and a 12-byte ICV for encryption and authentication.

aes128-gcm16

Configures ESP to use aes-gcm with a 128-bit key size and a 16-byte ICV for encryption and authentication.

aes192-gcm8

Configures ESP to use aes-gcm with a 192-bit key size and an 8-byte ICV for encryption and authentication.

aes192-gcm12

Configures ESP to use aes-gcm with a 192-bit key size and a 12-byte ICV for encryption and authentication.

aes192-gcm16

Configures ESP to use aes-gcm with a 192-bit key size and a 16-byte ICV for encryption and authentication.

aes256-gcm8

Configures ESP to use aes-gcm with a 256-bit key size and an 8-byte ICV for encryption and authentication.

aes256-gcm12

Configures ESP to use aes-gcm with a 256-bit key size and a 12-byte ICV for encryption and authentication.

aes128-gcm16

Configures ESP to use aes-gcm with a 256-bit key size and a 16-byte ICV for encryption and authentication.

null-aes128gmac

Configures ESP to use aes-gmac with a 128-bit key size for authentication only.

null-aes192gmac

Configures ESP to use aes-gmac with a 192-bit key size for authentication only.

null-aes256gmac

Configures ESP to use aes-gmac with a 256-bit key size for authentication only.

Platforms

7705 SAR Gen 2

esp-encryption-algorithm

Syntax

esp-encryption-algorithm {**aes128** | **aes256**}

no esp-encryption-algorithm

Context

[Tree] (config>grp-encryp>encryp-keygrp esp-encryption-algorithm)

Full Context

configure group-encryption encryption-keygroup esp-encryption-algorithm

Description

This command specifies the encryption algorithm used to perform encryption on the Encapsulating Security Payload (ESP) within NGE packets for services configured using this key group. All SPI entries must be deleted before the **no** form of the command may be entered or the **esp-encryption-algorithm** value changed from its current value.

The **no** form of the command resets the parameter to the default value.

Default

esp-encryption-algorithm aes128

Parameters

aes128

Configures the AES algorithm with a block size of 128 bits—a very strong algorithm choice.

aes256

Configures the AES algorithm with a block size of 256 bits—the strongest available version of AES.

Platforms

7705 SAR Gen 2

9.73 esp-ext-hdr

esp-ext-hdr

Syntax

esp-ext-hdr {**true** | **false**}

no esp-ext-hdr

Context

[Tree] (config>filter>ipv6-filter>entry>match esp-ext-hdr)

Full Context

configure filter ipv6-filter entry match esp-ext-hdr

Description

This command enables match on existence of ESP Extension Header in the IPv6 filter policy.

The **no** form of this command ignores ESP Extension Header presence/absence in a packet when evaluating match criteria of a given filter policy entry.

Default

no esp-ext-hdr

Parameters

true

Matches a packet with an ESP Extension Header.

false

Matches a packet without an ESP Extension Header.

Platforms

7705 SAR Gen 2

9.74 est

est

Syntax

est

Context

[Tree] (admin>certificate est)

Full Context

admin certificate est

Description

Commands in this context configure Enrollment over Secure Transport (EST) parameters.

Platforms

7705 SAR Gen 2

9.75 eth-tag

eth-tag

Syntax

eth-tag *tag-value*

no eth-tag

Context

[Tree] (config>service>epipe>bgp-evpn>remote-attachment-circuit eth-tag)

[Tree] (config>service>epipe>bgp-evpn>local-attachment-circuit eth-tag)

Full Context

configure service epipe bgp-evpn remote-attachment-circuit eth-tag

configure service epipe bgp-evpn local-attachment-circuit eth-tag

Description

This command configures the Ethernet tag value. When configured in the **local-attachment-circuit** context, the system uses the value in the advertised AD per-EVI route sent for the attachment circuit. When configured in the **remote-attachment-circuit** context the system compares that value with the eth-tag value of the imported AD per-EVI routes for the service. If there is a match, the system creates an EVPN destination for the Epipe.

Parameters

tag-value

Specifies the Ethernet tag value of the attachment circuit.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

9.76 ethernet

ethernet

Syntax

ethernet

Context

[Tree] (config>port ethernet)

Full Context

configure port ethernet

Description

This command the context to configure Ethernet port attributes.

This context can only be used when configuring Fast Ethernet, gigabit or 10-G Fast Ethernet or Ethernet LAN ports on an appropriate MDA.

Platforms

7705 SAR Gen 2

9.77 ethernet-ctag

ethernet-ctag

Syntax

[**no**] **ethernet-ctag**

Context

[Tree] (config>qos>sap-egress ethernet-ctag)

Full Context

configure qos sap-egress ethernet-ctag

Description

This command specifies that the top customer tag should be used for egress reclassification based on dot1p criteria. This command applies to all dot1p criteria configured in a given SAP egress QoS policy.

The **no** form of this command means that a service delimiting tag will be used for egress reclassification based on dot1p criteria.

Default

no ethernet-ctag

Platforms

7705 SAR Gen 2

9.78 etype

etype

Syntax

etype *etype-value*

no etype

Context

[Tree] (config>qos>sap-ingress>mac-criteria>entry>match etype)

Full Context

configure qos sap-ingress mac-criteria entry match etype

Description

Configures an Ethernet type II value to be used as a service ingress QoS policy match criterion.

The Ethernet type field is a 2-byte field used to identify the protocol carried by the Ethernet frame. For example, 0800 is used to identify the IPv4 packets.

The Ethernet type field is used by the Ethernet version-II frames. IEEE 802.3 Ethernet frames do not use the type field. For IEEE 802.3 frames, use the dsap, ssap, or snap-pid fields as match criteria.

The snap-pid field, etype field, ssap, and dsap fields are mutually exclusive and cannot be part of the same match criteria.

The no form of this command removes the previously entered etype field as the match criteria.

Default

no etype

Parameters

etype-value

The Ethernet type II frame Ethertype value to be used as a match criterion expressed in hexadecimal.

Values 0x0600 to 0xFFFF

Platforms

7705 SAR Gen 2

etype

Syntax

etype *0x0600xx0xffff*

no etype

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry>match etype)

Full Context

configure system security management-access-filter mac-filter entry match etype

Description

Configures an Ethernet type II Ethertype value to be used as a MAC filter match criterion.

The Ethernet type field is a two-byte field used to identify the protocol carried by the Ethernet frame. For example, 0800 is used to identify the IPv4 packets.

The Ethernet type field is used by the Ethernet version-II frames. IEEE 802.3 Ethernet frames do not use the type field. For IEEE 802.3 frames, use the dsap, ssap or snap-pid fields as match criteria.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and may not be part of the same match criteria. Refer to the *7705 SAR Gen 2 Router Configuration Guide* for information about MAC Match Criteria Exclusivity Rules fields that are exclusive based on the frame format.

The **no** form of this command removes the previously entered etype field as the match criteria.

Default

no etype

Parameters

ethernet-type

Specifies the Ethernet type II frame Ethertype value to be used as a match criterion expressed in hexadecimal.

Values 0x0600 to 0xFFFF

Platforms

7705 SAR Gen 2

9.79 event

event

Syntax

event *event-type* [**create**]

no event *event-type*

Context

[Tree] (config>card>mda event)

Full Context

configure card mda event

Description

This command allows the user to control the action to be taken when a specific hardware error event is raised against the target MDA.

If no event action has been created for a specific event type, then the hardware errors related to that event type are ignored by the management plane of the router.

The log event raised for any event type (for example, soft-error, memory-error) is tmnxEqHwEventDetected.

The **no** form of this command clears any action defined for the event.

Parameters

event-type

Specifies the event type, up to 32 characters.

Values soft-error — Defines the action to be taken when soft errors are detected on the MDA

internal-frame-loss — System detected frame loss in the traffic transiting the MDA.

memory-error — Provides the user options to handle MDA memory error events on MDAs. This feature is supported on FP2- and FP3-based Ethernet MDAs and IMMs.

data-link-error — Provides the user options to handle datapath link errors on an MDA.

create

Keyword used to create an event.

Platforms

7705 SAR Gen 2

event

Syntax

[**no**] **event**

Context

[Tree] (debug>router>ldp>if event)

[Tree] (debug>router>ldp>peer event)

Full Context

debug router ldp interface event

debug router ldp peer event

Description

This command configures debugging for specific LDP events.

Platforms

7705 SAR Gen 2

event

Syntax

[**no**] **event**

Context

[Tree] (debug>router>rsvp event)

[Tree] (debug>router>mpls event)

Full Context

debug router rsvp event

debug router mpls event

Description

This command enables debugging for specific events.

The **no** form of the command disables the debugging.

Platforms

7705 SAR Gen 2

event

Syntax

[**no**] **event**

Context

[Tree] (debug>router>ip event)

Full Context

debug router ip event

Description

This command enables debugging for specific IP events.

The **no** form of this command disables debugging for the specified IP events.

Platforms

7705 SAR Gen 2

event

Syntax

event *rmon-event-id* [*event-type*] [**description** *description-string*] [**owner** *owner-string*]

no event *rmon-event-id*

Context

[Tree] (config>system>thresholds>rmon event)

Full Context

configure system thresholds rmon event

Description

The event command configures an entry in the RMON-MIB event table. The event command controls the generation and notification of threshold crossing events configured with the alarm command. When a threshold crossing event is triggered, the **rmon>event** configuration optionally specifies if an entry in the RMON-MIB log table should be created to record the occurrence of the event. It may also specify that an SNMP notification (trap) should be generated for the event. The RMON-MIB defines two notifications for threshold crossing events: Rising Alarm and Falling Alarm.

Creating an event entry in the RMON-MIB log table does not create a corresponding entry in the SR OS event logs. However, when the **event-type** is set to trap, the generation of a Rising Alarm or Falling Alarm notification creates an entry in the SR OS event logs and that is distributed to all the SR OS log destinations that are configured: CONSOLE, session, memory, file, syslog, or SNMP trap destination.

The SR OS logger message includes a rising or falling threshold crossing event indicator, the sample type (absolute or delta), the sampled value, the threshold value, the RMON-alarm-id, the associated RMON-event-id and the sampled SNMP object identifier.

Use the **no** form of this command to remove an rmon-event-id from the configuration.

Parameters

rmon-event-id

Specifies an identifier for this event. Alarm ID values above 65400 are used for dynamic system threshold commands and should be avoided.

Values 1 to 65535

rmon-event-type

Specifies the type of notification action to be taken when this event occurs.

Values **log** — An entry is made in the RMON-MIB log table for each event occurrence.

This does **not** create an SR OS logger entry. The RMON-MIB log table entries can be viewed using the **show>system>thresholds** CLI command.

trap — An SR OS logger event is generated. The SR OS logger utility then distributes the notification of this event to its configured log destinations which may be CONSOLE, telnet session, memory log, cflash file, syslog, or SNMP trap destinations logs.

both — Both an entry in the RMON-MIB logTable and an SR OS logger event are generated.

none — No action is taken.

Default both

description-string

Specifies a user configurable string that can be used to identify the purpose of this event. This is an optional parameter and can be up to 80 characters long. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

owner-string

Specifies the owner string; the owner identifies the creator of this alarm. It defaults to "TiMOS CLI". This parameter is defined primarily to allow entries that have been created in the RMON-MIB alarmTable by remote SNMP managers to be saved and reloaded in a CLI configuration file. The owner will not normally be configured by CLI users and can be up 80 characters long.

Default TiMOS CLI

Configuration example:

```
event 5 rmon-event-type both description "alarm testing" owner "TiMOS CLI"
```

Platforms

7705 SAR Gen 2

event

Syntax

[**no**] **event** *application-id event-name-id*

Context

[Tree] (config>log>event-trigger event)

Full Context

configure log event-trigger event

Description

This command configures a specific log event as a trigger for one or more EHS handlers. Further matching criteria can be applied to only trigger certain handlers with certain instances of the log event.

The **no** form of this command removes the specified trigger event.

Parameters

application-id

Specifies the type of application that triggers the event.

Values adp, application_assurance, auto_prov, bfd, bgp, bier, bmp, calltrace, cflowd, chassis, cpmhwfilter, cpmhwqueue, debug, dhcp, dhcps, diameter, dot1x, dynsvc, efm_oam, elmi, ering, eth_cfm, etun, filter, fpe, gsmp, gtp, igmp, igmp_snooping, ip, ipfix, ipsec, ipsec_cpm, isis, l2tp, lag, ldap, ldp, li, lldp, logger, maffilter, macsec, mcac, mcpath, mc_redundancy, mgmt_core, mirror, mld, mld_snooping, mpls, mpls_tp, mpls_lmgr, mrp, msdp, nat, nge, ntp, oam, open_flow, ospf, pcap, pcep, pfcp, pim, pim_snooping, port, pppoe, pppoe_clnt, profile, ptp, pxc, python, qos, radius, rib_api, rip, rip_ng, route_next_hop, route_policy, rpki, rsvp, satellite, security, sflow, snmp, sr_mpls, sr_policy, srv6, stp, subscr_mgmt, sub_host_trk, svcmgr, system, telemetry, tip, tls, tree_sid, user, user_db, video, vrrp, vrtr, wlan_gw, wpp

event-name-id

Specifies the name or numerical identifier of the event.

Values 0 to 4294967295 | *event-name*: 32 characters max

Platforms

7705 SAR Gen 2

9.80 event-control

event-control

Syntax

event-control *application-id* [*event-name* | *event-number*] [**generate**] [*severity-level*] [**throttle**] [**specific-throttle-rate** *events-limit* **interval** *seconds* | **disable-specific-throttle**] [**repeat** | **no-repeat**]

event-control *application-id* [*event-name* | *event-number*] **suppress**

no event-control *application-id* [*event-name* | *event-number*]

Context

[Tree] (config>log event-control)

Full Context

configure log event-control

Description

This command is used to specify that a particular event or all events associated with an application is either generated or suppressed.

Events are generated by an application and contain an event number and description explaining the cause of the event. Each event has a default designation which directs it to be generated or suppressed.

Events are generated with a default severity level that can be modified by using the *severity-level* option.

Events that are suppressed by default are typically used for debugging purposes. Events are suppressed at the time the application requests the event's generation. No event log entry is generated regardless of the destination. While this feature can save processor resources, there may be a negative effect on the ability to troubleshoot problems if the logging entries are squelched. In reverse, indiscriminate application may cause excessive overhead.

The rate of event generation can be throttled by using the **throttle** parameter.

The **no** form of this command reverts the parameters to the default setting for events for the application or a specific event within the application. The severity, generate, suppress, and throttle options will also be reset to the initial values.

Default

Each event has a set of default settings. To display a list of all events and the current configuration use the **event-control** command.

Parameters

application-id

The application whose events are affected by this event control filter.

Values A valid application name. Use the **show log applications** command to display a list of valid application names. Examples of valid applications are **bgp**, **chassis**, **efm_oam**, **filter**, **security**, **system**, and **vrrp**.

event-name

To generate, suppress, or revert to default for a single event, enter the specific event short name up to 32 characters. If no event name is specified, the command applies to all events in the application. To display a list of all event short names use the **event-control** command.

event-number

To generate, suppress, or revert to default for a single event, enter the specific number. If no event number is specified, the command applies to all events in the application.

Values 0 to 4294967295

generate

Specifies that logger event is created when this event occurs. The generate keyword can be used with two optional parameters, *severity-level* and **throttle**.

Default generate

severity-level

An ASCII string representing the severity level to associate with the specified generated events

Default The system-assigned severity name

Values cleared, indeterminate, critical, major, minor, warning

throttle

Specifies whether or not events of this type will be throttled. By default, event throttling is on for most event types.

suppress

This keyword indicates that the specified events will not be logged. If the **suppress** keyword is not specified then the events are generated by default. For example on the 7705 SAR Gen 2, **event-control bgp suppress** will suppress all BGP events. If a log event is a raising event for a Facility Alarm, and the associated Facility Alarm is raised, then changing the log event to **suppress** clears the associated Facility Alarm.

Default generate

specific-throttle-rate *events-limit*

The log event throttling rate can be configured independently for each log event using this keyword. This specific-throttle-rate overrides the globally configured throttle rate (**config**>**log**>**throttle-rate**) for the specific log event.

Values 1 to 20000

interval *seconds*

Specifies the number of seconds that the specific throttling intervals lasts.

Values 1 to 1200

disable-specific-throttle

Specifies to disable the **specific-throttle-rate**.

repeat

Specifies that the log event should be repeated every minute until the underlying condition is cleared. Only supported for the following log events: BGP tBgpMaxNgPfxLmtThresholdReached and PORT tmnxEqPortEtherCrcAlarm (for **degrade** threshold only)

Platforms

7705 SAR Gen 2

9.81 event-damping

event-damping

Syntax

[**no**] **event-damping**

Context

[Tree] (config>log event-damping)

Full Context

configure log event-damping

Description

This command allows the user to set the event damping algorithm to suppress QoS or filter change events.

The **no** form of this command removes the event damping algorithm.

Note:
While this event damping is original behavior for some modules such as service manager, QoS, and filters, it can result in the NMS system database being out of sync because of missed change events. On the other hand, if the damping is disabled (**no event-damping**), it may take much longer to **exec** a large CLI configuration file after system bootup.

Platforms

7705 SAR Gen 2

9.82 event-handler

event-handler

Syntax

event-handler *event-handler*

no event-handler

Context

[Tree] (config>log>event-trigger>event>trigger-entry event-handler)

Full Context

configure log event-trigger event trigger-entry event-handler

Description

This command configures the event handler to be used for this trigger entry.

The **no** form of this command removes the event handler configuration.

Parameters

event-handler

Specifies the name of the event handler, up to 32 characters.

Platforms

7705 SAR Gen 2

event-handler

Syntax

event-handler

Context

[Tree] (config>system>security>cli-script>authorization event-handler)

Full Context

configure system security cli-script authorization event-handler

Description

Commands in this context configure authorization for the Event Handling System (EHS). EHS allows user-controlled programmatic exception handling by allowing a CLI script to be executed upon the detection of a log event.

Platforms
7705 SAR Gen 2

9.83 event-handling

event-handling

Syntax
event-handling

Context
[Tree] (config>log event-handling)

Full Context
configure log event-handling

Description
Commands in this context configure event handling within the Event Handler System (EHS).

Platforms
7705 SAR Gen 2

9.84 event-mon

event-mon

Syntax
event-mon

Context
[Tree] (config>oam-pm>session>meas-intvl event-mon)

Full Context
configure oam-pm session meas-interval event-mon

Description
This command enables the different threshold events on a specific measurement interval. Only one measurement interval with a configured OAM PM session can have events enabled using the **no shutdown** command.

Platforms
7705 SAR Gen 2

9.85 event-trigger

event-trigger

Syntax
event-trigger

Context
[Tree] (config>log event-trigger)

Full Context
configure log event-trigger

Description
Commands in this context configure log events as triggers for Event Handling System (EHS) handlers.

Platforms
7705 SAR Gen 2

9.86 event-type

event-type

Syntax
[**no**] **event-type** {**arp** | **config-change** | **oper-status-change** | **neighbor-discovery**}

Context
[Tree] (debug>service>id>sap event-type)

Full Context
debug service id sap event-type

Description
This command enables a particular debugging event type.

The **no** form of this command disables the event type debugging.

Parameters

arp

Displays ARP events.

config-change

Debugs configuration change events.

oper-status-change

Debugs service operational status changes.

neighbor-discovery

Displays the status of IPv6 neighbor discovery for the sap or the spoke-sdp.

Platforms

7705 SAR Gen 2

Output

The following output is an example of event-type information.

Output Example

```
A:bksim180# debug service id 1000 sap 1/7/1 event-type arp
DEBUG OUTPUT show on CLI is as follows:
3 2008/11/17 18:13:24.35 UTC MINOR: DEBUG #2001 Base Service 1000 SAP
1/7/1 "Service 1000 SAP 1/7/1:
RX: ARP_REQUEST (0x0001)
hwType     : 0x0001
prType     : 0x0800
hwLength   : 0x06
prLength   : 0x04
srcMac     : 8c:c7:01:07:00:03
destMac    : 00:00:00:00:00:00
srcIp      : 10.1.1.2
destIp     : 10.1.1.1
"

4 2008/11/17 18:13:24.35 UTC MINOR: DEBUG #2001 Base Service 1000
SAP 1/7/1 "Service 1000 SAP 1/7/1:
TX: ARP_RESPONSE (0x0002)
hwType     : 0x0001
prType     : 0x0800
hwLength   : 0x06
prLength   : 0x04
srcMac     : 00:03:0a:0a:0a:0a
destMac    : 8c:c7:01:07:00:03
srcIp      : 10.1.1.1
destIp     : 10.1.1.2
"
```

event-type

Syntax

[**no**] **event-type** {**config-change** | **oper-status-change** | **neighbor-discovery** | **control-channel-status**}

Context

[Tree] (debug>service>id>sdp event-type)

Full Context

debug service id sdp event-type

Description

This command enables a particular debugging event type.

The **no** form of this command disables the event type debugging.

Parameters

config-change

Debugs configuration change events.

oper-status-change

Debugs service operational status changes.

neighbor-discovery

Displays the status of IPv6 neighbor discovery for the sap or the spoke-sdp.

control-channel-status

Debugs control channel status events.

Platforms

7705 SAR Gen 2

event-type

Syntax

[**no**] **event-type** {**config-change** | **svc-oper-status-change** | **sap-oper-status-change** | **sdpbind-oper-status-change**}

Context

[Tree] (debug>service>id event-type)

Full Context

debug service id event-type

Description

This command enables a particular debugging event type. The **no** form of this command disables the event type debugging.

Parameters

config-change

Debugs configuration change events

svc-oper-status-change
Debugs service operational status changes

sap-oper-status-change
Debugs SAP operational status changes

sdpbind-oper-status-change
Debugs SDP operational status changes

Platforms
7705 SAR Gen 2

9.87 events

events

Syntax
[**no**] **events**

[**no**] **events interface** *ip-int-name* [**vrid** *virtual-router-id*]

[**no**] **events interface** *ip-int-name* **vrid** *virtual-router-id* **ipv6**

Context
[Tree] (debug>router>vrrp events)

Full Context
debug router vrrp events

Description
This command enables debugging for VRRP events.

The **no** form of the command disables debugging.

Parameters
ip-int-name
Displays the specified interface name.

virtual-router-id
Displays the specified VRID.

ipv6
Debugs the specified IPv6 VRRP interface.

Platforms
7705 SAR Gen 2

events

Syntax

events [**neighbor** *ip-address* | **group** *name*]

no events

Context

[Tree] (debug>router>bgp events)

Full Context

debug router bgp events

Description

This command logs all events changing the state of a BGP peer.

The **no** form of this command disables the debugging.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x [-interface] (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H
- d: [0 to 255]D
- interface: up to 32 characters for link local addresses

group *name*

Debugs only events affecting the specified peer group name, up to 64 characters, and associated neighbors.

Platforms

7705 SAR Gen 2

events

Syntax

[**no**] **events** [**neighbor** *ip-int-name* | *ip-addr*]

Context

[Tree] (debug>router>rip events)

Full Context

debug router rip events

Description

This command enables debugging for RIP events.

Parameters

ip-int-name* | *ip-address

Debugs the RIP events sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

events

Syntax

[**no**] **events** [**neighbor** *ip-int-name*]

Context

[Tree] (debug>router>ripng events)

Full Context

debug router ripng events

Description

This command enables debugging for RIPng events.

Parameters

ip-int-name

Debugs the RIPng events sent on the neighbor IP interface.

Platforms

7705 SAR Gen 2

9.88 evi

evi

Syntax

evi *value*

no evi

Context

[Tree] (config>service>vpls>bgp-evpn evi)

[Tree] (config>service>epipe>bgp-evpn evi)

Full Context

configure service vpls bgp-evpn evi

configure service epipe bgp-evpn evi

Description

This command allows the configuration of a 2-byte EVPN instance (EVI) unique in the system. It is used for the service-carving algorithm for multi-homing and auto-deriving route target and route distinguishers.

If not specified, the value is zero and no route distinguisher or route targets are auto-derived from it. If the *evi* value is specified and no other **route-distinguisher** or **route-target** is configured in the service, the following rules apply:

- the route distinguisher is derived from <system_ip>:evi
- the route target is derived from <autonomous-system>:evi

If VSI import and export policies are configured, the route target must be configured in the policies and those values take preference over the auto-derived route targets. If **bgp-ad>vpls-id** and **bgp-evpn>evi** are both configured on the same service, the VPLS ID auto-derived route target or route distinguisher takes precedence over the values auto-derived from the EVI. The operational route target for a service is displayed in the **show service id bgp** command.

The **no** form of this command sets the EVI value back to zero.

Parameters

value

Specifies the EVPN instance.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

evi

Syntax

evi *value*

no evi

Context

[Tree] (config>service>vprn>bgp-evpn>mpls evi)

Full Context

configure service vprn bgp-evpn mpls evi

Description

This command configures a 2-byte EVPN instance (EVI) unique in the system.

The router uses the EVI to identify the BGP EVPN instance in a VPRN (for the EVPN-IFL model) or an R-VPLS (for the EVPN-IFF model) that is associated with the Layer 3 Ethernet Segment (ES), for the purpose of IP Aliasing. This configuration is required on the PEs attached to the ES as well as on the remote PEs that need to create ES destinations to the multihoming Layer 3 ES.

The **no** form of this command removes the EVI value.

Default

no evi

Parameters

value

Specifies the EVPN instance.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

9.89 evi-three-byte-auto-rt

evi-three-byte-auto-rt

Syntax

[**no**] **evi-three-byte-auto-rt**

Context

[Tree] (config>service>vpls>bgp-evpn>mpls evi-three-byte-auto-rt)

[Tree] (config>service>epipe>bgp-evpn>mpls evi-three-byte-auto-rt)

Full Context

configure service vpls bgp-evpn mpls evi-three-byte-auto-rt

configure service epipe bgp-evpn mpls evi-three-byte-auto-rt

Description

This command specifies that the BGP-EVPN instance import and export route target is auto-derived as described in RFC 8365 (Global-Administrator:A/Type/D-ID/Service-ID).

Where:

- Global Administrator — is the configured 2-octet AS Number. If the configured ASN exceeds the 2 byte limit, the low order 16-bit value will be taken.
- A=0 (for auto-derivation)
- Type=4 (EVI-based route target)
- D-ID= [1..2] — encodes the BGP instance. This allows the auto-derivation of different route targets in multi-instance services. The value is inherited from the corresponding BGP instance.
- Service ID= 3-octet EVI

The **no** form of this command disallows the derivation of the route target.

Default

no evi-three-byte-auto-rt

Platforms

7705 SAR Gen 2

9.90 evpn

evpn

Syntax

evpn send *send-limit*

evpn send *send-limit* **receive** [**none**]

no evpn

Context

[Tree] (config>router>bgp>group>add-paths evpn)

[Tree] (config>router>bgp>add-paths evpn)

[Tree] (config>router>bgp>group>neighbor>add-paths evpn)

Full Context

configure router bgp group add-paths evpn

configure router bgp add-paths evpn

configure router bgp group neighbor add-paths evpn

Description

This command configures the Add-Paths capability for EVPN routes.

The **no** form of this command disables Add-Paths support for EVPN routes. This causes sessions that are established using Add-Paths for EVPN to go down and come back up without the Add-Paths capability.

Default

no evpn

Parameters

send-limit

Specifies the maximum number of EVPN paths to send.

Values 1 to 16, none, multipaths

receive

Keyword used to allow multiple EVPN paths per prefix from a peer.

none

Keyword used to specify that the router does not negotiate to receive multiple unlabeled unicast routes per EVPN prefix.

Platforms

7705 SAR Gen 2

evpn

Syntax

evpn

Context

[Tree] (config>service>vprn>if>vpls evpn)

[Tree] (config>service>ies>if>vpls evpn)

Full Context

configure service vprn interface vpls evpn

configure service ies interface vpls evpn

Description

Commands in this context configure EVPN parameters.

Platforms

7705 SAR Gen 2

9.91 evpn-etree-leaf-label

evpn-etree-leaf-label

Syntax

evpn-etree-leaf-label [[*32..524256*]]

no evpn-etree-leaf-label

Context

[Tree] (config>service>system>bgp-evpn evpn-etree-leaf-label)

Full Context

configure service system bgp-evpn evpn-etree-leaf-label

Description

This command enables EVPN Ethernet-Tree (E-Tree) VPLS services on the router (not B-VPLS). It allocates an E-Tree leaf label for the Provider Edge (PE) device and configures the ILM entry.

The command ensures that in-flight traffic can perform an ILM entry lookup at any time, and avoid the discards during **shutdown** or **no shutdown** services (or at least reduce the timing window so that it does not occur during normal operation or configuration).

The E-Tree leaf label can optionally be statically configured with a value. The label value must be in the static label range of the system.

Note:
The **evpn-etree-leaf-label** command must be configured to execute **bgp-evpn mpls no shutdown**.

The **no** form of this command removes the value from the configuration.

Default

no evpn-etree-leaf-label

Parameters

32..524256

Specifies the E-Tree leaf label

Values 32 to 524256

Platforms
7705 SAR Gen 2

9.92 evpn-link-bandwidth

evpn-link-bandwidth

Syntax
evpn-link-bandwidth

Context
[Tree] (config>service>vprn>bgp>group evpn-link-bandwidth)
[Tree] (config>service>vprn>bgp-evpn>mpls evpn-link-bandwidth)
[Tree] (config>service>vprn>bgp>group>neighbor evpn-link-bandwidth)

Full Context
configure service vprn bgp group evpn-link-bandwidth
configure service vprn bgp-evpn mpls evpn-link-bandwidth
configure service vprn bgp group neighbor evpn-link-bandwidth

Description
Commands in these contexts configure the EVPN link bandwidth.

Platforms
7705 SAR Gen 2

9.93 evpn-mpls

evpn-mpls

Syntax
[**no**] **evpn-mpls**

Context
[Tree] (debug>service>id>igmp-snooping evpn-mpls)

Full Context
debug service id igmp-snooping evpn-mpls

Description

This command shows IGMP packets for EVPN-MPLS destinations. The **no** form of this command disables the debugging for EVPN-MPLS destinations

Platforms

7705 SAR Gen 2

9.94 evpn-nd-advertise

evpn-nd-advertise

Syntax

evpn-nd-advertise {**host** | **router** | **router-host**}

Context

[Tree] (config>service>vpls>proxy-nd evpn-nd-advertise)

Full Context

configure service vpls proxy-nd evpn-nd-advertise

Description

This command enables the advertisement of static or dynamic entries that are learned as host, router, or host and router, (only one option is possible in a specified service). It also determines the R flag (host or router) when sending Neighbor Advertisement (NA) messages for existing EVPN entries in the proxy-ND table.

The **router-host** command option is only possible when the ARP/ND extended community is advertised along with the MAC/IP routes. It determines that both host and router (dynamic and static) entries are advertised in MAC/IP routes, with an indication whether the entry is host or router in the R flag. These EVPN entries are installed as host or router entries depending on the R flag of the route, and NA messages for them are sent with the proper host or router indication.

To modify this command you must shutdown the proxy ND.

```
configure service vpls proxy-nd shutdown
```

Default

evpn-nd-advertise router

Parameters

host

Enables the advertisement of static or dynamic entries that are learned as host.

router

Enables the advertisement of static or dynamic entries that are learned as routers.

router-host

Enables the advertisement of static or dynamic entries that are learned as router or host.

Platforms

7705 SAR Gen 2

9.95 evpn-route-tag

evpn-route-tag

Syntax

evpn-route-tag *tag*

no evpn-route-tag

Context

[Tree] (config>service>vpls>proxy-arp evpn-route-tag)

[Tree] (config>service>vpls>proxy-nd evpn-route-tag)

Full Context

configure service vpls proxy-arp evpn-route-tag

configure service vpls proxy-nd evpn-route-tag

Description

This command configures a local route tag that can be used on export policies to match MAC/IP routes generated by the proxy-ARP or proxy-ND module. For example, if a new active dynamic proxy-ARP entry is added to the proxy-ARP table and **evpn-route-tag** is 10, an export policy that matches on tag 10 and adds a site-of-origin community SOO-1, allows the router to advertise the MAC/IP route for the proxy-ARP entry with community SOO-1.

The **no** form of this command removes the route tag for the generated EVPN MAC/IP routes.

Parameters

tag

Specifies the route tag, in either decimal or hexadecimal form.

Values 1 to 255

Platforms

7705 SAR Gen 2

9.96 evpn-tunnel

evpn-tunnel

Syntax

evpn-tunnel [**ipv6-gateway-address** {**ip** | **mac**}] [**supplementary-broadcast-domain**]

no evpn-tunnel

Context

[Tree] (config>service>vprn>if>vpls evpn-tunnel)

Full Context

configure service vprn interface vpls evpn-tunnel

Description

This command sets the evpn-tunnel mode for the attached R-VPLS. When enabled for an IPv4 interface, no IPv4 address is required under the same interface. When enabled on an IPv6 interface, the **ipv6-gateway-address** parameter can be configured as **ip** or **mac**.

When configured as **evpn-tunnel ipv6-gateway-address ip** or simply **evpn-tunnel**, then:

- on transmission, the router populates the GW IP field of the route type 5 with a Link-Local-Address (LLA) if an explicit global IPv6 address is not configured. Otherwise, the configured IPv6 address is used.
- on reception of routes type 5 for IPv6 prefixes, only routes with non-zero GW IP are processed; the rest of the routes will be **treated-as-withdraw**.

When configured as **evpn-tunnel ipv6-gateway-address mac,** then:

- on transmission, the router sends routes type 5 with zero GW IP field, and a MAC extended community of the router, containing the VPRN interface MAC.
- on reception of IPv6 prefix routes, only routes with zero GW IP and non-zero router's MAC are processed; the rest of the routes will be **treated-as-withdraw**.

The **supplementary-broadcast-domain** option instructs the data path to exclude EVPN destinations in the Layer 3 lookup for packets coming from an RVPLS SAP and configures the entire set of VPRN as well as attached RVPLS services in OISM mode. Only one SBD RVPLS can exist in a given VPRN. In order to add or remove the **supplementary-broadcast-domain** option, the entire **evpn-tunnel** command must first be removed.

The configuration of **evpn-tunnel** without options is equivalent to the **ipv6-gateway-address ip** option.

The **no** form of this command disables the evpn-tunnel mode.

Default

no evpn-tunnel

Parameters

ipv6-gateway-address

Indicates whether the IPv6 Prefix route uses a GW IP or a GW MAC as gateway.

Values **ip, mac**

supplementary-broadcast-domain

Specifies to use the EVPN tunnel as a Supplementary Broadcast Domain (SBD). The SBD is used in EVPN OISM to advertise the SMET routes and receive the multicast traffic on egress PEs that are not attached to the source R-VPLS service.

Platforms

7705 SAR Gen 2

9.97 evpn-type

evpn-type

Syntax

evpn-type *type*

no evpn-type

Context

[Tree] (config>router>policy-options>policy-statement>entry>from evpn-type)

Full Context

configure router policy-options policy-statement entry from evpn-type

Description

This command matches BGP routes based on the EVPN route type. The route types supported in SR OS are the following:

- Type 1 or Auto-Discovery Ethernet Tag route, including both the AD per-ES and AD per-EVI routes Type 2 or MAC/IP route
- Type 2 or MAC/IP route
- Type 3 or IMET route, including Multicast Ethernet Tag
- Type 4 or ES (Ethernet Segment) route Type 5 of IP-prefix route, including IPv4 and IPv6 prefixes
- Type 6 or Selective Multicast Ethernet Tag route, including IPv4 and IPv6 multicast groups
- Type 7 or Multicast Join Synch route, including IPv4 and IPv6 multicast group
- Type 8 or Multicast Leave Synch route, including IPv4 and IPv6 multicast groups

The **no** form of this command removes the **evpn-type** matching.

Parameters

name

Specifies the EVPN route type.

Values 1 to 8

Platforms

7705 SAR Gen 2

9.98 exceed

exceed

Syntax

exceed

Context

[Tree] (config>qos>sap-egress>queue>drop-tail exceed)

Full Context

configure qos sap-egress queue drop-tail exceed

Description

Commands in this context configure the queue exceed drop tail parameters. The exceed drop tail defines the queue depth beyond which exceed-profile packets will not be accepted into the queue and will be discarded.

Platforms

7705 SAR Gen 2

exceed

Syntax

exceed

Context

[Tree] (cfg>qos>qgrps>egr>qgrp>queue>drop-tail exceed)

Full Context

configure qos queue-group-templates egress queue-group queue drop-tail exceed

Description

Commands in this context configure the queue exceed drop-tail parameters. The exceed drop tail defines the queue depth beyond which exceed-profile packets will not be accepted into the queue and will be discarded.

Platforms

7705 SAR Gen 2

9.99 exceed-action

exceed-action

Syntax

exceed-action {**discard** | **low-priority** | **none**}

Context

[Tree] (config>sys>security>dist-cpu-protection>policy>local-monitoring-policer exceed-action)

Full Context

configure system security dist-cpu-protection policy local-monitoring-policer exceed-action

Description

This command controls the action performed upon the extracted control packets when the configured policer rates are exceeded.

Default

exceed-action none

Parameters

discard

Discards packets that are nonconforming.

low-priority

Marks packets that are nonconforming as low-priority (discard eligible or out-profile). If there is congestion in the control plane of the SR OS then unmarked (green, hi-prio or in-profile) control packets are given preferential treatment.

none

no hold-down

Platforms

7705 SAR Gen 2

exceed-action

Syntax

exceed-action {**discard** [**hold-down** *seconds*] | **low-priority** [**hold-down** *seconds*] | **none**}

Context

[Tree] (config>sys>security>dist-cpu-protection>policy>static-policer exceed-action)

[Tree] (config>sys>security>dist-cpu-protection>policy>protocol>dynamic-parameters exceed-action)

Full Context

configure system security dist-cpu-protection policy static-policer exceed-action

configure system security dist-cpu-protection policy protocol dynamic-parameters exceed-action

Description

This command controls the action performed upon the extracted control packets when the configured policer rates are exceeded.

Default

exceed-action none

Parameters

discard

Discards packets that are nonconforming.

low-priority

Marks packets that are nonconforming as low-priority (for example, discard eligible or out-profile). If there is congestion in the control plane of the SR OS then unmarked (for example, green, hi-prio or in-profile) control packets are given preferential treatment.

hold-down *seconds*

When this optional parameter is specified, it causes the following "hold-down" behavior.

When the SR OS software detects that an enforcement policer has marked or discarded one or more packets (software may detect this some time after the packets are actually discarded), and an optional **hold-down** *seconds* value has been specified for the **exceed-action**, then the policer will be set into a "mark-all" or "drop-all" mode that cause the following:

- the policer state to be updated as normal
- all packets to be marked (if the action is "low-priority") or dropped (action = discard) regardless of the results of the policing decisions/actions/state.

The **hold-down** is cleared after approximately the configured time in seconds after it was set. The **hold-down** *seconds* option should be selected for protocols that receive more than one packet in a complete handshake/negotiation (for example, DHCP, PPP). **hold-down** is not applicable to a local monitoring policer. The "detection-time" will only start after any **hold-down** is complete. During the **hold-down** (and the detection-time), the policer

is considered as in an "exceed" state. The policer may re-enter the hold-down state if an exceed packet is detected during the detection-time countdown.

Configuring the **indefinite** parameter value will cause hold down to remain in place until the operator clears it manually using a tools command (**tools perform security dist-cpu-protection release-hold-down**) or removes the dist-cpu-protection policy from the object.

Configuring the **none** parameter value will disable hold down.

Values 1 to 10080, indefinite, none

Platforms

7705 SAR Gen 2

9.100 exceed-profile-octets-discarded-count

exceed-profile-octets-discarded-count

Syntax

[**no**] **exceed-profile-octets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>policer>e-counters exceed-profile-octets-discarded-count)

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters exceed-profile-octets-discarded-count)

Full Context

configure log accounting-policy custom-record policer e-counters exceed-profile-octets-discarded-count

configure log accounting-policy custom-record ref-policer e-counters exceed-profile-octets-discarded-count

Description

This command includes the exceed profile octets discarded count.

The **no** form of this command excludes the exceed profile octets discarded count.

Default

no exceed-profile-octets-discarded-count

Platforms

7705 SAR Gen 2

9.101 exceed-profile-octets-forwarded-count

exceed-profile-octets-forwarded-count

Syntax

[**no**] **exceed-profile-octets-forwarded-count**

Context

[Tree] (config>log>acct-policy>cr>policer>e-counters exceed-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters exceed-profile-octets-forwarded-count)

Full Context

configure log accounting-policy custom-record policer e-counters exceed-profile-octets-forwarded-count

configure log accounting-policy custom-record ref-policer e-counters exceed-profile-octets-forwarded-count

Description

This command includes the exceed profile octets forwarded count.

The **no** form of this command excludes the exceed profile octets forwarded count.

Default

no exceed-profile-octets-forwarded-count

Platforms

7705 SAR Gen 2

9.102 exceed-profile-octets-offered-count

exceed-profile-octets-offered-count

Syntax

[**no**] **exceed-profile-octets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>policer>e-counters exceed-profile-octets-offered-count)

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters exceed-profile-octets-offered-count)

Full Context

configure log accounting-policy custom-record policer e-counters exceed-profile-octets-offered-count

configure log accounting-policy custom-record ref-policer e-counters exceed-profile-octets-offered-count

Description

This command includes the exceed profile octets offered count.

The **no** form of this command excludes the exceed profile octets offered count.

Default

no exceed-profile-octets-offered-count

Platforms

7705 SAR Gen 2

9.103 exceed-profile-packets-discarded-count

exceed-profile-packets-discarded-count

Syntax

[**no**] **exceed-profile-packets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters exceed-profile-packets-discarded-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters exceed-profile-packets-discarded-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters exceed-profile-packets-discarded-count

configure log accounting-policy custom-record policer e-counters exceed-profile-packets-discarded-count

Description

This command includes the exceed profile packets discarded count.

The **no** form of this command excludes the exceed profile packets discarded count.

Default

no exceed-profile-packets-discarded-count

Platforms

7705 SAR Gen 2

9.104 exceed-profile-packets-forwarded-count

exceed-profile-packets-forwarded-count

Syntax

[**no**] **exceed-profile-packets-forwarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters exceed-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters exceed-profile-packets-forwarded-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters exceed-profile-packets-forwarded-count

configure log accounting-policy custom-record policer e-counters exceed-profile-packets-forwarded-count

Description

This command includes the exceed profile packets forwarded count.

The **no** form of this command excludes the exceed profile packets forwarded count.

Default

no exceed-profile-packets-forwarded-count

Platforms

7705 SAR Gen 2

9.105 exceed-profile-packets-offered-count

exceed-profile-packets-offered-count

Syntax

[**no**] **exceed-profile-packets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters exceed-profile-packets-offered-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters exceed-profile-packets-offered-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters exceed-profile-packets-offered-count

configure log accounting-policy custom-record policer e-counters exceed-profile-packets-offered-count

Description

This command includes the exceed profile packets offered count.

The **no** form of this command excludes the exceed profile packets offered count.

Default

no exceed-profile-packets-offered-count

Platforms

7705 SAR Gen 2

9.106 exception

exception

Syntax

[**no**] **exception**

Context

[Tree] (debug>service>id>stp exception)

Full Context

debug service id stp exception

Description

This command enables STP debugging for exceptions.

The **no** form of the command disables debugging.

Platforms

7705 SAR Gen 2

9.107 exclude

exclude

Syntax

exclude

Context

[Tree] (config>service>vprn>isis>loopfree-alternates exclude)

Full Context

configure service vprn isis loopfree-alternates exclude

Description

This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.

The user can exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority. The prefix tag will, however, be used in the main SPF.

Note:
Prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action of the **exclude** command, when not explicitly specified by the user in the prefix policy, is a "reject". Thus, regardless of whether the user has explicitly added the statement "default-action reject" to the prefix policy, a prefix that does not match any entry in the policy is accepted into LFA SPF.

The **no** form of this command deletes the exclude prefix policy.

Default

no exclude

Platforms

7705 SAR Gen 2

exclude

Syntax

exclude

Context

[Tree] (config>service>vprn>ospf3>loopfree-alternates exclude)

[Tree] (config>service>vprn>ospf>loopfree-alternates exclude)

Full Context

configure service vprn ospf3 loopfree-alternates exclude

configure service vprn ospf loopfree-alternates exclude

Description

This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.

The implementation already allows the user to exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority. The prefix tag will, however, be used in the main SPF.

Note:
Prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action of the **exclude** command, when not explicitly specified by the user in the prefix policy, is a "reject". Thus, regardless if the user did or did not explicitly add the statement "default-action reject" to the prefix policy, a prefix that did not match any entry in the policy will be accepted into LFA SPF.

The **no** form of this command deletes the exclude prefix policy.

Default

no exclude

Platforms

7705 SAR Gen 2

exclude

Syntax

exclude *group-name* [*group-name*]

no exclude [*group-name* [*group-name*]]

Context

[Tree] (config>router>mpls>lsp exclude)

[Tree] (config>router>mpls>lsp>secondary exclude)

[Tree] (config>router>mpls>lsp>primary exclude)

[Tree] (config>router>mpls>lsp-template exclude)

Full Context

configure router mpls lsp exclude

configure router mpls lsp secondary exclude

configure router mpls lsp primary exclude

configure router mpls lsp-template exclude

Description

This command specifies the admin groups to be excluded when an LSP is set up. Up to five groups per operation can be specified, up to 32 maximum. The admin groups are defined in the **config>router>if-attribute>admin-group** context.

Use the **no** form of this command to remove the exclude command.

Default

no exclude

Parameters

group-name

Specifies the existing group-name to be excluded when an LSP is set up.

Platforms

7705 SAR Gen 2

exclude

Syntax

[**no**] **exclude** *tag*

Context

[Tree] (config>router>admin-tags>route-admin-tag-policy exclude)

Full Context

configure router admin-tags route-admin-tag-policy exclude

Description

This configures an admin tag to be excluded when matching a route against an LSP.

Up to eight exclusion statements are supported per policy.

The **no** form of this command removes the admin tag from the exclude statement.

Parameters

tag

Specifies the value of the admin tag, up to 32 characters.

Platforms

7705 SAR Gen 2

exclude

Syntax

exclude

Context

[Tree] (config>router>fad>flex-algo exclude)

Full Context

configure router flexible-algorithm-definitions flex-algo exclude

Description

Commands in this context configure administrative groups that will be excluded from the flexible algorithm topology graph.

If the defined FAD includes administrative groups link in its exclude list, the specified links are excluded from the topology graph.

Platforms

7705 SAR Gen 2

exclude

Syntax

exclude

Context

[Tree] (config>router>isis>loopfree-alternates exclude)

Full Context

configure router isis loopfree-alternates exclude

Description

Commands in this context configure a prefix policy for excluding specific prefixes in the LFA calculation by ISIS or OSPF.

Platforms

7705 SAR Gen 2

exclude

Syntax

exclude

Context

[Tree] (config>router>ospf>loopfree-alternates exclude)

[Tree] (config>router>ospf3>loopfree-alternates exclude)

Full Context

configure router ospf loopfree-alternates exclude

configure router ospf3 loopfree-alternates exclude

Description

Commands in this context configure a prefix policy for excluding specific prefixes in the LFA calculation by ISIS or OSPF.

Platforms

7705 SAR Gen 2

9.108 exclude-addresses

exclude-addresses

Syntax

[**no**] **exclude-addresses** *start-ip-address* [*end-ip-address*]

Context

[Tree] (config>service>vprn>dhcp>server>pool>subnet exclude-addresses)

[Tree] (config>router>dhcp>server>pool>subnet exclude-addresses)

Full Context

configure service vprn dhcp local-dhcp-server pool subnet exclude-addresses

configure router dhcp local-dhcp-server pool subnet exclude-addresses

Description

This command specifies a range of IP addresses that excluded from the pool of IP addresses in this subnet.

The **no** form of the removes the parameters from the configuration.

Parameters

start-ip-address

Specifies the start address of this range to exclude. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

Values a.b.c.d

end-ip-address

Specifies the end address of this range to exclude. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

Values a.b.c.d

Platforms

7705 SAR Gen 2

9.109 exclude-from-avg

exclude-from-avg

Syntax

exclude-from-avg {**forward** | **backward** | **round-trip**} **bins** *bin-numbers*

no exclude-from-avg (**forward** | **backward** | **round-trip**}

Context

[Tree] (config>oam-pm>bin-group>bin-type exclude-from-avg)

Full Context

configure oam-pm bin-group bin-type exclude-from-avg

Description

This optional command allows the results from probes that map to the specified bins within the bin type to be excluded from the average calculation. Individual counters are incremented in the bin, but the average is not affected by the value of the excluded delay metric for the individual probes in this bin. The bin group does not allow this command to be added, modified, or deleted when a test is actively referencing the bin group. Sessions that reference the bin group must have the bin group and tests shut down before changes can be made.

The **no** form of this command removes the exclusion, and all bins are included in the average calculation.

Default

no exclude-from-avg forward

no exclude-from-avg backward

no exclude-from-avg round-trip

Parameters

forward

Specifies the forward direction bin.

backward

Specifies the backward direction bin.

round-trip

Specifies the round-trip direction bin.

bin-numbers

Specifies the bin numbers to be excluded from the average calculation. The values typically represent, but are not restricted to, the highest and lowest configured bins in order to eliminate outlying results that are not representative of network performance.

A hyphen can be entered between bin numbers to include a continuous sequence of bins; for example, entering 7-9 would specify bins 7, 8, and 9. Commas can be entered between bin numbers to include separate or non-continuous bins; for example, entering 0,8,9 would specify bins 0, 8, and 9. Both hyphens and commas can be used in this manner in the same configuration; for example, entering 0,7-9 would include bins 0, 7, 8, and 9. All bin numbers specified as part of this command must be configured. If a specified bin does not exist, the command fails.

Values 0 to 9

Platforms

7705 SAR Gen 2

9.110 exclude-group

exclude-group

Syntax

[**no**] **exclude-group** *ip-admin-group-name*

Context

[Tree] (config>router>route-next-hop-policy>template exclude-group)

Full Context

configure router route-next-hop-policy template exclude-group

Description

This command configures the admin group constraint into the route next-hop policy template.

Each group is entered individually. The **include-group** statement instructs the LFA SPF selection algorithm to pick up a subset of LFA next-hops among the links that belong to one or more of the specified admin groups. A link that does not belong to at least one of the admin-groups is excluded. However, a link can still be selected if it belongs to one of the groups in an include-group statement but also belongs to other groups that are not part of any include-group statement in the route next-hop policy.

The **pref** option is used to provide a relative preference for the admin group to select. A lower preference value means that LFA SPF will first attempt to select an LFA backup next-hop that is a member of the corresponding admin group. If none is found, then the admin group with the next highest preference value is evaluated. If no preference is configured for a given admin group name, then it is supposed to be the least preferred, that is, numerically the highest preference value.

When evaluating multiple **include-group** statements within the same preference, any link that belongs to one or more of the included admin groups can be selected as an LFA next-hop. There is no relative preference based on how many of those included admin groups the link is a member of.

The **exclude-group** statement simply prunes all links belonging to the specified admin group before making the LFA backup next-hop selection for a prefix.

If the same group name is part of both include and exclude statements, the exclude statement will win. It other words, the exclude statement can be viewed as having an implicit preference value of zero (0).

The admin-group criteria are applied before running the LFA next-hop selection algorithm.

The **no** form deletes the admin group constraint from the route next-hop policy template.

Parameters

ip-admin-group-name

Specifies the name of the group, up to 32 characters.

Platforms

7705 SAR Gen 2

9.111 exclude-mac-policy

exclude-mac-policy

Syntax

exclude-mac-policy *mac-policy-id*

no exclude-mac-policy

Context

[Tree] (config>port>ethernet>dot1x>macsec exclude-mac-policy)

Full Context

configure port ethernet dot1x macsec exclude-mac-policy

Description

This command specifies the MAC policy to be excluded from MACsec encryption.

The **no** form of this command removes the policy from the MACsec and allows all destination MAC addresses.

Default

no exclude-mac-policy

Parameters

mac-policy-id

Specifies the MAC policy to exclude from the configuration.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

9.112 exclude-node

exclude-node

Syntax

exclude-node *ip-address*

no exclude-node

Context

[Tree] (config>router>mpls>lsp exclude-node)

Full Context

configure router mpls lsp exclude-node

Description

This command enables the option to include XRO object in the bypass LSP PATH message object. The exclude-node option is required for manual bypass LSP with XRO to FRR protect ABR node in a multi-vendor network deployment. This command must be configured on the PLR node that protects the ABR node. The ABR node IP address must be configured as exclude-node.

Default

no exclude-node

Platforms

7705 SAR Gen 2

9.113 exclude-prefix

exclude-prefix

Syntax

[**no**] **exclude-prefix** *ipv6-prefix/prefix-length*

Context

[Tree] (config>service>vprn>dhcp6>server>pool exclude-prefix)

[Tree] (config>router>dhcp6>server>pool exclude-prefix)

Full Context

configure service vprn dhcp6 local-dhcp-server pool exclude-prefix

configure router dhcp6 local-dhcp-server pool exclude-prefix

Description

This command defines a prefix that to be excluded from available prefix in the pool for DHCP6. The typical use case is to exclude the interface address.

- A held lease is deleted if it got excluded by an exclude prefix.
- An exclude range can never exclude only a part of an existing lease. If for example a /63 PD is assigned, an exclude of /64 which belongs to this /63 cannot be configured.
- A single exclude prefix can never exclude a whole include prefix.
- When applying or removing an exclude prefix, the threshold stats are adjusted to reflect the actual address space and its usage.

The **no** form of this command removes the prefix that is to be excluded from available prefix in the pool.

Parameters

ipv6-prefix/prefix-length

Specifies an IPv6 prefix and prefix length.

Values

ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)	
	x:x:x:x:x:x:d.d.d.d	
	x - [0 to FFFF]H	
	d - [0 to 255]D	
prefix-length	0 to 128	

Platforms

7705 SAR Gen 2

9.114 exclude-protocol

exclude-protocol

Syntax

[**no**] **exclude-protocol** {*protocol-name*}

Context

[Tree] (config>port>ethernet>dot1x>macsec exclude-protocol)

Full Context

configure port ethernet dot1x macsec exclude-protocol

Description

Specifies protocols whose packets are not secured using Media Access Control Security (MACsec) when MACsec is enabled on a port.

When this option is enabled in a connectivity association that is attached to an interface, MACsec is not enabled for all packets of the specified protocols that are sent and received on the link.

When this option is enabled on a port where MACsec is configured, packets of the specified protocols are sent and accepted in cleartext.

The **no** form of this command secures the packets of the specified protocol.

Default

no exclude-protocol

Parameters

protocol-name

Specifies the protocol name.

Values **cdp**, **lacp**, **lldp**, **eapol-start**, **efm-oam**, **eth-cfm**, **ptp**, **ubfd**

Platforms

7705 SAR Gen 2

9.115 exclusive-lock-time

exclusive-lock-time

Syntax

exclusive-lock-time *seconds*

no exclusive-lock

Context

[Tree] (config>router>policy-options exclusive-lock-time)

Full Context

configure router policy-options exclusive-lock-time

Description

This command specifies the inactivity timer for the exclusive lock time for policy editing. When a session is idle for greater than this time, the lock is removed and the configuration changes is aborted.

Default

exclusive-lock-time 300

Parameters

seconds

Specifies the duration the session with exclusive lock may be inactive.

Values Values: 1 to 3600

Platforms

7705 SAR Gen 2

9.116 exec

exec

Syntax

exec [**-echo**] [**-syntax**] {*file-name* | *eof-marker-string*} [**-argument** [*256 chars max*] [[*256 chars max*]]

Context

[Tree] (exec)

Full Context

exec

Description

This command executes the contents of a text file as if they were CLI commands entered at the console.

exec commands do not have **no** versions.

Related Commands:

boot-bad-exec: Use this command to configure a URL for a CLI script to exec following a failed configuration boot.

boot-good-exec: Use this command to configure a URL for a CLI script to exec following a successful configuration boot.

stdin can be used as the source of commands for the **exec** command. When **stdin** is used as the **exec** command input, the command list is terminated with <Ctrl-C>, "EOF<Return>" or "*eof_string*<Return>".

If an error occurs entering an exec file sourced from stdin, all commands after the command returning the error will be silently ignored. The **exec** command will indicate the command error line number when the stdin input is terminated with an end-of-file input.

Example:

Assume the *test.cfg* file has the following commands:

echo $(1)

echo $(2)

echo $(3)

Enter the following command:

exec test.cfg –arguments 10 20 30

The output from this command will be:

```
10
20
30
```

Parameters

-echo

Echoes the contents of the **exec** file to the session screen as it executes.

Default echo disabled

-syntax

Performs a syntax check of the file without executing the commands. Syntax checking will be able to find invalid commands and keywords, but it will not be able to validate erroneous user- supplied parameters.

Default execute file commands

file-name

Specifies the text file with CLI commands to execute, up to 256 characters.

eof-marker-string

Specifies the ASCII printable string used to indicate the end of the exec file when stdin is used as the exec file source. <Ctrl-C> and "EOF" can always be used to terminate an exec file sourced from stdin up to 254 characters.

Default EOF

-argument

Specifies up to five arguments, each up to 256 characters.

Platforms

7705 SAR Gen 2

9.117 exit

exit

Syntax

exit [**all**]

Context

[Tree] (exit)

Full Context

exit

Description

This command returns to the context from which the current level was entered. For example, to navigate to the current level on a context by context basis, then the **exit** command only moves the cursor back one level.

```
A:ALA-1# configure
A:ALA-1>config# router
A:ALA-1>config>router# ospf
A:ALA-1>config>router>ospf# exit
A:ALA-1>config>router# exit
A:ALA-1>config# exit
```

When navigating to the current level by entering a command string, the **exit** command returns the cursor to the context in which the command was initially entered.

```
A:ALA-1# configure router ospf
A:ALA-1>config>router>ospf# exit
A:ALA-1#
```

The **exit all** command moves the cursor all the way back to the root level.

```
A:ALA-1# configure
A:ALA-1>config# router
A:ALA-1>config>router# ospf
A:ALA-1>config>router>ospf# exit all
A:ALA-1#
```

Parameters

all

Exits back to the root CLI context.

Platforms

7705 SAR Gen 2

9.118 expire-time

expire-time

Syntax

expire-time {*seconds* | **forever**}

Context

[Tree] (config>system>script-control>script-policy expire-time)

Full Context

configure system script-control script-policy expire-time

Description

This command is used to configure the maximum amount of time to keep the run history status entry from a script run.

Default

expire-time 3600

Parameters

seconds

Specifies the time to keep the run history status entry, in seconds.

Values 0 to 21474836

Default 3600 (1 hour)

forever

Specifies to keep the run history status entry indefinitely.

Platforms

7705 SAR Gen 2

9.119 exponential-backoff

exponential-backoff

Syntax

[**no**] **exponential-backoff**

Context

[Tree] (config>system>login-control exponential-backoff)

Full Context

configure system login-control exponential-backoff

Description

This command enables the exponential-backoff of the login prompt. The exponential-backoff command is used to deter dictionary attacks, when a malicious user can gain access to the CLI by using a script to try **admin** with any conceivable password.

The **no** form of this command disables exponential-backoff.

Default

no exponential-backoff

Platforms

7705 SAR Gen 2

9.120 exponential-backoff-retry

exponential-backoff-retry

Syntax

exponential-backoff-retry

no exponential-backoff-retry

Context

[Tree] (config>router>mpls exponential-backoff-retry)

Full Context

configure router mpls exponential-backoff-retry

Description

This command enables the use of an exponential back-off timer when re-trying an LSP. When an LSP path establishment attempt fails, the path is put into retry procedures and a new attempt will be performed at the expiry of the user-configurable retry timer (config>router>mpls>lsp>retry-timer). By default, the retry time is constant for every attempt. The exponential back-off timer procedures will double the value of the user configured retry timer value at every failure of the attempt to adjust to the potential network congestion that caused the failure. An LSP establishment fails if no Resv message was received and the Path message retry timer expired or a PathErr message was received before the timer expired.

Platforms

7705 SAR Gen 2

9.121 export

export

Syntax

export *plcy-or-long-expr* [*plcy-or-expr*]

no export

Context

[Tree] (config>service>vprn>bgp>group>neighbor export)

[Tree] (config>service>vprn>bgp>group export)

[Tree] (config>service>vprn>bgp export)

Full Context

configure service vprn bgp group neighbor export

configure service vprn bgp group export

configure service vprn bgp export

Description

This command is used to specify route policies that control how outbound routes transmitted to certain peers are handled. Route policies are configured in the **config>router>policy-options** context.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in a peer-group) or neighbor level (only applies to the specified peer). The most specific level is used.

The **export** command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine the modifications of each route and the final action to accept or reject the route.

Only one of the 15 objects referenced by the **export** command can be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters while the remaining 14 objects have a maximum length of 64 characters each.

When multiple **export** commands are issued, the last command entered overrides the previous command.

When an export policy is not specified, BGP-learned routes are advertised by default; non-BGP routes are not advertised.

The **no** form of this command removes the policy association.

Default

no export

Parameters

plcy-or-long-expr

Specifies the route policy name, up to 64 characters in length, or a policy logical expression, up to 255 characters in length.

plcy-or-expr

Specifies the route policy name, up to 64 characters in length, or a policy logical expression, up to 255 characters in length.

Platforms

7705 SAR Gen 2

export

Syntax

[**no**] **export** *policy-name* [*policy-name* **...up to 5 max**]

Context

[Tree] (config>service>vprn>isis export)

Full Context

configure service vprn isis export

Description

This command configures export routing policies that determine the routes exported from the routing table to IS-IS.

If no export policy is defined, non IS-IS routes are not exported from the routing table manager to IS-IS.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered overrides the previous command. A maximum of five policy names can be specified.

If an **aggregate** command is also configured in the **config>router** context, then the aggregation is applied before the export policy is applied.

Routing policies are created in the **config>router>policy-options** context.

The **no** form of this command removes the specified *policy-name* or all policies from the configuration if no *policy-name* is specified.

Default

no export — No export policy name is specified.

Parameters

policy-name

The export policy name. Up to five *policy-name* arguments can be specified.

Platforms

7705 SAR Gen 2

export

Syntax

export *policy-name* [*policy-name*]

no export

Context

[Tree] (config>service>vprn>ospf>area export)

[Tree] (config>service>vprn>ospf3>area export)

Full Context

configure service vprn ospf area export

configure service vprn ospf3 area export

Description

This command configures ABR export policies to filter OSPFv2 Type 3 Summary-LSAs or OSPFv3 Inter-Area-Prefix-LSA between areas, in to only permit the export of specified routes into an area.

This command cannot be used in OSPF area 0.

The **no** form of this command reverts to the default value.

Default

no export

Parameters

policy-name

Specifies the export route policy name. A maximum of five policy names may be specified. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), enclose the entire string in double quotes.

The specified policy names must be predefined and already exist in the system.

Platforms

7705 SAR Gen 2

export

Syntax

export *policy-name* [*policy-name*]

no export

Context

[Tree] (config>service>vprn>ospf3 export)

[Tree] (config>service>vprn>ospf export)

Full Context

configure service vprn ospf3 export

configure service vprn ospf export

Description

This command associates export route policies to determine which routes are exported from the route table to OSPF. Export polices are only in effect if OSPF is configured as an ASBR.

If no export policy is specified, non-OSPF routes are not exported from the routing table manager to OSPF.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of this command removes all policies from the configuration.

Default

no export — No export route policies specified.

Parameters

policy-name

Specifies the export route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

The specified policy name(s) must be predefined and already exist in the system.

Platforms

7705 SAR Gen 2

export

Syntax

export *policy-name* [*policy-name*...(up to 5 max)]

no export

Context

[Tree] (config>service>vprn>ripng>group>neighbor export)

[Tree] (config>service>vprn>ripng>group export)

[Tree] (config>service>vprn>ripng export)

[Tree] (config>service>vprn>rip>group export)

[Tree] (config>service>vprn>rip>group>neighbor export)

[Tree] (config>service>vprn>rip export)

Full Context

configure service vprn ripng group neighbor export

configure service vprn ripng group export

configure service vprn ripng export

configure service vprn rip group export

configure service vprn rip group neighbor export

configure service vprn rip export

Description

This command specifies the export route policies used to determine routes that are exported to RIP. If no export policy is specified, non-RIP routes will not be exported from the routing table manager to RIP; RIP-learned routes will be exported to RIP neighbors.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of this command removes all policies from the configuration.

Default

no export

Parameters

policy-name

The export route policy name. Allowed values are any string up to 32 characters in length and composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the string must be enclosed between double quotes. The specified names must already be defined.

Platforms

7705 SAR Gen 2

export

Syntax

export *policy-name* [*policy-name*]

no export

Context

[Tree] (config>router>ldp export)

Full Context

configure router ldp export

Description

This command specifies the export route policies used to determine which routes are exported to LDP. Policies are configured in the **config>router>policy-options** context.

If no export policy is specified, non-LDP routes will not be exported from the routing table manager to LDP. LDP-learned routes will be exported to LDP neighbors. Present implementation of export policy (outbound filtering) can be used "only" to add FECs for label propagation. The export policy does not control propagation of FECs that an LSR receives from its neighbors.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of 5 policy names can be specified.

The **no** form of this command removes all policies from the configuration.

Default

no export — No export route policies specified.

Parameters

policy-name

Specifies up to five export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

Platforms

7705 SAR Gen 2

export

Syntax

export type {*type*} **input** *filename* **output** *url-string* **format** *output-format* [**password** [*32 chars max*]] [**pkey** *filename*]

Context

[Tree] (admin>certificate export)

Full Context

admin certificate export

Description

This command performs certificate operations.

Parameters

url-string

Specifies the local CF card url of the file.

Values

url-string	<local-url> [up to 99 characters]
local-url	<cflash-id>/<file-path>
cflash-id	cf1: | cf2: | cf3:

type

Specifies the type of input file.

Values cert, key, crl

format

Specifies the format of output file.

Values pkcs10, pkcs12, pkcs7-der, pkcs7-pem, pem, der

Platforms

7705 SAR Gen 2

export

Syntax

export *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]

no export [*plcy-or-long-expr*]

Context

[Tree] (config>router>bgp>group export)

[Tree] (config>router>bgp>group>neighbor export)

[Tree] (config>router>bgp export)

Full Context

configure router bgp group export

configure router bgp group neighbor export

configure router bgp export

Description

This command specifies route policies that control the handling of outbound routes transmitted to all peers. Route policies are configured in the **config>router>policy-options** context.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific level is used.

The export command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine the modifications of each route and the final action to accept or reject the route.

Only one of the 15 objects referenced by the command can be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters; the remaining 14 objects have a maximum length of 64 characters each.

When multiple export commands are issued, the last command entered overrides the previous command.

When an export policy is not specified, BGP-learned routes are advertised by default and non-BGP routes are not advertised.

The **no** form of this command removes the policy association.

Default

no export

Parameters

plcy-or-long-expr

Specifies the route policy name (up to 64 characters) or a policy logical expression (up to 255 characters long). Allowed values are any string up to 255 characters long composed of

printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

plcy-or-expr

Specifies up to 14 route policy names (up to 64 characters each) or a policy logical expression (up to 64 characters long). Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

export

Syntax

[**no**] **export** *policy-name* [*policy-name*]

Context

[Tree] (config>router>isis export)

Full Context

configure router isis export

Description

This command configures export routing policies that determine the routes exported from the routing table to IS-IS.

If no export policy is defined, non IS-IS routes are not exported from the routing table manager to IS-IS.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered overrides the previous command. A maximum of five policy names can be specified.

If an **aggregate** command is also configured in the **config>router** context, then the aggregation is applied before the export policy is applied.

Routing policies are created in the **config>router>policy-options** context.

The **no** form of this command removes the specified *policy-name* or all policies from the configuration if no *policy-name* is specified.

Parameters

policy-name

Specifies up to five export policy names.

Platforms

7705 SAR Gen 2

export

Syntax

export *policy-name* [*policy-name*]

no export

Context

[Tree] (config>router>ospf3 export)

[Tree] (config>router>ospf export)

Full Context

configure router ospf3 export

configure router ospf export

Description

This command associates export route policies to determine which routes are exported from the route table to OSPF. Export polices are only in effect if OSPF is configured as an ASBR.

If no export policy is specified, non-OSPF routes are not exported from the routing table manager to OSPF.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of this command removes all policies from the configuration.

Default

no export

Parameters

policy-name

Specifies up to 5 export route policy names. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. The specified names must already be defined.

Platforms

7705 SAR Gen 2

export

Syntax

[**no**] **export** *policy-name* [*policy-name*]

Context

[Tree] (config>router>ospf>area export)

[Tree] (config>router>ospf3>area export)

Full Context

configure router ospf area export

configure router ospf3 area export

Description

This command configures ABR export policies to filter OSPFv2 Type 3 Summary-LSAs or OSPFv3 Inter-Area-Prefix-LSA between areas, in order to only permit the specified routes from being exported into an area.

This command cannot be used in OSPF area 0.

The **no** form of this command reverts to the default value.

Default

no export

Parameters

policy-name

Specifies up to five export route policy names. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. The specified names must already be defined.

Platforms

7705 SAR Gen 2

export

Syntax

export *policy-name* [*policy-name*]

no export

Context

[Tree] (config>router>rip>group export)

[Tree] (config>router>rip>group>neighbor export)

[Tree] (config>router>ripng>group>neighbor export)

[Tree] (config>router>ripng export)

[Tree] (config>router>rip export)

[Tree] (config>router>ripng>group export)

Full Context

configure router rip group export

configure router rip group neighbor export

configure router ripng group neighbor export

configure router ripng export

configure router rip export

configure router ripng group export

Description

This command specifies the export route policies used to determine which routes are exported to RIP.

If no export policy is specified, non-RIP routes will not be exported from the routing table manager to RIP. RIP-learned routes will be exported to RIP neighbors.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of the command removes all policies from the configuration.

Default

no export

Parameters

policy-name

Specifies up to five export route policy names. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on.), the entire string must be enclosed within double quotes.

The specified names must already be defined.

Platforms

7705 SAR Gen 2

export

Syntax

[**no**] **export**

Context

[Tree] (config>service>vprn>bgp-ipvpn>attribute-set export)

Full Context

configure service vprn bgp-ipvpn attribute-set export

Description

This command configures the router to add an ATTR_SET path attribute to all VPN-IP routes that come from the VRF export of BGP routes advertised by PE-CE peers of the VPRN. This attribute contains an exact copy of all BGP path attributes (post-import policy) of the PE-CE BGP route, excluding the NEXT_HOP, MP_REACH, and MP_UNREACH attributes, as well as the AS4_PATH or AS4_AGGREGATOR attributes. The origin AS in the ATTR_SET encodes the ASN (or confederation ID, if configured) of the exporting VPRN service. Neither the VRF export policy nor a regular BGP export policy is allowed to modify the contents of the ATTR_SET.

The **no** form of this command configures the router to not add an ATTR_SET path attribute to VPN-IP routes exported by the VPRN. Nokia recommends using the **no** form of this command, unless there is a requirement for the VPRN to deliver an independent domain Layer 3 VPN service.

Default

no export

Platforms

7705 SAR Gen 2

9.122 export-addresses

export-addresses

Syntax

export-addresses *policy-name* [*policy-name*]

no export-addresses

Context

[Tree] (config>router>ldp>session-params>peer export-addresses)

Full Context

configure router ldp session-parameters peer export-addresses

Description

This command specifies the export prefix policy to local addresses advertised to this peer.

Policies are configured in the **config>router>policy-options** context. A maximum of five policy names can be specified.

The **no** form of this command removes the policy from the configuration.

Parameters

policy-name

Specifies up to five export-prefix route policy names. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes.

If the string contains spaces, use double quotes to delimit the start and end of the string. The specified name(s) must already be defined.

Platforms

7705 SAR Gen 2

9.123 export-grt

export-grt

Syntax

export-grt *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]

no export-grt

Context

[Tree] (config>service>vprn>grt-lookup export-grt)

Full Context

configure service vprn grt-lookup export-grt

Description

This command uses the route policy to determine which routes are exported from the VRF to the GRT along with all the forwarding information. These entries are marked as BGP-VPN routes in the GRT. For proper routing to occur from the GRT to the VRF, the routes must be in the GRT.

Default

no export-grt

Parameters

plcy-or-long-expr

Specifies the route policy name, up to 64 characters, or a policy logical expression, up to 255 characters.

plcy-or-expr

Specifies the route policy name (up to 64 characters) or a policy logical expression (up to 255 characters). Up to four policy names or logical expressions can be specified in a single statement.

Platforms

7705 SAR Gen 2

9.124 export-inactive-bgp

export-inactive-bgp

Syntax

[**no**] **export-inactive-bgp**

Context

[Tree] (config>service>vprn export-inactive-bgp)

Full Context

configure service vprn export-inactive-bgp

Description

This command allows the preferred BGP route learned by a VPRN to be exported as the VPN route, even when if it is inactive in the route table because a preferred BGP VPRN route from another PE is present. This overrides the default state in which the VPRN cannot export an inactive BGP route.

For the BGP route to be exported, the VRF export policy must accept it.

This command applies to both MPLS VPN and SRv6 VPN routes. In SRv6 VPN routes the advertised instruction is an End.DT, while in MPLS VPN routes the advertised label is a per-next-hop label.

This "best-external" type of route advertisement is useful in active/standby multi-homing scenarios because it ensures that all PEs know about the backup path provided by the standby PE.

Default

no export-inactive-bgp

Platforms

7705 SAR Gen 2

9.125 export-inactive-bgp-enhanced

export-inactive-bgp-enhanced

Syntax

[**no**] **export-inactive-bgp-enhanced**

Context

[Tree] (config>service>vprn export-inactive-bgp-enhanced)

Full Context

configure service vprn export-inactive-bgp-enhanced

Description

This command configures the router to allow a BGP route that is inactive (because a better non-BGP route for the same prefix is present) to be exportable as a VPN-IP route.

A BGP route learned from a VPRN BGP peer is exportable as a VPN-IP route, only if it is the best route for the prefix and is installed in the route table of the VPRN. If the **export-inactive-bgp** command is enabled in the VPRN configuration, this rule is relaxed, and the best inactive VPRN BGP route is exportable as a VPN-IP route, provided that the active installed route for the prefix is an imported VPN-IP route.

The rule described in the preceding paragraph can be relaxed even further by enabling this command. When this command is enabled, the best inactive VPRN BGP route (best amongst all routes received from all CEs) is exportable as a VPN-IP route, regardless of the route type of the active installed route.

The configuration of this command overrides the **export-inactive-bgp** command. If this command is already enabled, do not enable the **export-inactive-bgp** command.

The **no** form of this command disables the router from allowing an inactive BGP route in the presence of a better non-BGP route to be exportable as a VPN-IP route.

Default

no export-inactive-bgp-enhanced

Platforms

7705 SAR Gen 2

9.126 export-limit

export-limit

Syntax

export-limit *num-routes*

no export-limit

Context

[Tree] (config>service>vprn>ospf export-limit)

[Tree] (config>service>vprn>grt-lookup export-limit)

[Tree] (config>service>vprn>ospf3 export-limit)

Full Context

configure service vprn ospf export-limit

configure service vprn grt-lookup export-limit

configure service vprn ospf3 export-limit

Description

This command limits the total number of routes exported from the VRF to the GRT. Configuring **export-limit 0** disables the maximum limit for routes exported from the VRF to the GRT.

The **no** form of this command sets the export-limit to a default of five (5).

Default

export-limit 5

Parameters

num-routes

Specifies the maximum number of routes that can be exported. Configuring a num-routes value in a range of 1 to 1000 limits the number of routes to the specified value.

Values 0 to 1000

Platforms

7705 SAR Gen 2

export-limit

Syntax

export-limit *number* [**log** *percentage*]

no export-limit

Context

[Tree] (config>service>vprn>rip export-limit)

[Tree] (config>service>vprn>ripng export-limit)

Full Context

configure service vprn rip export-limit

configure service vprn ripng export-limit

Description

This command configures the maximum number of routes (prefixes) that can be exported into RIP from the route table.

The **no** form of this command removes the parameters from the configuration.

Default

no export-limit

Parameters

number

Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.

Values 1 to 4294967295

log *percentage*

Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.

Values 1 to 100

Platforms

7705 SAR Gen 2

export-limit

Syntax

export-limit *number* [**log** *percentage*]

no export-limit

Context

[Tree] (config>service>vprn>isis export-limit)

Full Context

configure service vprn isis export-limit

Description

This command configures the maximum number of routes (prefixes) that can be exported into IS-IS from the route table for the VPRN instance.

The **no** form of this command removes the parameters from the configuration.

Default

no export-limit - The export limit for routes or prefixes is disabled.

Parameters

number

Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.

Values 1 to 4294967295

log *percentage*

Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.

Values 1 to 100

Platforms

7705 SAR Gen 2

export-limit

Syntax

export-limit *number* [**log** *percentage*]

no export-limit

Context

[Tree] (config>router>isis export-limit)

Full Context

configure router isis export-limit

Description

This command configures the maximum number of routes (prefixes) that can be exported into IS-IS from the route table. After the maximum is reached, a warning log message is sent and additional routes are ignored.

The **no** form of this command removes the parameters from the configuration.

Parameters

number

Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.

Values 1 to 4294967295

percentage

Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.

Values 1 to 100

Platforms

7705 SAR Gen 2

export-limit

Syntax

export-limit *number* [**log** *percentage*]

no export-limit

Context

[Tree] (config>router>ospf3 export-limit)

[Tree] (config>router>ospf export-limit)

Full Context

configure router ospf3 export-limit

configure router ospf export-limit

Description

This command configures the maximum number of routes (prefixes) that can be exported into OSPF from the route table. After the maximum is reached, a warning log message is sent and additional routes are ignored.

The **no** form of this command removes the parameters from the configuration.

Default

no export-limit

Parameters

number

Specifies the maximum number of routes (prefixes) that can be exported into OSPF from the route table.

Values 1 to 4294967295

percentage

Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.

Values 1 to 100

Platforms

7705 SAR Gen 2

export-limit

Syntax

export-limit *number* [**log** *percentage*]

no export-limit

Context

[Tree] (config>router>ripng export-limit)

[Tree] (config>router>rip export-limit)

Full Context

configure router ripng export-limit

configure router rip export-limit

Description

This command configures the maximum number of routes (prefixes) that can be exported into RIP from the route table.

The **no** form of the command removes the parameters from the configuration.

Default

no export-limit

Parameters

number

Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.

Values 1 to 4294967295

percentage

Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.

Values 1 to 100

Platforms

7705 SAR Gen 2

9.127 export-prefixes

export-prefixes

Syntax

[**no**] **export-prefixes** *policy-name*

Context

[Tree] (config>router>ldp>session-params>peer export-prefixes)

Full Context

configure router ldp session-parameters peer export-prefixes

Description

This command specifies the export route policy used to determine which prefixes received from other LDP and T-LDP peers are re-distributed to this LDP peer via the LDP/T-LDP session to this peer. A prefix that is filtered out (deny) is not exported. A prefix that is filtered in (accept) will be exported.

If no export policy is specified, all FEC prefixes learned will be exported to this LDP peer. This policy is applied in addition to the global LDP policy and targeted session policy.

Policies are configured in the **config>router>policy-options** context. A maximum of five policy names can be specified. Peer address has to be the peer LSR-ID address.

The **no** form of this command removes the policy from the configuration.

Default

no export-prefixes - no export route policy is specified

Parameters

policy-name

Specifies up to five export-prefix route policy names. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string. The specified name(s) must already be defined.

Platforms

7705 SAR Gen 2

export-prefixes

Syntax

export-prefixes *policy-name* [*policy-name*]

no export-prefixes

Context

[Tree] (config>router>ldp>targeted-session export-prefixes)

Full Context

configure router ldp targeted-session export-prefixes

Description

This command specifies the export route policy used to determine which FEC prefix label bindings are exported from a targeted LDP session. A route that is filtered out (deny) will not be exported. A route that is filtered in (accept) will be exported.

If no export policy is specified, all bindings learned through a targeted LDP session will be exported to all targeted LDP peers. This policy is applied in addition to the global LDP policy.

Policies are configured in the **config>router>policy-options** context. A maximum of five policy names can be specified.

The **no** form of this command removes the policy from the configuration.

Parameters

policy-name

Specifies up to five export policy names. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

9.128 export-tunnel-table

export-tunnel-table

Syntax

export-tunnel-table *policy-name* [*policy-name*...(up to 5 max)]

no export-tunnel-table

Context

[Tree] (config>router>ldp export-tunnel-table)

Full Context

configure router ldp export-tunnel-table

Description

This command enables exports BGP label route and SR tunnels from the TTM into LDP for the purpose of stitching an LDP FEC to a BGP or SR tunnel for the same destination prefix.

To enable route stitching between LDP and BGP, separately configure tunnel table route export policies in both protocols and enable the advertisement of RFC 8277, *Using BGP to Bind MPLS Labels to Address Prefixes*, formatted labeled routes for prefixes learned from LDP FECs.

The BGP route export policy instructs BGP to listen to LDP route entries in the CPM Tunnel Table. If a /32 LDP FEC prefix matches an entry in the export policy, BGP originates a BGP labeled route, stitches it to the LDP FEC, and re-distributes the BGP labeled route to its Interior Border Gateway Protocol (IBGP) neighbors.

Using the following commands to add LDP FEC prefixes with the **from protocol ldp** statement in the existing BGP export policy configuration at the global level, peer-group level, or peer level:

- **config>router>bgp>export** *policy-name*
- **config>router>bgp>group>export** *policy-name*
- **config>router>bgp>group>neighbor>export** *policy-name*

To indicate to BGP to evaluate the entries with the **from protocol ldp** statement in the export policy when applied to a specific BGP neighbor, use commands:

- **config>router>bgp>group>neighbor>family label-ipv4** and
- **config>router>bgp>group>neighbor>advertise-ldp-prefix**

Without the latter configuration, only core IPv4 routes learned from RTM are advertised as BGP labeled routes to the neighbor. No stitching of LDP FEC to the BGP labeled route will be performed for this neighbor even if the same prefix was learned from LDP.

The LDP tunnel table route export policy instructs LDP to listen to BGP route entries in the CPM Tunnel Table. If a /32 BGP labeled route matches a prefix entry in the export policy, LDP originates an LDP FEC for the prefix, stitches it to the BGP labeled route, and re-distributes the LDP FEC to its IBGP neighbors.

The user can add BGP labeled route prefixes with the **from protocol bgp** statement in the configuration of the LDP tunnel table export policy. The **from protocol** statement is applied only when the protocol value is **ldp**. Policy entries with protocol values of **rsvp**, **bgp**, or any value other than **ldp** are ignored at the time the policy is applied to LDP.

In the LDP-to-SR data path direction, LDP listens to SR tunnel entries in the TTM. The user can restrict the export of SR tunnels to LDP from a specific prefix list. The user can also restrict the export to a specific IGP instance by optionally specifying the instance ID in the "from protocol” statement. The statement has an effect only when the protocol value is **isis** or **bgp**. Policy entries with any other protocol value are ignored at the time the policy is applied. If the user configures multiple **from protocol** statements in the same policy or does not include the **from protocol** statement but adds a default action of accept, then LDP will follow the TTM selection rules to select a tunnel to which it will stitch the LDP ILM:

1. LDP selects the tunnel from the lowest TTM preference protocol.
2. If two or more of IS-IS or OSPF protocol instances and BGP protocol have the same preference, then LDP selects the protocol using the default TTM protocol preference.
3. Within the same IGP protocol, LDP selects the lowest instance ID.

If an LDP FEC primary next-hop cannot be resolved using an RTM route and a SR tunnel of type SR-ISIS to the same destination prefix matches a prefix entry in the export policy, LDP programs an LDP ILM and stitches it to the SR node-SID tunnel endpoint. LDP also originates an FEC for the prefix and re-distributes

it to its LDP peers. When an LDP FEC is stitched to a SR tunnel, packets forwarded benefit from the protection of the LFA/remote LFA backup next-hop of the SR tunnel.

When resolving a FEC, LDP will prefer RTM over TTM when both resolutions are possible. That is, swapping the LDP ILM to a LDP NHLFE is preferred over stitching it to an SR tunnel endpoint.

Nokia recommends that the user should enable the bfd-enable option on the interfaces in LDP, IGP instance, and BGP contexts to speed up failure detection and activation of the SR LFA/remote-LFA backup next-hop or the BGP backup, depending on the stitching operation.

This feature is limited to IPv4 /32 prefixes in LDP, BGP and SR.

The **no** form of this command disables the export of BGP and SR tunnels to LDP.

Default

no export-tunnel-table

Parameters

policy-name

Specifies up to five export-tunnel-table route policy names. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string. The specified name(s) must already be defined.

Platforms

7705 SAR Gen 2

export-tunnel-table

Syntax

export-tunnel-table ldp

no export-tunnel-table

Context

[Tree] (config>router>isis>segment-routing export-tunnel-table)

Full Context

configure router isis segment-routing export-tunnel-table

Description

This command exports the LDP tunnels to an IGP instance for the purpose of stitching a SR tunnel to a LDP FEC for the same destination IPv4 /32 prefix.

In the SR-to-LDP data path direction, the SR mapping server provides a global policy for the prefixes corresponding to the LDP FECs the SR stitches to.

When this command is enabled in the segment-routing context of an IGP instance, IGP listens to LDP tunnel entries in the TTM. Whenever a LDP tunnel destination matches a prefix for which IGP received a prefix-SID sub-TLV from a mapping server, it instructs the SR module to program the SR ILM and to

stitch it to the LDP tunnel endpoint. The LDP FEC can be resolved via a static route, a IS-IS instance, or an OSPF instance.

When an SR tunnel is stitched to a LDP FEC, packets forwarded will benefit from the protection of the LFA backup next-hop of the LDP FEC.

When resolving a node SID, IGP will prefer resolution of prefix SID received in a IP Reach TLV over a prefix SID received via the mapping server. That is, swapping the SR ILM to a SR NHLFE is preferred over stitching it to a LDP tunnel endpoint.

Nokia recommends that the user should enable the bfd-enable option on the interfaces in both LDP and IGP instance contexts to speed up the failure detection and the activation of the LFA/remote-LFA backup next-hop in either direction of the stitching.

This feature is limited to IPv4 /32 prefixes in both LDP and SR.

The **no** form of this command disables the exporting of LDP tunnels to the IGP instance.

Default

no export-tunnel-table

Parameters

ldp

Exports LDP tunnels from the tunnel table into an IGP instance.

Platforms

7705 SAR Gen 2

export-tunnel-table

Syntax

[**no**] **export-tunnel-table ldp**

Context

[Tree] (config>router>ospf>segm-rtng export-tunnel-table)

Full Context

configure router ospf segment-routing export-tunnel-table

Description

This command enables exporting, to an IGP instance, the LDP tunnels for the purpose of stitching a SR tunnel to a LDP FEC for the same destination IPv4 /32 prefix.

In the SR-to-LDP data path direction, the SR mapping server provides a global policy for the prefixes corresponding to the LDP FECs that the SR stitches to.

When this command is enabled in the segment-routing context of an IGP instance, IGP listens to LDP tunnel entries in the TTM. Whenever a LDP tunnel destination matches a prefix for which IGP received a prefix-SID sub-TLV from a mapping server, it instructs the SR module to program the SR ILM and to

stitch it to the LDP tunnel endpoint. The LDP FEC can be resolved via a static route, a IS-IS instance, or an OSPF instance.

When an SR tunnel is stitched to a LDP FEC, packets forwarded will benefit from the protection of the LFA backup next hop of the LDP FEC.

When resolving a node SID, IGP will prefer resolution of prefix SID received in a IP Reach TLV over a prefix SID received via the mapping server. In other words, the swapping of the SR ILM to a SR NHLFE is preferred over stitching it to a LDP tunnel endpoint.

It is recommended to enable the **bfd-enable** option on the interfaces in both LDP and IGP instance contexts, to speed up the failure detection and the activation of the LFA/remote-LFA backup next hop in either direction of the stitching.

This feature is limited to IPv4 /32 prefixes in both LDP and SR.

The **no** form of this command disables the exporting of LDP tunnels to the IGP instance.

Platforms

7705 SAR Gen 2

9.129 export-v6-limit

export-v6-limit

Syntax

export-v6-limit *num-routes*

no export-v6-limit

Context

[Tree] (config>service>vprn>grt-lookup export-v6-limit)

Full Context

configure service vprn grt-lookup export-v6-limit

Description

This command limits the total number of IPv6 routes exported from the VPRN to the GRT. Configuring **export-v6-limit 0** disables the maximum limit for IPv6 routes exported from the VPRN to the GRT.

The **no** form of this command sets the export-limit to a default of 5.

Default

export-v6-limit 5

Parameters

num-routes

Specifies the maximum number of IPv6 routes that can be exported. Configuring a *num-routes* value in a range of 1 to 1000 limits the number of IPv6 routes to the specified value.

Values 0 to 1000

Platforms

7705 SAR Gen 2

9.130 expression

expression

Syntax

expression *regular-expression*

no expression

Context

[Tree] (config>router>policy-options>as-path expression)

Full Context

configure router policy-options as-path expression

Description

This command configures a route policy AS path regular expression statement to use in the route policy entries.

An AS path in a BGP route matches an AS path regular expression, if the path matches the pattern of the regular expression. A regular expression incorporates terms and operators that use the terms. An individual AS number is an elementary term in the AS path regular expression. More complex terms can be built from elementary terms. The following are key operators supported by SR OS:

- .
- *
- ?
- {n}
- {m,n}
- {m, }

To reverse the match criteria when specifying a list of ranges or single values using square brackets, use the non-match operator (^) before the elements within the square brackets.

The **no** form of this command deletes the AS path regular expression statement.

Parameters

regular-expression

The AS path regular expression. Allowed values are any string up to 255 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at signs (@); for example, "@variable@”.

null

The AS path expressed as an empty regular expression string.

Platforms

7705 SAR Gen 2

expression

Syntax

expression *expression* [**exact**]

no expression

Context

[Tree] (config>router>policy-options>community expression)

Full Context

configure router policy-options community expression

Description

This command creates a logical expression to match a route policy community.

The **no** form of this command deletes the logical expression.

Default

no expression

Parameters

expression *expression*

Specifies a logical expression containing terms and operators. It can contain sub-expressions enclosed in round brackets.

Values up to 900 characters

<expression> is one of the following: <expression> {AND| OR} <expression> [NOT] (<expression>) [NOT] <comm-id>

For example:

from community expression "[community list A] OR ([community list B] AND [community list C])"

exact

All the communities indicated by the expression must be present in the route in order for a match to occur.

Platforms

7705 SAR Gen 2

9.131 extended-community

extended-community

Syntax

[**no**] **extended-community**

Context

[Tree] (config>router>bgp>group>neighbor>outbound-route-filtering extended-community)

[Tree] (config>router>bgp>group>outbound-route-filtering extended-community)

[Tree] (config>router>bgp>outbound-route-filtering extended-community)

Full Context

configure router bgp group neighbor outbound-route-filtering extended-community

configure router bgp group outbound-route-filtering extended-community

configure router bgp outbound-route-filtering extended-community

Description

The extended-community command opens the configuration tree for sending or accepting extended-community based BGP filters.

For the **no** version of the command to work, all sub-commands (**send-orf**, **accept-orf**) must be removed first.

Default

no extended-community

Platforms

7705 SAR Gen 2

9.132 extended-lsa

extended-lsa

Syntax

extended-lsa {**sparse** | **only**}

no extended-lsa

Context

[Tree] (config>router>ospf3 extended-lsa)

Full Context

configure router ospf3 extended-lsa

Description

This command configures the use of extended LSA format in OSPFv3, as described in *draft-ietf-ospf-ospfv3-lsa-extend*.

Prior to this feature, SR OS used the fixed format LSA to carry the prefix and link information as described in RFC 5340, *OSPF for IPv6*. The fixed format is not extensible and the TLV format of the extended LSA must be used.

With this feature, the default mode of operation for OSPFv3 is referred to as **sparse** mode, meaning that the router will always advertise the fixed format for existing LSAs and will add the TLV-based extended LSA only when it needs to advertise new sub-TLVs. This mode of operation is similar to the way OSPFv2 advertises the segment routing information. It sends the prefix in the original fixed-format prefix LSA and then follows with the extended prefix TLV which is sent in an extended prefix opaque LSA containing the prefix SID sub-TLV.

The **extended-lsa only** value enables the full extended LSA mode. This causes all existing and new LSAs to use the extended LSA format.

The OSPFv3 instance must first be shut down before the user can change the mode of operation since the protocol must flush all LSAs and re-establish all adjacencies.

The **no** form of this command at the OSPFv3 instance level reverts the OSPFv3 instance to the default **sparse** mode of operation.

Default

extended-lsa sparse

Parameters

sparse

Enables the sparse mode of operation in an OSPFv3 instance.

only

Enables the full extended LSA mode of operation in an OSPFv3 instance.

Platforms

7705 SAR Gen 2

extended-lsa

Syntax

extended-lsa {**inherit** | **only**}

no extended-lsa

Context

[Tree] (config>router>ospf3>area extended-lsa)

Full Context

configure router ospf3 area extended-lsa

Description

This command configures the use of extended LSA format in a OSPFv3 area as described in *draft-ietf-ospf-ospfv3-lsa-extend*.

By default, the area inherits the instance-level configuration. The latter defaults to the **sparse** mode of operation. The **extended-lsa only** value enables the full extended LSA mode, which causes all existing and new LSAs to use the extended LSA format.

The OSPFv3 instance must first be shut down before the user can change the mode of operation since the protocol must flush all LSAs and reestablish all adjacencies.

The **no** form of this command at the area level returns the area to the default mode of inheriting the mode from the OSPFv3 instance level.

Default

extended-lsa inherit

Parameters

inherit

Configures the area to inherit the mode of operation enabled at the OSPFv3 instance level.

only

Enables the full extended LSA mode of operation in an OSPFv3 area.

Platforms

7705 SAR Gen 2

9.133 extended-nh-encoding

extended-nh-encoding

Syntax

extended-nh-encoding [**ipv4**]

no extended-nh-encoding

Context

[Tree] (config>service>vprn>bgp extended-nh-encoding)

[Tree] (config>service>vprn>bgp>group extended-nh-encoding)

[Tree] (config>service>vprn>bgp>group>neighbor extended-nh-encoding)

Full Context

configure service vprn bgp extended-nh-encoding

configure service vprn bgp group extended-nh-encoding

configure service vprn bgp group neighbor extended-nh-encoding

Description

This command configures BGP to advertise (at session OPEN) the capability to receive IPv4 or IPv4 routes with IPv4 or IPv6 next hops from the VPRN BGP peers included in the scope of the command. These peers should not send these routes unless they receive the capability. If the SR OS router receives an IPv4 route from a peer to which it did not advertise the necessary capability, the UPDATE message will be considered malformed and causes either a session reset or treat as withdraw behavior depending on the error handling settings.

The **no** form of this command causes the sending of an extended NH encoding BGP capability to the associated BGP peers to be inherited from a higher configuration level or disabled (if configured at the BGP level).

Default

no extended-nh-encoding

Parameters

ipv4

Specifies that the command should be applied to unlabeled unicast IPv4 routes.

Platforms

7705 SAR Gen 2

extended-nh-encoding

Syntax

extended-nh-encoding [**label-ipv4**] [**vpn-ipv4**] [**ipv4**]

no extended-nh-encoding

Context

[Tree] (config>router>bgp>group extended-nh-encoding)

[Tree] (config>router>bgp>group>neighbor extended-nh-encoding)

[Tree] (config>router>bgp extended-nh-encoding)

Full Context

configure router bgp group extended-nh-encoding

configure router bgp group neighbor extended-nh-encoding

configure router bgp extended-nh-encoding

Description

This command configures BGP to advertise (at session OPEN) the capability to receive label IPv4, VPN IPv4 routes, or IPv6 next hops from the peers. These peers should not send such routes unless they receive notification of this capability. If the SR OS router receives a label IPv4 or VPN IPv4 route from a peer to which it did not advertise the necessary capability, the UPDATE message will be considered malformed and this will cause either session reset or **treat-as-withdraw** behavior depending on the error handling settings.

The **no** form of this command causes the sending of an extended NH encoding BGP capability to the associated BGP peers to be inherited from a higher configuration level or disabled (if configured at the BGP level).

Default

no extended-nh-encoding

Parameters

label-ipv4

Instructs BGP to advertise an extended NH encoding capability for NLRI AFI=1, NLRI SAFI=4, and next-hop AFI=2.

vpn-ipv4

Instructs BGP to advertise an extended NH encoding capability for NLRI AFI=1, NLRI SAFI=128, and next-hop AFI=2.

ipv4

Instructs BGP to advertise an extended NH encoding capability for NLRI AFI=1, NLRI SAFI=1 and next-hop AFI=2.

Platforms

7705 SAR Gen 2

9.134 extended-sequence-number

extended-sequence-number

Syntax

[no] extended-sequence-number

Context

[Tree] (config>ipsec>ipsec-transform extended-sequence-number)

Full Context

configure ipsec ipsec-transform extended-sequence-number

Description

This command enables 64-bit extended sequence numbering support. This numbering is used for high throughput CHILD_SA to avoid frequent rekeying caused by sequence numbering wrap around.

The **no** form of this command disables extended sequence numbering support. Only 32-bit sequence numbering is supported.

Default

no extended-seq-number

Platforms

7705 SAR Gen 2

9.135 external

external

Syntax

[**no**] **external**

Context

[Tree] (config>router>policy-options>policy-statement>entry>from external)

Full Context

configure router policy-options policy-statement entry from external

Description

This command specifies the external route matching criteria for the entry.

Default

no external

Platforms

7705 SAR Gen 2

9.136 external-db-overflow

external-db-overflow

Syntax

external-db-overflow *limit interval*

no external-db-overflow

Context

[Tree] (config>service>vprn>ospf external-db-overflow)

[Tree] (config>service>vprn>ospf3 external-db-overflow)

Full Context

configure service vprn ospf external-db-overflow

configure service vprn ospf3 external-db-overflow

Description

This command enables limits on the number of non-default AS-external-LSA entries that can be stored in the LSDB and specifies a wait timer before processing these after the limit is exceeded.

The *limit* value specifies the maximum number of non-default AS-external-LSA entries that can be stored in the link-state database (LSDB). Placing a limit on the non-default AS-external-LSAs in the LSDB protects the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reach or exceed the *limit*, the table is in an overflow state. When in an overflow state, the router does not originate any new AS-external-LSAs and it withdraws all self-originated non-default external LSAs.

The *interval* specifies the amount of time to wait after an overflow state before regenerating and processing non-default AS-external-LSAs. The waiting period acts like a dampening period, which prevents the router from continuously running Shortest Path First (SPF) calculations caused by the excessive number of non-default AS-external LSAs.

The **external-db-overflow** must be set identically on all routers attached to any regular OSPF area. OSPF stub areas and not-so-stubby areas (NSSAs) are excluded.

The **no** form of this command disables limiting the number of non-default AS-external-LSA entries.

Default

no external-db-overflow — No limit on non-default AS-external-LSA entries.

Parameters

limit

The maximum number of non-default AS-external-LSA entries that can be stored in the LSDB before going into an overflow state expressed as a decimal integer.

Values -1 to 2147483647

Note:
Setting a value of -1 is equivalent to **no external-db-overflow**.

interval

The number of seconds after entering an overflow state before attempting to process non-default AS-external-LSAs expressed as a decimal integer.

Values 0 to 2147483647

Platforms

7705 SAR Gen 2

external-db-overflow

Syntax

external-db-overflow *limit interval*

no external-db-overflow

Context

[Tree] (config>router>ospf3 external-db-overflow)

[Tree] (config>router>ospf external-db-overflow)

Full Context

configure router ospf3 external-db-overflow

configure router ospf external-db-overflow

Description

This command enables limits on the number of non-default AS-external-LSA entries that can be stored in the LSDB and specifies a wait timer before processing these after the limit is exceeded.

The *limit* value specifies the maximum number of non-default AS-external-LSA entries that can be stored in the link-state database (LSDB). Placing a limit on the non-default AS-external-LSAs in the LSDB protects the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reach or exceed the *limit*, the table is in an overflow state. When in an overflow state, the router will not originate any new AS-external-LSAs. In fact, it withdraws all the self-originated non-default external LSAs.

The *interval* specifies the amount of time to wait after an overflow state before regenerating and processing non-default AS-external-LSAs. The waiting period acts like a dampening period preventing the router from continuously running Shortest Path First (SPF) calculations caused by the excessive number of non-default AS-external LSAs.

The **external-db-overflow** must be set identically on all routers attached to any regular OSPF area. OSPF stub areas and not-so-stubby areas (NSSAs) are excluded.

The **no** form of this command disables limiting the number of non-default AS-external-LSA entries.

Default

no external-db-overflow

Parameters

limit

Specifies the maximum number of non-default AS-external-LSA entries that can be stored in the LSDB before going into an overflow state expressed as a decimal integer.

Values 0 to 2147483647

interval

The number of seconds after entering an overflow state before attempting to process non-default AS-external-LSAs expressed as a decimal integer.

Values 0 to 2147483647

Platforms

7705 SAR Gen 2

9.137 external-preference

external-preference

Syntax

external-preference *preference*

no external-preference

Context

[Tree] (config>service>vprn>isis>level external-preference)

Full Context

configure service vprn isis level external-preference

Description

This command configures the external route preference for the IS-IS level.

The **external-preference** command configures the preference level of either IS-IS level 1 or IS-IS level 2 external routes. By default, the preferences are as listed in the table below.

A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference decides the route to use.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is dependent on the default preference table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of the route to use is determined by the configuration of the **ecmp** in the **config>router** context.

Default

Default preferences are listed in Table 33: Default Preferences.

Table 33: Default Preferences

Route Type	Preference	Configurable
Direct attached	0	No
Static route	5	Yes
MPLS	7	—
OSPF internal routes	10	No
IS-IS Level 1 internal	15	Yes
IS-IS Level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS Level 1 external	160	Yes
IS-IS Level 2 external	165	Yes
BGP	170	Yes
BGP	170	Yes

Note:

1. Internal preferences are changed using the **preference** command in the **config>router>isis>level** *level-number* context.

Parameters

preference

The preference for external routes at this level as expressed.

Values 1 to 255

Platforms

7705 SAR Gen 2

external-preference

Syntax

external-preference *preference*

no external-preference

Context

[Tree] (config>service>vprn>ospf3 external-preference)

[Tree] (config>service>vprn>ospf external-preference)

Full Context

configure service vprn ospf3 external-preference

configure service vprn ospf external-preference

Description

This command configures the preference for OSPF external routes.

A route can be learned by the router from different protocols, in which case the costs are not comparable. If this occurs, preference is used to decide which route is used.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is per the default preference table as defined in Table 34: Default External Route Preferences . If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of which route to use is determined by the configuration of the **ecmp** in the **config**>**router** context.

The **no** form of this command reverts to the default value.

Table 34: Default External Route Preferences

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes

Route Type	Preference	Configurable
OSPF internal	10	Yes[2]
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes

Default

external-preference 150 — OSPF external routes have a default preference of 150.

Parameters

preference

The preference for external routes expressed as a decimal integer.

Values 1 to 255

Platforms

7705 SAR Gen 2

external-preference

Syntax

external-preference *preference*

no external-preference

Context

[Tree] (config>router>isis>level external-preference)

Full Context

configure router isis level external-preference

Description

This command configures the external route preference for the IS-IS level.

[2] Preference for OSPF internal routes is configured with the **preference** command.

The **external-preference** command configures the preference level of either IS-IS level 1 or IS-IS level 2 external routes. By default, the preferences are as listed in the table below.

A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference decides the route to use.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is dependent on the default preference table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of the route to use is determined by the configuration of the **ecmp** in the **config>router** context.

The **no** form of this command reverts to the default value.

Default

external-preference (Level 1) — 160

external-preference (Level 2) — 165

Parameters

preference

Specifies the preference for external routes at this level as expressed.

Default preferences are listed in the following table.

Table 35: Default External Route Preferences

Route Type	Preference	Configurable
Direct attached	0	—
Static-route	5	Yes
OSPF internal routes	10	—
IS-IS Level 1 internal	15	Yes[3]
IS-IS Level 2 internal	18	Yes[3]
OSPF external	150	Yes
IS-IS Level 1 external	160	Yes
IS-IS Level 2 external	165	Yes
BGP	170	Yes

Values 1 to 255

[3] Internal preferences are changed using the preference command in the **config>router>isis>level** *level-number* context.

Platforms

7705 SAR Gen 2

external-preference

Syntax

external-preference *preference*

no external-preference

Context

[Tree] (config>router>ospf3 external-preference)

[Tree] (config>router>ospf external-preference)

Full Context

configure router ospf3 external-preference

configure router ospf external-preference

Description

This command configures the preference for OSPF external routes.

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in Table 36: Route Preference Defaults by Route Type . If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the **config**>**router** context.

The **no** form of this command reverts to the default value.

Default

external-preference 150

Parameters

preference

Specifies the preference for external routes expressed as a decimal integer. Defaults for different route types are listed in Table 36: Route Preference Defaults by Route Type .

Table 36: Route Preference Defaults by Route Type

Route Type	Preference	Configurable
Direct attached	0	No

Route Type	Preference	Configurable
Static routes	5	Yes
OSPF internal	10	Yes[4]
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

Values 1 to 255

Platforms

7705 SAR Gen 2

9.138 extranet

extranet

Syntax

extranet [**detail**]

no extranet

Context

[Tree] (debug>router>pim extranet)

Full Context

debug router pim extranet

Description

This command enables debugging for extranet PIM.

The **no** form of this command disables PIM extranet debugging.

[4] Preference for OSPF internal routes is configured with the **preference** command.

Parameters

detail

Debugs detailed extranet PIM information.

Platforms

7705 SAR Gen 2

10 f Commands

10.1 facility

facility

Syntax

facility *syslog-facility*

no facility

Context

[Tree] (config>service>vprn>log>syslog facility)

Full Context

configure service vprn log syslog facility

Description

This command configures the facility code for messages sent to the syslog target host.

Multiple syslog IDs can be created with the same target host but each syslog ID can only have one facility code. If multiple facility codes are entered, the last *facility-code* entered overwrites the previous facility-code.

If multiple facilities need to be generated for a single syslog target host, then multiple **log-id** entries must be created, each with its own filter criteria to select the events to be sent to the syslog target host with a given facility code.

The **no** form of this command reverts to the default value.

Default

local7 — Syslog entries are sent with the local7 facility code.

Parameters

syslog-facility

Specifies syslog facility name represents a specific numeric facility code. The code should be entered in accordance with the syslog RFC. However, the software does not validate if the facility code configured is appropriate for the event type being sent to the syslog target host.

Values kernel, user, mail, systemd, auth, syslogd, printer, netnews, uucp, cron, authpriv, ftp, ntp, logaudit, logalert, cron2, local0, local1, local2, local3, local4, local5, local6, local7

Valid responses per RFC 3164, *The BSD syslog Protocol,* are listed in Table 37: Syslog Facility Codes.

Table 37: Syslog Facility Codes

Numerical Code	Facility Code
0	kernel
1	user
2	mail
3	systemd
4	auth
5	syslogd
6	printer
7	net-news
8	uucp
9	cron
10	auth-priv
11	ftp
12	ntp
13	log-audit
14	log-alert
15	cron2
16	local0
17	local1
18	local2
19	local3
20	local4
21	local5
22	local6
23	local7

Values: 0 to 23

Platforms

7705 SAR Gen 2

facility

Syntax

facility *syslog-facility*

no facility

Context

[Tree] (config>log>syslog facility)

Full Context

configure log syslog facility

Description

This command configures the facility code for messages sent to the syslog target host.

Multiple syslog IDs can be created with the same target host but each syslog ID can only have one facility code. If multiple facility codes are entered, the last *facility-code* entered overwrites the previous facility-code.

If multiple facilities need to be generated for a single syslog target host, then multiple **log-id** entries must be created, each with its own filter criteria to select the events to be sent to the syslog target host with a given facility code.

The **no** form of this command reverts to the default value.

Default

facility local7

Parameters

syslog-facility

Specifies a syslog facility name which represents a specific numeric facility code. The code must be entered in accordance with the syslog RFC. However, the software does not validate if the facility code configured is appropriate for the event type being sent to the syslog target host.

Values kernel, user, mail, systemd, auth, syslogd, printer, netnews, uucp, cron, authpriv, ftp, ntp, logaudit, logalert, cron2, local0, local1, local2, local3, local4, local5, local6, local7

Valid responses per RFC 3164, *The BSD syslog Protocol,* are listed in Table 38: Syslog Protocol Valid Responses.

Table 38: Syslog Protocol Valid Responses

Numerical Code	Facility Code
0	kernel
1	user
2	mail
3	systemd
4	auth
5	syslogd
6	printer
7	net-news
8	uucp
9	cron
10	auth-priv
11	ftp
12	ntp
13	log-audit
14	log-alert
15	cron2
16	local0
17	local1
18	local2
19	local3
20	local4
21	local5
22	local6
23	local7

Platforms

7705 SAR Gen 2

10.2 fail-on-error

fail-on-error

Syntax

[**no**] **fail-on-error**

Context

[Tree] (config>card fail-on-error)

Full Context

configure card fail-on-error

Description

This command controls the behavior of the card when any one of a specific set of card level errors is encountered in the system. When the **fail-on-error** command is enabled, and any one (or more) of the specific errors is detected, then the Operational State of the card is set to Failed. This Failed state will persist until the clear card command is issued (reset) or the card is removed and re-inserted (re-seat). If the condition persists after re-seating the card, then Nokia support should be contacted for further investigation.

Enabling **fail-on-error** is only recommended when the network is designed to be able to route traffic around a failed card (redundant cards, nodes or other paths exist).

The list of specific errors includes:

- CHASSIS event ID# 2063 – tmnxEqCardPChipMemoryEvent
- CHASSIS event ID# 2076 – tmnxEqCardPChipCamEvent
- CHASSIS event ID# 2059 – tmnxEqCardPChipError (for ingress Ethernet only)
- CHASSIS event ID# 2098 – tmnxEqCardQChipBufMemoryEvent
- CHASSIS event ID# 2099 – tmnxEqCardQChipStatsMemoryEvent
- CHASSIS event ID# 2101 – tmnxEqCardQChipIntMemoryEvent
- CHASSIS event ID# 2103 – tmnxEqCardChipIfCellEvent

On platforms without independent IOM/IMM and CPM cards, the node is rebooted if fail-on-error is enabled and one of the card level errors is encountered.

The tmnxEqCardPChipError is only considered as a trigger for card fail-on-error for ingress FCS errors (not egress FCS errors), and only for Ethernet MDAs or IMMs.

Note that upon the detection of the event/error in the system, the reporting of the event (logs) and the **fail-on-error** behavior of the card are independent. Log event control configuration will determine whether the events are reported in logs (or SNMP traps, and so on) and the **fail-on-error** configuration will determine the behavior of the card. This implies that the card can be configured to **fail-on-error** even if the events are suppressed (some may be suppressed in the system by default). In order to facilitate post-failure analysis, Nokia recommends that you enable the reporting of the specific events/errors (**configure log event-control**) when **fail-on-error** is enabled.

Default

no fail-on-error

Platforms

7705 SAR Gen 2

fail-on-error

Syntax

[**no**] **fail-on-error**

Context

[Tree] (config>card>mda fail-on-error)

Full Context

configure card mda fail-on-error

Description

This command enables the fail-on-error feature. If an MDA is experiencing too many Egress XPL Errors, this feature causes the MDA to fail. This can force an APS switchover or **traffic re-route**. The purpose of this feature is to avoid situations where traffic is forced to use a physical link that suffers from errors but is still technically operational.

The feature uses values configured in the **config>card>mda>egress-xpl** context. When this feature is enabled on a MDA, if *window* consecutive minutes pass in which the MDA experiences more than *threshold* Egress XPL Errors per minute, then the MDA will be put in the *failed* state.

The **no** form of this command disables the feature on the MDA.

Platforms

7705 SAR Gen 2

10.3 failed-threshold

failed-threshold

Syntax

failed-threshold [1 to 1000]

failed-threshold all

Context

[Tree] (config>service>vpls>site failed-threshold)

Full Context

configure service vpls site failed-threshold

Description

This command defines the number of objects should be down for the site to be declared down. Both administrative and operational status must be evaluated and if at least one is down, the related object is declared down.

Default

failed-threshold all

Parameters

1 to 1000

Specifies the threshold for the site to be declared down.

Platforms

7705 SAR Gen 2

10.4 failover

failover

Syntax

failover

Context

[Tree] (config>service>vprn>dhcp>server>pool failover)

[Tree] (config>router>dhcp6>server>pool failover)

[Tree] (config>router>dhcp>server>pool failover)

[Tree] (config>service>vprn>dhcp6>server failover)

[Tree] (config>service>vprn>dhcp6>server>pool failover)

[Tree] (config>router>dhcp6>server failover)

[Tree] (config>service>vprn>dhcp>server failover)

[Tree] (config>router>dhcp>server failover)

Full Context

configure service vprn dhcp local-dhcp-server pool failover

configure router dhcp6 local-dhcp-server pool failover

configure router dhcp local-dhcp-server pool failover

configure service vprn dhcp6 local-dhcp-server failover

configure service vprn dhcp6 local-dhcp-server pool failover

configure router dhcp6 local-dhcp-server failover

configure service vprn dhcp local-dhcp-server failover

configure router dhcp local-dhcp-server failover

Description

Commands in this context configure failover parameters.

Platforms

7705 SAR Gen 2

10.5 fallback-path-computation-method

fallback-path-computation-method

Syntax

fallback-path-computation-method {**none** | **local-cspf**}

no fallback-path-computation-method

Context

[Tree] (config>router>mpls>lsp fallback-path-computation-method)

[Tree] (config>router>mpls>lsp-template fallback-path-computation-method)

Full Context

configure router mpls lsp fallback-path-computation-method

configure router mpls lsp-template fallback-path-computation-method

Description

This command specifies the fallback path computation method used if all configured PCEs are down or the signaling overload and the redelegation timer has expired. This method is used regardless of whether the LSP is PCE-controlled and PCE-computed, or just PCE-computed.

The **no** form of this command removes the fallback path computation method used.

Default

fallback-path-computation-method none

Parameters

none

Specifies to fall back to using the named path for RSVP-TE LSPs.

local-cspf

Specifies to fall back to using local CSPF computation.

Platforms

7705 SAR Gen 2

10.6 family

family

Syntax

family *family*

Context

[Tree] (config>service>vprn>bgp>convergence family)

Full Context

configure service vprn bgp convergence family

Description

This command specifies the convergence family used for route convergence.

Parameters

family

Specifies the convergence family used for route convergence

Values ipv4, ipv6

Platforms

7705 SAR Gen 2

family

Syntax

[**no**] **family** {**ipv4** | **ipv6** | **label-ipv4** | **flow-ipv4** | **flow-ipv6**}

Context

[Tree] (config>service>vprn>bgp>group>graceful-restart>long-lived family)

[Tree] (config>service>vprn>bgp>graceful-restart>long-lived family)

[Tree] (config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived family)

Full Context

configure service vprn bgp group graceful-restart long-lived family

configure service vprn bgp graceful-restart long-lived family

configure service vprn bgp group neighbor graceful-restart long-lived family

Description

This command configures family-specific LLGR parameters for BGP peers.

Default

no family

Parameters

ipv4

Specifies the IPv4 family.

ipv6

Specifies the IPv6 family.

label-ipv4

Specifies the label IPv4 family.

flow-ipv4

Specifies the flow IPv4 family.

flow-ipv6

Specifies the flow IPv6 family.

Platforms

7705 SAR Gen 2

family

Syntax

family [**ipv4**] [**label-ipv4**] [**ipv6**] [**mcast-ipv4**] [**flow-ipv4**] [**mcast-ipv6**] [**flow-ipv6**]

no family

Context

[Tree] (config>service>vprn>bgp family)

[Tree] (config>service>vprn>bgp>group>neighbor family)

[Tree] (config>service>vprn>bgp>group family)

Full Context

configure service vprn bgp family

configure service vprn bgp group neighbor family

configure service vprn bgp group family

Description

This command configures the set of BGP address families (AFI plus SAFI) to be supported by the applicable VPRN BGP sessions.

The **no** form of this command restores the default, which is equivalent to configuring unlabeled IPv4 unicast routes (AFI 1, SAFI 1) only.

Default

no family

Parameters

ipv4

Keyword to advertise support for the IPv4 unicast (unlabeled) address family.

label-ipv4

Keyword to advertise support for the IPv4 unicast (labeled) address family.

ipv6

Keyword to advertise support for the IPv6 unicast (unlabeled) address family.

mcast-ipv4

Keyword to advertise support for the IPv4 multicast SAFI address family.

flow-ipv4

Keyword to advertise support for the IPv4 FlowSpec address family.

mcast-ipv6

Keyword to advertise support for the IPv6 multicast SAFI address family.

flow-ipv6

Keyword to advertise support for the IPv6 FlowSpec address family.

Platforms

7705 SAR Gen 2

family

Syntax

family [**ipv4** | **ipv6**]

Context

[Tree] (config>router>mpls>lsp-template family)

Full Context

configure router mpls lsp-template family

Description

This command specifies if the lsp-template is for use in IPv4 or IPv6 SR-TE LSP.

This command is optional in a IPv4 SR-TE auto-LSP but must be set to **ipv6** value in a IPv6 SR-TE auto-LSP. By default, this command is set to **ipv4** value for backward compatibility.

When establishing both IPv4 and IPv6 SR-TE mesh auto-LSPs with the same parameters and constraints, a separate LSP template of type **mesh-p2p-srte** must be configured for each address family with the **family** CLI leaf set to the IPv4 or IPv6 value. SR-TE one-hop auto-LSPs can only be established for either IPv4 or IPv6 family, but not both. The **family** leaf in the LSP template of type **one-hop-p2p-srte** should be set to the desired IP family value.

The **no** form of this command reverts to the default value.

Default

family ipv4

Parameters

ipv4

Specifies the lsp-template is for use in IPv4 SR-TE LSP.

ipv6

Specifies the lsp-template is for use in IPv6 SR-TE LSP.

Platforms

7705 SAR Gen 2

family

Syntax

family *family*

Context

[Tree] (config>router>bgp>convergence family)

Full Context

configure router bgp convergence family

Description

This command configures the IP family used for route convergence.

Parameters

family

Specifies the convergence family.

Values ipv4, ipv6, vpn-ipv4, vpn-ipv6, label-ipv4, label-ipv6

Platforms

7705 SAR Gen 2

family

Syntax

family [**ipv4**] [**label-ipv4**] [**vpn-ipv4**] [**ipv6**] [**label-ipv6**] [**vpn-ipv6**] [**mcast-ipv4**] [**l2-vpn**] [**mvpn-ipv4**] [**mvpn-ipv6**] [**mdt-safi**] [**ms-pw**] [**flow-ipv4**] [**flow-ipv6**] [**route-target**] [**mcast-vpn-ipv4**] [**evpn**] [**bgp-ls**] [**mcast-ipv6**] [**mcast-vpn-ipv6**] [**sr-policy-ipv4**] [**sr-policy-ipv6**] [**flow-vpn-ipv4**] [**flow-vpn-ipv6**]

no family

Context

[Tree] (config>router>bgp family)

[Tree] (config>router>bgp>group>neighbor family)

[Tree] (config>router>bgp>group family)

Full Context

configure router bgp family

configure router bgp group neighbor family

configure router bgp group family

Description

This command configures the set of BGP address families (AFI/SAFI) to be supported by the base router BGP sessions.

The **no** form of this command restores the default, which corresponds to unlabeled IPv4 unicast routes (AFI 1, SAFI 1) only.

Default

family ipv4

Parameters

ipv4

Keyword to advertise MP-BGP support for the IPv4 unicast (unlabeled) address family.

label-ipv4

Keyword to advertise MP-BGP support for the IPv4 unicast (labeled) address family.

vpn-ipv4

Keyword to advertise MP-BGP support for the IPv4 VPN (SAFI 128) address family.

ipv6

Keyword to advertise MP-BGP support for the IPv6 unicast (unlabeled) address family.

label-ipv6

Keyword to advertise MP-BGP support for the IPv6 unicast (labeled) address family.

vpn-ipv6

Keyword to advertise MP-BGP support for the IPv6 VPN (SAFI 128) address family.

mcast-ipv4

Keyword to advertise MP-BGP support for the IPv4 multicast SAFI address family.

l2-vpn

Keyword to advertise MP-BGP support for the L2 VPN address family.

mvpn-ipv4

Keyword to advertise MP-BGP support for the IPv4 multicast VPN address family.

mvpn-ipv6

Keyword to advertise MP-BGP support for the IPv6 multicast VPN address family.

mdt-safi

Keyword to advertise MP-BGP support for the MDT SAFI address family.

ms-pw

Keyword to advertise MP-BGP support for the multi-segment pseudowire address family.

flow-ipv4

Keyword to advertise MP-BGP support for the IPv4 FlowSpec address family.

flow-ipv6

Keyword to advertise MP-BGP support for the IPv6 FlowSpec address family.

route-target

Keyword to advertise MP-BGP support for RT constraint routes.

mcast-vpn-ipv4

Keyword to advertise MP-BGP support for the IPv4 VPN multicast (SAFI 129) address family.

evpn

Keyword to advertise MP-BGP support for the EVPN address family.

bgp-ls

Keyword to advertise MP-BGP support for the BGP-LS address family.

mcast-ipv6

Keyword to advertise MP-BGP support for the IPv6 multicast SAFI address family.

mcast-vpn-ipv6

Keyword to advertise MP-BGP support for the IPv6 multicast routes from a VPRN over the provider network. This family is only applicable in the base BGP routing context.

sr-policy-ipv4

Keyword to advertise MP-BGP support for AFI1/SAFI73 IP address families for BGP routes that encode a segment-routing policy to an IPv4 destination.

sr-policy-ipv6

Keyword to advertise MP-BGP support for AF12/SAF173 IP address families for BGP routes that encode a segment-routing policy to an IPv6 destination.

flow-vpn-ipv4

Keyword to advertise support for the FlowSpec-VPN IPv4 address family (AFI 1, SAFI 134).

flow-vpn-ipv6

Keyword to advertise support for the FlowSpec-VPN IPv6 address family (AFI 2, SAFI 134).

Platforms

7705 SAR Gen 2

family

Syntax

[**no**] **family** {**ipv4** | **ipv6** | **label-ipv4** | **label-ipv6** | **vpn-ipv4** | **vpn-ipv6** | **l2-vpn** | **route-target** | **flow-ipv4** | **flow-ipv6** | **flow-vpn-ipv4** | **flow-vpn-ipv6**}

Context

[Tree] (config>router>bgp>group>neighbor>graceful-restart>long-lived family)

[Tree] (config>router>bgp>graceful-restart>long-lived family)

[Tree] (config>router>bgp>group>graceful-restart>long-lived family)

Full Context

configure router bgp group neighbor graceful-restart long-lived family

configure router bgp graceful-restart long-lived family

configure router bgp group graceful-restart long-lived family

Description

This command configures family-specific LLGR parameters for BGP peers.

The **no** form of this command deletes the context.

Default

no family

Parameters

ipv4

Keyword to specify the IPv4 family.

ipv6

Keyword to specify the IPv6 family.

label-ipv4

Keyword to specify the label IPv4 family.

label-ipv6
Keyword to specify the label IPv6 family.

vpn-ipv4
Keyword to specify the VPN IPv4 family.

vpn-ipv6
Keyword to specify the VPN IPv6 family.

l2-vpn
Keyword to specify the Layer 2 VPN family.

route-target
Keyword to specify the route target family.

flow-ipv4
Keyword to specify the flow IPv4 family.

flow-ipv6
Keyword to specify the flow IPv6 family.

flow-vpn-ipv4
Keyword to specify the FlowSpec-VPN IPv4 address family.

flow-vpn-ipv6
Keyword to specify the FlowSpec-VPN IPv6 address family.

Platforms
7705 SAR Gen 2

family

Syntax
family {**label-ipv4** | **label-ipv6** | **vpn**}

Context
[Tree] (config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel family)

Full Context
configure router bgp next-hop-resolution labeled-routes transport-tunnel family

Description
This command configures the address family context for configuring next-hop resolution of BGP label routes.

Parameters
label-ipv4
Enters the context for configuring next-hop-resolution options for labeled-unicast IPv4 routes.

label-ipv6

Enters the context for configuring next-hop-resolution options for labeled-unicast IPv6 routes.

vpn

Enters the context for configuring next-hop-resolution options for VPN-IPv4 and VPN-IPv6 routes when they are not imported into any VPRN service.

Platforms

7705 SAR Gen 2

family

Syntax

family {**ipv4** | **ipv6**}

Context

[Tree] (config>router>bgp>next-hop-res>shortcut-tunnel family)

Full Context

configure router bgp next-hop-resolution shortcut-tunnel family

Description

This command creates the context to configure next-hop resolution of unlabeled IPv4 or unlabeled IPv6 routes by certain tunnel types in the tunnel table.

Parameters

ipv4

Specifies that the configuration applies to unlabeled IPv4 BGP routes.

ipv6

Specifies that the configuration applies to unlabeled IPv6 BGP routes.

Platforms

7705 SAR Gen 2

family

Syntax

family {**ipv4** | **ipv6** | **srv4** | **srv6**}

Context

[Tree] (config>router>isis>igp-shortcut>tunnel-next-hop family)

Full Context

configure router isis igp-shortcut tunnel-next-hop family

Description

Commands in this context configure the resolution of IGP IPv4 and IGP IPv6 prefix families, as well as SR-ISIS IPv4 and SR-ISIS IPv6 tunnel families using IGP shortcuts.

Parameters

ipv4

Selects the IPv4 address family.

ipv6

Selects the IPv6 address family.

srv4

Selects the SR-ISIS IPv4 tunnel family.

srv6

Selects the SR-ISIS IPv6 tunnel family.

Platforms

7705 SAR Gen 2

family

Syntax

family {**ipv4** | **ipv6**}

no family

Context

[Tree] (config>router>isis>segment-routing>adjacency-set family)

Full Context

configure router isis segment-routing adjacency-set family

Description

This command specifies the address family of an adjacency set in IS-IS.

The **no** form of this command reverts to the default.

Default

family ipv4

Parameters

ipv4

Specifies a family of IPv4.

ipv6

Specifies a family of IPv6.

Platforms

7705 SAR Gen 2

family

Syntax

family {**ipv4** | **srv4**}

Context

[Tree] (config>router>ospf>igp-shortcut>tunnel-next-hop family)

Full Context

configure router ospf igp-shortcut tunnel-next-hop family

Description

Commands in this context configure the resolution of the IGP IPv4 prefix family or SR-OSPF IPv4 tunnel using IGP shortcuts.

Parameters

ipv4

Selects the IPv4 address family.

srv4

Selects the SR-OSPF IPv4 tunnel family.

Platforms

7705 SAR Gen 2

family

Syntax

family ipv6

Context

[Tree] (config>router>ospf3>igp-shortcut>tunnel-next-hop family)

Full Context

configure router ospf3 igp-shortcut tunnel-next-hop family

Description

Commands in this context configure the resolution of the IGP IPv6 prefix family using IGP shortcuts.

Parameters

ipv6

Selects the IPv6 address family.

Platforms

7705 SAR Gen 2

family

Syntax

family [**ipv4**] [**label-ipv4**] [**vpn-ipv4**] [**ipv6**] [**label-ipv6**] [**vpn-ipv6**] [**mcast-ipv4**] [**l2-vpn**] [**mvpn-ipv4**] [**mvpn-ipv6**] [**mdt-safi**] [**ms-pw**] [**flow-ipv4**] [**flow-ipv6**] [**route-target**] [**mcast-vpn-ipv4**] [**evpn**] [**bgp-ls**] [**mcast-ipv6**] [**mcast-vpn-ipv6**] [**sr-policy-ipv4**] [**sr-policy-ipv6**] [**flow-vpn-ipv4**] [**flow-vpn-ipv6**]

no family

Context

[Tree] (config>router>policy-options>policy-statement>entry>from family)

Full Context

configure router policy-options policy-statement entry from family

Description

This command specifies address families as matching conditions.

The **no** form of the command configures the router to use the default value.

Default

no family

Parameters

ipv4

Keyword to match routes belonging to the IPv4 unicast (unlabeled) address family.

label-ipv4

Keyword to match routes belonging to the IPv4 unicast (labeled) address family.

vpn-ipv4

Keyword to match routes belonging to the IPv4 VPN (SAFI 128) address family.

ipv6

Keyword to match routes belonging to the IPv6 unicast (unlabeled) address family.

label-ipv6

Keyword to match routes belonging to the IPv6 unicast (labeled) address family.

vpn-ipv6

Keyword to match routes belonging to the IPv6 VPN (SAFI 128) address family.

mcast-ipv4

Keyword to match routes belonging to the IPv4 multicast SAFI address family.

l2-vpn

Keyword to match routes belonging to the L2 VPN address family.

mvpn-ipv4

Keyword to match routes belonging to the IPv4 multicast VPN address family.

mvpn-ipv6

Keyword to match routes belonging to the IPv6 multicast VPN address family.

mdt-safi

Keyword to match routes belonging to the MDT SAFI address family.

ms-pw

Keyword to match routes belonging to the multi-segment pseudowire address family.

flow-ipv4

Keyword to match routes belonging to the IPv4 FlowSpec address family.

flow-ipv6

Keyword to match routes belonging to the IPv6 FlowSpec address family.

route-target

Keyword to match routes belonging to the address family for RT constrain routes.

mcast-vpn-ipv4

Keyword to match routes belonging to the IPv4 VPN multicast (SAFI 129) address family.

evpn

Keyword to match routes belonging to the EVPN address family.

bgp-ls

Keyword to advertise the BGP-LS address family to the associated BGP neighbors.

mcast-ipv6

Keyword to match routes belonging to the IPv6 multicast SAFI address family.

mcast-vpn-ipv6

Keyword to match routes belonging to the IPv6 multicast routes from a VPRN over the provider network. This family is only applicable in the base BGP routing context.

sr-policy-ipv4

Keyword to match routes belonging to the segment routing policy IPv4 address family (AFI1/SAFI73).

sr-policy-ipv6

Keyword to match routes belonging to the segment routing policy IPv6 address family (AFI2/SAFI73).

flow-vpn-ipv4

Keyword to match routes belonging to the FlowSpec-VPN IPv4 address family (AFI 1, SAFI 134).

flow-vpn-ipv6

Keyword to match routes belonging to the FlowSpec-VPN IPv6 address family (AFI 2, SAFI 134).

Platforms

7705 SAR Gen 2

10.7 far-end

far-end

Syntax

far-end *ip-address* [**vc-id** *vc-id*] [{**ing-svc-label** *ingress-vc-label* | **tldp**}] [**icb**]

no far-end *ip-address*

Context

[Tree] (config>mirror>mirror-dest>remote-source far-end)

Full Context

configure mirror mirror-dest remote-source far-end

Description

This command is used on a destination router in a remote mirroring solution. See the description for the **remote-source** command for additional information.

When using L2TPv3, MPLS-TP or LDP IPv6 LSP SDPs in the remote mirroring solution, the destination node should be configured with **remote-src>spoke-sdp** entries. For all other types of SDPs, r**emote-source>far-end** entries are used.

Up to 50 far-end entries can be specified.

The **no** form of this command removes the IP address from the remote source configuration.

Parameters

ip-address

Specifies the service IP address (system IP address) of the remote device sending mirrored traffic to this mirror destination service. If 0.0.0.0 is specified, any remote is allowed to send to this service.

Values 1.0.0.1 to 223.255.255.254

vc-id

Specifies the virtual circuit identifier of the remote source. For mirror services, the *vc-id* defaults to the *service-id*. However, if the *vc-id* is being used by another service a unique vc-id is required to create an SDP binding. For this purpose the mirror service SDP bindings accepts *vc-ids*. This VC ID must match the VC ID used on the spoke SDP that is configured on the source router.

ingress-vc-label

Specifies the ingress service label for mirrored service traffic on the **far end** device for manually configured mirror service labels.

The defined *ing-svc-label* is entered into the ingress service label table which causes ingress packet with that service label to be handled by this mirror destination service.

The specified *ing-svc-label* must not have been used for any other service ID and must match the egress service label being used on the spoke SDP that is configured on the source router. It must be within the range specified for manually configured service labels defined on this router. It may be reused for other far end addresses on this *mirror-dest-service-id*.

Values 2048 to 18431

tldp

Specifies that the label is obtained through signaling via the LDP.

icb

Specifies that the remote source is an inter-chassis backup SDP binding.

Platforms

7705 SAR Gen 2

far-end

Syntax

far-end node-id *node-id* [**global-id** *global-id*]

far-end [*ip-address* | *ipv6-address*]

no far-end *ip-address* | *ipv6-address*

Context

[Tree] (config>service>sdp far-end)

Full Context

configure service sdp far-end

Description

This command configures the system IP address of the far-end destination router for the service destination point (SDP) that is the termination point for a service.

The far-end IP address must be explicitly configured. The destination IP address must be that of an SR OS and for a GRE SDP it must match the system IP address of the far end router.

If the SDP uses GRE for the destination encapsulation, the IP address is checked against other GRE SDPs to verify uniqueness. If the IP address is not unique within the configured GRE SDPs, an error is generated and the IP address is not associated with the SDP. The local device may not know whether the IP address is actually a system IP interface address on the far-end device.

If the SDP uses MPLS encapsulation, the **far-end** address is used to check LSP names when added to the SDP. If the " **to** IP address" defined within the LSP configuration does not exactly match the SDP **far-end** address, the LSP will not be added to the SDP and an error will be generated. Alternatively, an SDP that uses MPLS can have an MPLS-TP node with an MPLS-TP node-id and (optionally) a global ID. In this case, the SDP must use an MPLS-TP LSP and the SDP **signaling** parameter must be set to **off**.

An SDP cannot be administratively enabled until a **far-end** *ip-address* or MPLS-TP node-id is defined. The SDP is operational when it is administratively enabled (**no shutdown**) and the **far-end** *ip-address* is contained in the IGP routing table as a host route. OSPF ABRs should not summarize host routes between areas. This can cause SDPs to become operationally down. Static host routes (direct and indirect) can be defined in the local device to alleviate this issue.

On a tunnel configured as SDP with delivery type of eth-gre-bridged, this command designates L2oGRE tunnel end points. This is the only configuration option allowed for this type of SDP.

The **no** form of this command removes the currently configured destination IP address for the SDP. The *ip-address* parameter is not specified and will generate an error if used in the **no far-end** command. The SDP must be administratively disabled using the **config service sdp shutdown** command before the **no far-end** command can be executed. Removing the far-end IP address will cause all *lsp-name* associations with the SDP to be removed.

Parameters

far-end

Specifies the far-end termination point for the GRE tunnel.

ip-address* | *ipv6-address

Specifies a IPv4 or IPv6 address of the far-end SR OS for the SDP in dotted decimal notation.

node-id

Specifies the MPLS-TP Node ID of the far-end system for the SDP, either in dotted decimal notation (a.b.c.d) or an unsigned 32-bit integer (1 to 4294967295). This parameter is mandatory for an SDP using an MPLS-TP LSP.

global-id

Specifies a MPLS-TP Global ID of the far-end system for the SDP, in an unsigned 32-bit integer (0 to 4294967295). This parameter is optional for an SDP using an MPLS-TP LSP. If not entered, a default value for the Global ID of '0' is used. A global ID of '0' indicates that the far-end node is in the same domain as the local node. The user must explicitly configure a Global ID if its value is non-zero.

Platforms

7705 SAR Gen 2

10.8 fast-leave

fast-leave

Syntax

[**no**] **fast-leave**

Context

[Tree] (config>service>vpls>sap>mld-snooping fast-leave)

[Tree] (config>service>vpls>sap>igmp-snooping fast-leave)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping fast-leave)

[Tree] (config>service>vpls>spoke-sdp>mld-snooping fast-leave)

[Tree] (config>service>vpls>mesh-sdp>mld-snooping fast-leave)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping fast-leave)

Full Context

configure service vpls sap mld-snooping fast-leave

configure service vpls sap igmp-snooping fast-leave

configure service vpls spoke-sdp igmp-snooping fast-leave

configure service vpls spoke-sdp mld-snooping fast-leave

configure service vpls mesh-sdp mld-snooping fast-leave

configure service vpls mesh-sdp igmp-snooping fast-leave

Description

This command enables fast leave.

When IGMP fast leave processing is enabled, the 7705 SAR Gen 2 immediately removes a SAP or SDP from the IP multicast group when it detects an IGMP leave message on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a leave message from the forwarding table without first sending out group-specific queries to the SAP or SDP, which speeds up the process of changing channels.

Fast leave should only be enabled when there is a single receiver present on the SAP or SDP.

When fast leave is enabled, the configured **last-member-query-interval** value is ignored.

Default

no fast-leave

Platforms

7705 SAR Gen 2

fast-leave

Syntax

[**no**] **fast-leave**

Context

[Tree] (config>service>pw-template>igmp-snooping fast-leave)

Full Context

configure service pw-template igmp-snooping fast-leave

Description

This command enables fast leave.

When IGMP fast leave processing is enabled, the 7705 SAR Gen 2 will immediately remove a SAP or SDP from the IP multicast group when it detects an IGMP **leave** on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a **leave** from the forwarding table without first sending out group-specific queries to the SAP or SDP, and thus speeds up the process of changing channels (zapping).

Fast leave should only be enabled when there is a single receiver present on the SAP or SDP.

When fast leave is enabled, the configured last-member-query-interval value is ignored.

Default

no fast-leave

Platforms

7705 SAR Gen 2

10.9 fast-reroute

fast-reroute

Syntax

fast-reroute [**backup-sr-tunnel**]

no fast-reroute

Context

[Tree] (config>router>ldp fast-reroute)

Full Context

configure router ldp fast-reroute

Description

This command enables LDP Fast-Reroute (FRR) procedures. When enabled, LDP uses both the primary next-hop and LFA next-hop, when available, for resolving the next-hop of an LDP FEC against the corresponding prefix in the routing table. This will result in LDP programming a primary NHLFE and a backup NHLFE into the forwarding engine for each next-hop of a FEC prefix for the purpose of forwarding packets over the LDP FEC.

When any of the following events occurs, LDP instructs in the fast path the forwarding engines to enable the backup NHLFE for each FEC next-hop impacted by this event:

- An LDP interface goes operationally down, or is admin shutdown.
- An LDP session to a peer went down as the result of the Hello or Keep-Alive timer expiring.
- The TCP connection used by a link LDP session to a peer went down, due say to next-hop tracking of the LDP transport address in RTM, which brings down the LDP session.
- A BFD session, enabled on a T-LDP session to a peer, times-out and as a result the link LDP session to the same peer and which uses the same TCP connection as the T-LDP session goes also down.
- A BFD session enabled on the LDP interface to a directly connected peer, times out and brings down the link LDP session to this peer.

The **tunnel-down-dump-time** option or the **label-withdrawal-delay** option, when enabled, does not cause the corresponding timer to be activated for a FEC as long as a backup NHLFE is still available.

Because LDP can detect the loss of a neighbor/next-hop independently, it is possible that it switches to the LFA next-hop while IGP is still using the primary next-hop. Also, when the interface for the previous primary next-hop is restored, IGP may re-converge before LDP completed the FEC exchange with it neighbor over that interface. This may cause LDP to de-program the LFA next-hop from the FEC and blackhole traffic. In order to avoid this situation, it is recommended to enable IGP-LDP synchronization on the LDP interface.

When the SPF computation determines there is more than one primary next-hop for a prefix, it will not program any LFA next-hop in RTM. Thus, the LDP FEC will resolve to the multiple primary next-hops that provide the required protection.

The **backup-sr-tunnel** option enables the use of SR tunnel, as a remote LFA or TI-LFA backup tunnel next-hop by an LDP FEC.

As a pre-requisite, the user must enable the stitching of LDP and SR in the LDP-to-SR direction. That is because the LSR must perform the stitching of the LDP ILM to SR tunnel when the primary LDP next-hop of the FEC fails. Thus LDP must listen to SR tunnels programmed by the IGP in TTM but the mapping server feature is not required.

Assuming the following:

- the **backup-sr-tunnel** option is enabled in LDP
- the {**loopfree-alternates remote-lfa**} and/or the {**loopfree-alternates ti-lfa**} option is enabled in the IGP instance
- LDP was able to resolve the primary next-hop of the LDP FEC in RTM

IGP SPF will run both the base LFA and the TI-LFA algorithms and if it does not find a backup next-hop for a prefix of an LDP FEC, it will also run the remote LFA algorithm. If IGP finds a TI-LFA or a remote LFA tunnel next-hop, LDP programs the primary next-hop of the FEC using a LDP NHLFE and programs the LFA backup next-hop using a LDP NHLFE pointing to the SR tunnel endpoint. Note that the LDP packet is not "tunneled" over the SR tunnel. The LDP label is actually stitched to the segment routing label stack. LDP points both the LDP ILM and the LTN to the backup LDP NHLFE which itself uses the SR tunnel endpoint.

The behavior of the feature is thus similar to the LDP-to-SR stitching feature, except the behavior is augmented to allow the stitching of an LDP ILM/LTN to a SR tunnel also when the primary LDP next-hop of the FEC fails.

If the LDP FEC primary next-hop failed and LDP has pre-programmed a remote LFA or TI-LFA next-hop with a LDP backup NHLFE pointing to SR tunnel, the LDP ILM/LTN switches to it. Note that if for some reason the failure impacted only the LDP tunnel primary next-hop but not the SR tunnel primary next-hop, the LDP backup NHLFE will effectively point to the primary next-hop of the SR tunnel and traffic of the LDP ILM/LTN will follow this path instead of the TI-LFA or remote LFA next-hop of the SR tunnel until the latter is activated.

This feature is limited to IPv4 /32 prefixes in both LDP and SR.

The **no** form of this command disables the use of SR tunnels as backups for LDP FECs and disables LDP FRR.

Default

no fast-reroute

Platforms

7705 SAR Gen 2

fast-reroute

Syntax

fast-reroute *frr-method*

no fast-reroute

Context

[Tree] (config>router>mpls>lsp fast-reroute)

[Tree] (config>router>mpls>lsp-template fast-reroute)

Full Context

configure router mpls lsp fast-reroute

configure router mpls lsp-template fast-reroute

Description

This command creates a pre-computed detour LSP from each node in the path of the LSP. In case of failure of a link or LSP between two nodes, traffic is immediately rerouted on the pre-computed detour LSP, thus avoiding packet-loss.

When **fast-reroute** is enabled, each node along the path of the LSP tries to establish a detour LSP as follows:

- Each upstream node sets up a detour LSP that avoids only the immediate downstream node, and merges back on to the actual path of the LSP as soon as possible.

  If it is not possible to set up a detour LSP that avoids the immediate downstream node, a detour can be set up to the downstream node on a different interface.

- The detour LSP may take one or more hops (see **config>router>mpls>lsp hop-limit**, **config>router>mpls>lsp>primary-p2mp-instance hop-limit**) before merging back on to the main LSP path.
- When the upstream node detects a downstream link or node failure, the ingress router switches traffic to a standby path if one was set up for the LSP.

Fast reroute is available only for the primary path. No configuration is required on the transit hops of the LSP. The ingress router will signal all intermediate routers using RSVP to set up their detours. TE must be enabled for fast-reroute to work.

If an LSP is configured with **fast-reroute** *frr-method* specified but does not enable CSPF, then global revertive will not be available for the LSP to recover.

The **no** form of the **fast-reroute** command removes the detour LSP from each node on the primary path. This command will also remove configuration information about the hop-limit and the bandwidth for the detour routes.

The **no** form of **fast-reroute hop-limit** command reverts to the default value.

Note:
A one-to-one detour backup LSP cannot be used at the PLR for ABR node protection. As a result, a PLR node does not signal a one-to-one detour LSP for ABR protection. In addition, the ABR node rejects a Path message that it has received from a third-party implementation configured with a detour object and a loose ERO next-hop. The Path message is rejected regardless of whether the **cspf-on-loose-hop** command is enabled on the node. When the router transits ABR for the detour path, the router rejects the signaling of an inter-area detour backup LSP.

Default

no fast-reroute — When fast-reroute is specified, the default fast-reroute method is one-to-one.

Parameters

frr-method

Configures the fast-reroute method.

Values **one-to-one** — In the one-to-one technique, a label switched path is established which intersects the original LSP somewhere downstream of the point of link or node failure. For each LSP which is backed up, a separate backup LSP is established.

Values **facility** — This option, sometimes called many-to-one, takes advantage of the MPLS label stack. Instead of creating a separate LSP for every backed-up LSP, a single LSP is created which serves to backup up a set of LSPs. This LSP tunnel is called a bypass tunnel.

The bypass tunnel must intersect the path of the original LSP(s) somewhere downstream of the point of local repair (PLR). Naturally, this constrains the set of LSPs being backed-up through that bypass tunnel to those that pass through a common downstream node. All LSPs which pass through the PLR and through this common node which do not also use the facilities involved in the bypass tunnel are candidates for this set of LSPs.

Platforms

7705 SAR Gen 2

10.10 fc

fc

Syntax

fc *fc-name* **class-type** *ct-number*

no fc fc-name

Context

[Tree] (config>router>rsvp>diffserv-te fc)

Full Context

configure router rsvp diffserv-te fc

Description

This command maps one or more system forwarding classes to a Diff-Serv Class Type (CT). The default mapping is shown in Table 39: Forwarding Classes Mapping.

Table 39: Forwarding Classes Mapping

FC ID	FC Name	FC Designation	Class Type (CT)
7	Network Control	NC	7
6	High-1	H1	6
5	Expedited	EF	5
4	High-2	H2	4
3	Low-1	L1	3
2	Assured	AF	2
1	Low-2	L2	1
0	Best Effort	BE	0

The **no** form of this command reverts to the default mapping for the forwarding class name.

Parameters

class-type *ct-number*

The Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.

Values 0 to 7

Platforms

7705 SAR Gen 2

fc

Syntax

fc *fc-name*

no fc

Context

[Tree] (config>mirror>mirror-dest fc)

Full Context

configure mirror mirror-dest fc

Description

This command specifies a forwarding class for all mirrored packets transmitted to the destination SAP or SDP overriding the default (be) forwarding class. All packets are sent with the same class of service to minimize out-of-sequence issues. The mirrored packet does not inherit the forwarding class of the original packet.

When the destination is on a SAP, a single egress queue is created that pulls buffers from the buffer pool associated with the *fc-name*.

When the destination is on an SDP, the *fc-name* defines the DiffServ-based egress queue that is used to reach the destination. The *fc-name* also defines the encoded forwarding class of the encapsulation.

The FC configuration also affects how mirrored packets are treated at the ingress queuing point on the line cards. One ingress queue is used per mirror destination (service) and that is an expedited queue if the configured FC is expedited (one of nc, h1, ef or h2). The ingress mirror queues have no CIR, but a line-rate PIR.

The **no** form of this command reverts the mirror-dest service ID forwarding class to the default forwarding class.

Default

The best effort (be) forwarding class is associated with the mirror-dest service ID.

Parameters

fc-name

The name of the forwarding class with which to associate mirrored service traffic. The forwarding class name must already be defined within the system. If the *fc-name* does not exist, an error is returned and the **fc** command has no effect. If the *fc-name* does exist, the forwarding class associated with *fc-name* overrides the default forwarding class.

Values be, l2, af, l1, h2, ef, h1, nc

Platforms

7705 SAR Gen 2

fc

Syntax

fc *fc-name*

no fc

Context

[Tree] (config>saa>test>type-multi-line>lsp-ping fc)

[Tree] (config>saa>test>type-multi-line>lsp-ping>sr-policy fc)

Full Context

configure saa test type-multi-line lsp-ping fc

configure saa test type-multi-line lsp-ping sr-policy fc

Description

This command specifies the FC and profile parameters that are used to indicate the forwarding class and profile of the MPLS echo request packet.

The **no** form of this command reverts to the default value.

Default

fc be

Parameters

fc-name

Specifies the forwarding class name.

Values be, l2, af, l1, h2, ef, h1, nc

Default be

Platforms

7705 SAR Gen 2

fc

Syntax

fc *fc-name*

no fc

Context

[Tree] (config>oam-pm>session>ip fc)

Full Context

configure oam-pm session ip fc

Description

This command sets the forwarding class designation for TWAMP Light packets that are sent through the node and exposed to the various QoS functions on the network element.

The **no** form of this command restores the default value.

Default

fc be

Parameters

fc-name

Specifies the forwarding class name.

Values be, l2, af, l1, h2, ef, h1, nc

Default be

Platforms

7705 SAR Gen 2

fc

Syntax

fc *fc-name* [**create**]

no fc *fc-name*

Context

[Tree] (config>qos>sap-ingress fc)

Full Context

configure qos sap-ingress fc

Description

The **fc** command creates a class or subclass instance of the forwarding class fc-name. When the *fc-name* is created, classification actions can be applied and the subclass can be used in match classification criteria. Attempting to use an undefined subclass in a classification command will result in an execution error and the command will fail.

The **no** form of this command removes all the explicit queue mappings for *fc-name* forwarding types. The queue mappings revert to the default queues for *fc-name*. To successfully remove a subclass, all associations with the subclass in the classification commands within the policy must first be removed or diverted to another forwarding class or subclass.

Parameters

fc-name

The parameter subclass-name is optional and must be defined using a dot separated notation with a preceding valid system-wide forwarding class name. Creating a subclass follows normal naming conventions. Up to sixteen ASCII characters may be used. If the same sub-name is used with two or more forwarding class names, each is considered a different instance of subclass. A subclass must always be specified with its preceding forwarding class name. When a forwarding class is created or specified without the optional subclass, the parent forwarding class is assumed.

Within the SAP ingress QoS policy, up to 56 subclasses may be created. Each of the 56 subclasses may be created within any of the eight parental forwarding classes. When the limit of 56 is reached, any further subclass creations will fail and the subclass will not exist.

Successfully creating a subclass places the CLI within the context of the subclass for further subclass parameter definitions. Within the subclass context, commands may be executed that define subclass priority (within the parent forwarding class queue mapping), subclass color aware profile settings, subclass in-profile and out-of-profile precedence or DSCP markings.

The *subclass-name* parameter is optional and used with the *fc-name* parameter to define a pre-existing subclass. The *fc-name* and *subclass-name* parameters must be separated by a period (.). If *subclass-name* does not exist in the context of *fc-name*, an error will occur. If *subclass-name* is removed using the **no fc** *fc-name.subclass-name* **force** command, the **default-fc** command will automatically drop the *subclass-name* and only use *fc-name* (the parent forwarding class for the subclass) as the forwarding class.

Values

fc: *class*[.*subclass*]

class: be, l2, af, l1, h2, ef, h1, nc

subclass: 29 characters max

create

Required parameter when creating a SAP QoS ingress policy forwarding class.

Platforms

7705 SAR Gen 2

fc

Syntax

fc *fc-name* [**create**]

no fc *fc-name*

Context

[Tree] (config>qos>sap-egress fc)

Full Context

configure qos sap-egress fc

Description

The **fc** fc-*name* node within the SAP egress QoS policy is used to contain the explicitly defined queue mapping and dot1p marking commands for *fc-name*. When the mapping for *fc-name* points to the default queue and the dot1p marking is not defined, the node for *fc-name* is not displayed in the **show configuration** or **save configuration** output unless the detail option is specified.

The **no** form of this command removes the explicit queue mapping and dot1p marking commands for *fc-name*. The queue mapping reverts to the default queue for *fc-name* and the dot1p marking (if appropriate) uses the default of 0.

Default

no fc

Parameters

fc-name

This parameter specifies that the forwarding class queue mapping or dot1p marking is to be edited. The value given for *fc-name* must be one of the predefined forwarding classes in the system.

Values be, l2, af, l1, h2, ef, h1, nc

Platforms

7705 SAR Gen 2

fc

Syntax

fc *fc-name*

no fc

Context

[Tree] (config>qos>network>ingress fc)

Full Context

configure qos network ingress fc

Description

This command is used to enter the CLI node to configure QoS parameters for the specified forwarding class. The **fc** command overrides the default parameters for that forwarding class from the values defined in the network default policy.

The **no** form of this command removes the forwarding class name configuration. The forwarding class reverts to the parameters defined in the default network policy.

Parameters

fc-name

The case-sensitive, system-defined forwarding class name for which policy entries will be created.

Values be, l2, af, l1, h2, ef, h1, nc

Platforms

7705 SAR Gen 2

fc

Syntax

fc *fc-name*

no fc

Context

[Tree] (config>qos>network>egress fc)

Full Context

configure qos network egress fc

Description

This command is used to enter the CLI node to configure QoS parameters for the specified forwarding class. The FC name represents a CLI parent node that contains parameters describing the egress marking criteria of packets flowing through it. This command overrides the default parameters for that forwarding class from the values defined in the network default policy. It can also be used to redirect packets to a policer or queue in a network egress queue group instance.

The **no** form of this command removes the forwarding class name configuration. The forwarding class reverts to the parameters defined in the default network policy.

Parameters

fc-name

The case-sensitive, system-defined forwarding class name for which policy entries will be created.

Values be, l2, af, l1, h2, ef, h1, nc

Platforms

7705 SAR Gen 2

fc

Syntax

fc *fc-name* [**create**]

no fc *fc-name*

Context

[Tree] (config>qos>network-queue fc)

Full Context

configure qos network-queue fc

Description

The **fc** command is used to enter the forwarding class mapping context for the given *fc-name*. Each forwarding class maps by default to queues 1 (unicast) and 9 (multipoint).

Parameters

fc-name

A valid forwarding class must be specified as *fc-name* when the **fc** command is executed. When the **fc** *fc-name* command is successfully executed, the system will enter the specified forwarding class context where the **queue** *queue-id* command may be executed.

Values be, l2, af, l1, h2, ef, h1, nc

create

Required parameter when creating an FC node.

Platforms

7705 SAR Gen 2

fc

Syntax

fc *fc-name* [**create**]

no fc *fc-name*

Context

[Tree] (cfg>qos>qgrps>egr>qgrp fc)

Full Context

configure qos queue-group-templates egress queue-group fc

Description

The **fc** command is used to enter the forwarding class mapping context for the given fc-name. Each forwarding class has a default mapping depending on the egress queue group template. The system-created policer-output-queue template contains queues 1 and 2 by default with queue 1 being best-effort and queue 2 expedited. Forwarding classes be, l1, af and l2 all map to queue 1 by default. Forwarding classes h1, ef, h2 and nc all map to queue 2 by default. More queues may be created within the policer-output-queues template and the default forwarding classes may be changed to any defined queue within the template.

When all other user-defined egress queue group templates are created, only queue 1 (best-effort) exists and all forwarding classes are mapped to that queue. Other queues may be created and the forwarding classes may be changed to any defined queue within the template.

Besides the default mappings within the templates, the egress queue group template forwarding class queue mappings operate the same as the forwarding class mappings in a sap-egress QoS policy.

The template forwarding class mappings are the default mechanism for mapping egress policed traffic to a queue within an egress port queue group associated with the template. If a queue-id is explicitly specified in the QoS policy forwarding class policer mapping, and that queue exists within the queue group, the template forwarding class mapping is ignored.

Egress policed subscriber traffic works in a slightly different way. The subscriber and subscriber host support destination and organization strings are used to identify the egress port queue group. In this instance, the forwarding class mappings are always used and any queue overrides in the QoS policy are ignored. If neither string exists for the subscriber host, the egress queue group queue-id can be derived from either the QoS policy policer mapping or the template forwarding class queue mappings.

The **no** form of this command is used to return the specified forwarding class to its default template queue mapping.

Parameters

fc-name

A valid forwarding class must be specified as *fc-name* when the **fc** command is executed. When the **fc** *fc-name* command is successfully executed, the system will enter the specified forwarding class context where the **queue** *queue-id* command may be executed.

Values **be**, **l1**, **af**, **l2**, **h1**, **ef**, **h2**, **nc**

Platforms

7705 SAR Gen 2

10.11 fdb-table-high-wmark

fdb-table-high-wmark

Syntax

[**no**] **fdb-table-high-wmark** *high-water-mark*

Context

[Tree] (config>service>template>vpls-template fdb-table-high-wmark)

[Tree] (config>service>vpls fdb-table-high-wmark)

Full Context

configure service template vpls-template fdb-table-high-wmark

configure service vpls fdb-table-high-wmark

Description

This command specifies the value to send logs and traps when the threshold is reached.

The **no** form of this command reverts to the default value.

Default

fdb-table-high-wmark 95

Parameters

high-water-mark

Specifies the value as a percentage.

Values 0 to 100

Platforms

7705 SAR Gen 2

10.12 fdb-table-low-wmark

fdb-table-low-wmark

Syntax

[**no**] **fdb-table-low-wmark** *low-water-mark*

Context

[Tree] (config>service>template>vpls-template fdb-table-low-wmark)

[Tree] (config>service>vpls fdb-table-low-wmark)

Full Context

configure service template vpls-template fdb-table-low-wmark

configure service vpls fdb-table-low-wmark

Description

This command specifies the value to send logs and traps when the threshold is reached.

The **no** form of this command reverts to the default value.

Default

fdb-table-low-wmark 90

Parameters

low-water-mark

Specifies the value as a percentage.

Values 0 to 100

Platforms

7705 SAR Gen 2

10.13 fdb-table-size

fdb-table-size

Syntax

fdb-table-size *table-size*

no fdb-table-size [*table-size*]

Context

[Tree] (config>service>vpls fdb-table-size)

[Tree] (config>service>template>vpls-template fdb-table-size)

Full Context

configure service vpls fdb-table-size

configure service template vpls-template fdb-table-size

Description

This command specifies the maximum number of MAC entries in the forwarding database (FDB) for the VPLS instance on this node.

The **fdb-table-size** specifies the maximum number of forwarding database entries for both learned and static MAC addresses for the VPLS instance.

The **no** form of this command returns the maximum FDB table size to default.

Default

fdb-table-size 250

Parameters

table-size

Specifies the number of entries permitted in the forwarding database for this VPLS instance.

Values 7705 SAR Gen 2: 1 to 32767

Platforms

7705 SAR Gen 2

fdb-table-size

Syntax

fdb-table-size *table-size*

no fdb-table-size

Context

[Tree] (config>service>system fdb-table-size)

Full Context

configure service system fdb-table-size

Description

This command configures the maximum system FDB table size, which is dependent on the chassis type. CPMs with at least 16 GB of memory are required when exceeding 500k MAC addresses in a system. The table size cannot be reduced below its default value, which is also chassis-dependent.

The maximum system FDB table size also limits the maximum FDB table size of any card within the system.

The **no** version of this command sets the table size to its default.

The command default depends on the chassis type and available memory.

Parameters

table-size

Specifies the maximum system FDB table size.

Values 32767 to 32767

Platforms

7705 SAR Gen 2

10.14 fec-limit

fec-limit

Syntax

fec-limit *limit* [**log-only**] [**threshold** *percentage*]

no fec-limit

Context

[Tree] (config>router>ldp>session-params>peer fec-limit)

Full Context

configure router ldp session-parameters peer fec-limit

Description

This command configures a limit on the number of FECs which an LSR will accept from a given peer and add into the LDP label database. The limit applies to the aggregate count of all FEC types including service FEC. Once the limit is reached, any FEC received will be released back to the peer. This behavior is different from the per-peer import policy which will still accept the FEC into the label database but will not resolve it.

When the FEC limit for a peer is reached, the LSR performs the following actions:

1. Generates a trap and a syslog message.

2. Generates a LDP notification message with the LSR overload status TLV, for each LDP FEC type including service FEC, to this peer only if this peer advertised support for the LSR overload sub-TLV via the LSR Overload Protection Capability TLV at session initialization.
3. Releases, with LDP Status Code of "No_Label_Resources", any new FEC, including service FEC, from this peer which exceeds the limit.

If a legitimate FEC is released back to a peer, while the FEC limit was exceeded, the user must have a means to replay that FEC back to the router LSR once the condition clears. This is done automatically if the peer is an SR OS-based router and supports the LDP overload status TLV (SR OS 11.0R5 and higher). Third-party peer implementations must support the LDP overload status TLV or provide a manual command to replay the FEC.

The **threshold** option allows to set a threshold value when a trap and an syslog message are generated as a warning to the user in addition to when the limit is reached. The default value for the threshold when not configured is 90%.

The **log-only** option causes a trap and syslog message to be generated when reaching the threshold and limit. However, LDP labels are not released back to the peer.

If the user decreases the limit value such that it is lower than the current number of FECs accepted from the peer, the LDP LSR raises the trap for exceeding the limit. In addition, it will set overload for peers which signaled support for LDP overload protection capability TLV. However, no existing resolved FECs from the peer which does not support the overload protection capability TLV should be de-programmed or released.

A different trap is released when crossing the threshold in the upward direction, when reaching the FEC limit, and when crossing the threshold in the downward direction. However the same trap will not be generated more often than 2 minutes apart if the number of FECs oscillates around the threshold or the FEC limit.

Default

no fec-limit

Parameters

limit

Specifies the aggregate count of FECs of all types which can be accepted from this LDP peer.

log-only

Specifies that only a trap and syslog message are generated when reaching the threshold and limit. However, LDP labels are not released back to the peer.

percentage

Specifies the threshold value (as a percentage) that triggers a warning syslog message and trap to be sent.

Platforms

7705 SAR Gen 2

10.15 fec-originate

fec-originate

Syntax

fec-originate *ip-prefix/mask* [**advertised-label** *in-label*] [**swap-label** *out-label*] **interface interface-name**

fec-originate *ip-prefix/mask* [**advertised-label** *in-label*] **next-hop** *ip-address* [**swap-label** *out-label*]

fec-originate *ip-prefix/mask* [**advertised-label** *in-label*] **next-hop** *ip-address* [**swap-label** *out-label*] **interface interface-name**

fec-originate *ip-prefix/mask* [**advertised-label** *in-label*] **pop**

no fec-originate *ip-prefix/mask* **interface interface-name**

no fec-originate *ip-prefix/mask* **next-hop** *ip-address*

no fec-originate *ip-prefix/mask* **next-hop** *ip-address* **interface interface-name**

no fec-originate *ip-prefix/mask* **pop**

Context

[Tree] (config>router>ldp fec-originate)

Full Context

configure router ldp fec-originate

Description

This command defines a way to originate a FEC (with a swap action) for which the LSR is not egress, or to originate a FEC (with a pop action) for which the LSR is egress.

Parameters

ip-prefix/mask

Specifies information for the specified IP prefix and mask length.

Values

ipv4-prefix - a.b.c.d	
ipv4-prefix-le - [0..32]	
ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0..FFFF]H
	d - [0..255]D
	ipv6-prefix-le - [0..128]

next-hop

Specifies the IP address of the next hop of the prefix.

advertised-label

Specifies the label advertised to the upstream peer. If not configured, then the label advertised should be from the label pool. If the configured static label is not available then the IP prefix is not advertised.

out-label

Specifies the LSR to swap the label. If configured, then the LSR should swap the label with the configured swap-label. If not configured, then the default action is pop if the next-hop parameter is not defined.

The next-hop, advertised-label, swap-label parameters are all optional. If next-hop is configured but no swap label specified, it will be a swap with label 3, such as, pop and forward to the next-hop. If the next-hop and swap-label are configured, then it is a regular swap. If no parameters are specified, a pop and route is performed.

Values 16 to 1048575

in-label

Specifies the number of labels to send to the peer associated with this FEC.

Values 32 to 1023

pop

Specifies to pop the label and transmit without the label.

interface *interface-name*

Specifies the name of the interface the label for the originated FEC is swapped to. For an unnumbered interface, this parameter is mandatory since there is no address for the next-hop. For a numbered interface, it is optional.

Platforms

7705 SAR Gen 2

10.16 fec-type-capability

fec-type-capability

Syntax

fec-type-capability

Context

[Tree] (config>router>ldp>if-params>if>ipv4 fec-type-capability)

[Tree] (config>router>ldp>if-params>if>ipv6 fec-type-capability)

[Tree] (config>router>ldp>session-params>peer fec-type-capability)

Full Context

configure router ldp interface-parameters interface ipv4 fec-type-capability

configure router ldp interface-parameters interface ipv6 fec-type-capability

configure router ldp session-parameters peer fec-type-capability

Description

This command enables or disables the advertisement of a FEC type on a given LDP session or Hello adjacency to a peer.

Platforms

7705 SAR Gen 2

10.17 fec129-cisco-interop

fec129-cisco-interop

Syntax

[**no**] **fec129-cisco-interop**

Context

[Tree] (config>router>ldp>session-params>peer fec129-cisco-interop)

Full Context

configure router ldp session-parameters peer fec129-cisco-interop

Description

This command specifies whether LDP will provide translation between non-compliant FEC 129 formats of Cisco. Peer LDP sessions must be manually configured towards the non-compliant Cisco PEs.

When enabled, Cisco non-compliant format will be used to send and interpret received label release messages that is the FEC129 SAII and TAII fields will be reversed.

When the disabled, Cisco non-compliant format will not be used or supported. Peer address has to be the peer LSR-ID address.

The **no** form of this command returns the default.

Default

no fec129-cisco-interop

Platforms

7705 SAR Gen 2

10.18 fib-priority

fib-priority

Syntax

fib-priority {**high** | **standard**}

Context

[Tree] (config>service>vprn fib-priority)

Full Context

configure service vprn fib-priority

Description

This command specifies the FIB priority for VPRN BGP routes.

Parameters

high

Specifies high FIB priority for VPRN.

standard

Specifies standard FIB priority for VPRN.

Platforms

7705 SAR Gen 2

fib-priority

Syntax

fib-priority {**high** | **standard**}

Context

[Tree] (config>router fib-priority)

Full Context

configure router fib-priority

Description

This command specifies the FIB priority for VPRN BGP routes.

Default

fib-priority standard

Parameters

high

Specifies the high FIB priority.

standard

Specifies the standard FIB priority.

Platforms

7705 SAR Gen 2

10.19 fib-telemetry

fib-telemetry

Syntax

[**no**] **fib-telemetry**

Context

[Tree] (config>router fib-telemetry)

Full Context

configure router fib-telemetry

Description

This command enables the collection of extra state information related to the forwarding table state of certain IP routes, TTM tunnels, and MPLS LFIB entries. This extra state can be retrieved by gNMI telemetry subscriptions targeted to the following YANG paths:

- /state/router/route-fib
- /state/router/tunnel-fib
- /state/router/label-fib

If this command is not configured, no information is displayed by the following **show** commands:

- **show>router>fib-telemetry>route**
- **show>router>fib-telemetry>tunnel**

The **no** form of this command disables the collection of this extra state.

Default

no fib-telemetry

Platforms

7705 SAR Gen 2

10.20 file

file

Syntax

file

Context

[Tree] (file)

Full Context

file

Description

Specifies the context to enter and perform file system operations. When entering the **file** context, the prompt changes to reflect the present working directory. Navigating the file system with the **cd ..** command results in a changed prompt.

The **exit all** command leaves the file system/file operation context and returns to the operational root CLI context. The state of the present working directory is maintained for the CLI session. Entering the **file** command returns the cursor to the working directory where the **exit** command was issued.

Platforms

7705 SAR Gen 2

10.21 file-id

file-id

Syntax

[**no**] **file-id** *file-id* [**name** *file-policy-name*]

Context

[Tree] (config>log file-id)

Full Context

configure log file-id

Description

This command creates the context to configure a file policy that is used as the destination for an event log or billing (accounting) file.

This command defines the file location and characteristics that are to be used as the destination for a log event message stream or accounting/billing information. The file defined in this context is subsequently specified in the **to** command under **log-id** or **accounting-policy** to direct specific logging or billing source streams to the file destination.

A file policy can only be assigned to either *one* **log-id** or *one* **accounting-policy**. It cannot be reused for multiple instances. A file policy and associated file definition must exist for each log and billing file that must be stored in the file system.

A file is created when the file policy defined in this command is selected as the destination type for a specific log or accounting record. Log files are collected in a "log" directory. Accounting files are collected in an "act" directory.

The file names for a log are created by the system as summarized in Table 40: Log File Names.

Table 40: Log File Names

File Type	File Name
Log File	log*llff-timestamp*
Accounting File	act*aaff-timestamp*

Where:

- *ll* is the *log-id*
- *aa* is the accounting *policy-id*
- *ff* is the file-id
- The *timestamp* is the actual timestamp when the file is created. The format for the timestamp is *yyyymmdd-hhmmss* where:
  - *yyyy* is the year (for example, 2006)
  - *mm* is the month number (for example, 12 for December)
  - *dd* is the day of the month (for example, 03 for the 3rd of the month)
  - *hh* is the hour of the day in 24 hour format (for example, 04 for 4 a.m.)
  - *mm* is the minutes (for example, 30 for 30 minutes past the hour)
  - *ss* is the number of seconds (for example, 14 for 14 seconds)
- The accounting file is compressed and has a gz extension.

When initialized, each file contains:

- The *log-id* description.
- The time the file was opened.
- The reason the file was created.
- If the event log file was closed properly, the sequence number of the last event stored on the log is recorded.

If the process of writing to a log file fails (for example, the compact flash card is full) and if a backup location is not specified or fails, the log file will not become operational even if the compact flash card is replaced. Enter either a **clear log** command or a **shutdown/no shutdown** command to reinitialize the file.

If the primary location fails (for example, the compact flash card fills up during the write process), a trap is sent and logging continues to the specified backup location. This can result in truncated files in different locations.

The **no** form of this command removes the file policy from the configuration. A file policy can only be removed from the configuration if the policy is not the designated output for a log destination. The actual log or accounting file remain on the file system when a file policy is deleted.

Parameters

file-id

The file identification number for the file policy, expressed as a decimal integer.

Values 1 to 99

name file-policy-name

Configures an optional file policy name, up to 64 characters, that can be used to refer to the file policy after it is created. If the name begins with a numerical digit (from 1 to 9), the name is a number from 1 to 99.

Platforms

7705 SAR Gen 2

10.22 file-storage-control

file-storage-control

Syntax

file-storage-control

Context

[Tree] (config>log file-storage-control)

Full Context

configure log file-storage-control

Description

Commands in this context configure the total size limit of log and accounting files on each storage device on the active CPM.

Platforms

7705 SAR Gen 2

10.23 file-transmission-profile

file-transmission-profile

Syntax

file-transmission-profile *name* [**create**]

no file-transmission-profile

Context

[Tree] (config>system file-transmission-profile)

Full Context

configure system file-transmission-profile

Description

This command creates a new file transmission profile or enters the configuration context of an existing file-transmission-profile.

The **file-transmission-profile** context defines transport parameters for protocol such as HTTP, include routing instance, source address, timeout value, and so on.

The **no** form of the command removes the profile name from the configuration.

Default

no file-transmission-profile

Parameters

name

Specifies the file transmission profile name, up to 32 characters.

create

Keyword used to create the transmission profile. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

file-transmission-profile

Syntax

file-transmission-profile *profile-name*

no file-transmission-profile

Context

[Tree] (config>system>security>pki>ca-prof>auto-crl-update>crl-urls>url-entry file-transmission-profile)

Full Context

configure system security pki ca-profile auto-crl-update crl-urls url-entry file-transmission-profile

Description

This command specifies the file-transmission-profile for the **url-entry**. When the system downloads a CRL from the configured URL in the **url-entry** it will use the transportation parameter configured in the **file-transmission-profile**. **auto-crl-update** supports Base/Management/VPRN routing instance. **vpls-management** is not supported. In case of VPRN, the HTTP server port can only be 80 or 8080.

The **no** form of this command removes the specified profile name.

Default

no file-transmission-profile

Parameters

profile-name

Specifies the name of the file transmission profile to be matched up to 32 characters. The profile name is configured in the **config>system>file-transmission-profile** context.

Platforms

7705 SAR Gen 2

10.24 file-url

file-url

Syntax

file-url *file-url*

no file-url

Context

[Tree] (config>mirror>mirror-dest>pcap file-url)

Full Context

configure mirror mirror-dest pcap file-url

Description

This command specifies a file URL for the FTP or TFTP server, including the filename for packet capture transfer. After the file URL is entered, the system attempts to establish a connection and creates a file using the filename specified. The command prompt displays an error and rejects the file URL if the

session establishment fails, if write privilege to remote server fails, or if the session experiences a sudden termination. If the FTP or TFTP server is unreachable, the command prompt is halted for further input until the retires are timed out after 24 seconds (after four attempts of about six seconds each). This command overwrites any file on the FTP or TFTP server with the same filename.

The **no** form of this command removes the *file-url* instance and stops the packet capture and file transfer session.

Parameters

file-url

Specifies the URL for the file to direct the search.

Values [*local-url* | *remote-url*]

where:

- *local-url* — [*cflash-id/*] [*file-path*]
  180 chars max, including *cflash-id*
  directory length 99 chars max each
- *remote-url* — [{ftp://| tftp://} *login*:*pswd*@*remote-locn*/][*file-path*]
  180 chars max
  directory length 99 chars max each
  where: *remote-locn* — [*hostname* | *ipv4-address* | *ipv6-address*]

ipv4-address	a.b.c.d
ipv6-address	x:x:x:x:x:x:x:x[-interface]
	x:x:x:x:x:x:d.d.d.d[-interface]
	x - [0..FFFF]H
	d - [0..255]D
	interface - 32 chars max, for link local addresses
cflash-id	cf1:| cf1-A:| cf1-B:| cf2:| cf2-A:| cf2-B:| cf3:| cf3-A:| cf3-B:

Platforms

7705 SAR Gen 2

10.25 filter

filter

Syntax

filter ip *ip-filter-id*

filter ipv6 *ipv6-filter-id*

no filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]

Context

[Tree] (config>service>vprn>nw-if>egress filter)

[Tree] (config>service>vprn>if>spoke-sdp>ingress filter)

[Tree] (config>service>vprn>if>spoke-sdp>egress filter)

Full Context

configure service vprn network-interface egress filter

configure service vprn interface spoke-sdp ingress filter

configure service vprn interface spoke-sdp egress filter

Description

This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface. An IP filter policy can be associated with spoke SDPs. Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria.

The filter command is used to associate a filter policy with a specified ip-filter-id with an ingress or egress SAP. The ip-filter-id must already be defined before the filter command is executed. If the filter policy does not exist, the operation fails and an error message returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use scope command within the filter definition to change the scope to local or global. The default scope of a filter is local.

Parameters

ip *ip-filter-id*

Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 to 65535

Platforms

7705 SAR Gen 2

filter

Syntax

filter ip *ip-filter-id*

filter ipv6 *ipv6-filter-id*

filter mac *mac-filter-id*

no filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]

Context

[Tree] (config>service>ies>if>sap>egress filter)

[Tree] (config>service>ies>if>sap>ingress filter)

Full Context

configure service ies interface sap egress filter

configure service ies interface sap ingress filter

Description

This command associates a filter policy with an ingress or egress Service Access Point (SAP). Filter policies control the forwarding and dropping of packets based on the matching criteria. MAC filters are only allowed on Epipe and Virtual Private LAN Service (VPLS) SAPs.

The **filter** command is used to associate a filter policy with a specified *ip-filter-id* or *ipv6-filter-id* with an ingress or egress SAP. The filter policy must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation fails and an error message returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to the match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Parameters

ip *ip-filter-id*

Specifies the ID for the IP filter policy and corresponds to a previously created IP filter policy in the **config>filter>ip-filter** context.

Values 1 to 65535

ipv6 *ipv6-filter-id*

Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

mac *mac-filter-id*

Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 to 65535

Platforms

7705 SAR Gen 2

filter

Syntax

filter ip *ip-filter-id*

filter ipv6 *ipv6-filter-id*

filter mac *mac-filter-id*

no filter [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]

Context

[Tree] (config>service>vpls>sap>ingress filter)

[Tree] (config>service>vpls>spoke-sdp>egress filter)

[Tree] (config>service>vpls>sap>egress filter)

[Tree] (config>service>vpls>spoke-sdp>ingress filter)

[Tree] (config>service>vpls>mesh-sdp>ingress filter)

[Tree] (config>service>vpls>mesh-sdp>egress filter)

Full Context

configure service vpls sap ingress filter

configure service vpls spoke-sdp egress filter

configure service vpls sap egress filter

configure service vpls spoke-sdp ingress filter

configure service vpls mesh-sdp ingress filter

configure service vpls mesh-sdp egress filter

Description

This command associates an IP filter policy or MAC filter policy with an ingress or egress Service Access Point (SAP) or IP interface.

Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria. There are two types of filter policies: IP and MAC. Only one type may be applied to a SAP at a time.

The **filter** command is used to associate a filter policy with a specified *filter ID* with an ingress or egress SAP. The *filter ID* must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Parameters

ip *ip-filter-id*

Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 to 65535

ipv6 *ipv6-filter-id*

Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

mac *mac-filter-id*

Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 to 65535

Platforms

7705 SAR Gen 2

filter

Syntax

filter [**ip** *ip-filter-id*]

filter [**ipv6** *ipv6-filter-id*]

filter [**mac** *mac-filter-id*]

no filter [**ip** *ip-filter-id*]

no filter [**ipv6** *ipv6-filter-id*]

no filter [**mac** *mac-filter-id*]

Context

[Tree] (config>service>epipe>sap>ingress filter)

[Tree] (config>service>epipe>spoke-sdp>ingress filter)

[Tree] (config>service>epipe>spoke-sdp>egress filter)

[Tree] (config>service>epipe>sap>egress filter)

Full Context

configure service epipe sap ingress filter

configure service epipe spoke-sdp ingress filter

configure service epipe spoke-sdp egress filter

configure service epipe sap egress filter

Description

This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface.

Filter policies control the forwarding and dropping of packets based on IP matching criteria. Only one filter can be applied to a SAP at a time.

The **filter** command is used to associate a filter policy with a specified *filter-id* with an ingress or egress SAP. The *filter-id* must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

IPv6 filters are only supported by the 7705 SAR Gen 2 but are not supported on a Layer 2 SAP that is configured with QoS MAC criteria. Also, MAC filters are not supported on a Layer 2 SAP that is configured with QoS IPv6 criteria.

Parameters

ip-filter-id

Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 to 65535

ipv6-filter-id

Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

mac-filter-id

Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 to 65535

Platforms

7705 SAR Gen 2

filter

Syntax

filter ip *ip-filter-id*

filter ipv6 *ipv6-filter-id*

no filter

Context

[Tree] (config>service>ies>if>spoke-sdp>ingress filter)

[Tree] (config>service>ies>if>spoke-sdp>egress filter)

Full Context

configure service ies interface spoke-sdp ingress filter

configure service ies interface spoke-sdp egress filter

Description

This command associates an IP filter policy filter policy with an ingress or egress spoke SDP.

Filter policies control the forwarding and dropping of packets based on matching criteria.

MAC filters are only allowed on Epipe and Virtual Private LAN Service (VPLS) SAPs.

The **filter** command is used to associate a filter policy with a specified *ip-filter-id* with an ingress or egress spoke SDP. The *ip-filter-id* must already be defined in the **config>filter** context before the **filter** command is executed. If the filter policy does not exist, the operation fails and an error message returned.

In general, filters applied to SAPs or spoke SDPs (ingress or egress) apply to all packets on the SAP or spoke SDPs. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Parameters

ip

Keyword indicating the filter policy is an IP filter.

ip-filter-id

The filter name acts as the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy. The filter ID must already exist within the created IP filters.

Platforms

7705 SAR Gen 2

filter

Syntax

filter ip *ip-filter-id*

filter ipv6 *ipv6-filter-id*

no filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]

Context

[Tree] (config>service>vprn>if>sap>egress filter)

[Tree] (config>service>vprn>if>sap>ingress filter)

Full Context

configure service vprn interface sap egress filter

configure service vprn interface sap ingress filter

Description

This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface. Filter policies control the forwarding and dropping of packets based on IP matching criteria.

The **filter** command is used to associate a filter policy with a specified filter ID with an ingress or egress SAP. The filter ID must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation fails and an error message returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Parameters

ip *ip-filter-id*

Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values ip-filter-id: 1 to 65535

name: up to 64 characters

ipv6 *ipv6-filter-id*

Specifies IPv6 filter policy. The filter ID must already exist within the created IP filters.

Values ip-filter-id: 1 to 65535

name: up to 64 characters

Platforms

7705 SAR Gen 2

filter

Syntax

filter *filter-id* [**name** *filter-name*]

no filter *filter-id*

Context

[Tree] (config>service>vprn>log filter)

[Tree] (config>service>vprn>log>log-id filter)

Full Context

configure service vprn log filter

configure service vprn log log-id filter

Description

This command creates a context for an event filter. An event filter specifies whether to forward or drop an event or trap based on the match criteria.

Filters are configured in the **filter** *filter-id* context and then applied to a log in the **log-id** *log-id* context. Only events for the configured log source streams destined to the log ID where the filter is applied are filtered.

Changes made to an existing filter using any of the sub-commands are immediately applied to the destinations where the filter is applied.

By default, no event filters are defined. Event filters must be explicitly configured.

The **no** form of this command removes the filter association from log IDs, which causes those logs to forward all events.

Default

No event filters are defined.

Parameters

filter-id

Specifies the unique filter ID.

Values 1 to 1500

name filter-name

Configures an optional filter name, up to 64 characters, that can be used to refer to the filter after it is created.

Platforms

7705 SAR Gen 2

filter

Syntax

filter ip *ip-filter-id*

filter ipv6 *ipv6-filter-id*

no filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]

Context

[Tree] (config>service>vprn>network>ingress filter)

Full Context

configure service vprn network ingress filter

Description

This command configures a network ingress filter for IPv4 or IPv6 traffic arriving over explicitly defined spokes or auto-bind network interfaces for the VPRN service.

The **no** form of this command removes an IPv4, IPv6, or both filters.

Default

no filter

Parameters

ip-filter-id/ipv6-filter-id

Specifies an existing IP/IPv6 filter policy of a scope template.

Values 1 to 65535, *name*

name: 64 characters maximum

Platforms

7705 SAR Gen 2

filter

Syntax

filter ip *ip-filter-id*

no filter

Context

[Tree] (config>service>vprn>ipmirrorif>spoke-sdp>ingress filter)

Full Context

configure service vprn ip-mirror-interface spoke-sdp ingress filter

Description

This command places a filter on the IP mirror interface spoke SDP. It is recommended to configure this filter with a PBR filter to redirect the mirror traffic to the proper egress interface.

The **no** form of this command removes the filter ID from the configuration.

Parameters

ip-filter-id

Specifies the IP filter ID.

Values 1 to 65525 or a name, up to 64 characters.

Platforms

7705 SAR Gen 2

filter

Syntax

filter ip *ip-filter-id*

filter ipv6 *ipv6-filter-id*

no filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]

Context

[Tree] (config>router>if>egress filter)

[Tree] (config>router>if>ingress filter)

Full Context

configure router interface egress filter

configure router interface ingress filter

Description

This command associates an IP filter policy with an IP interface.

Filter policies control packet forwarding and dropping based on IP match criteria.

The *ip-filter-id* must have been preconfigured before this **filter** command is executed. If the filter ID does not exist, an error occurs.

Only one filter ID can be specified.

The **no** form of this command removes the filter policy association with the IP interface.

Default

no filter

Parameters

ip-filter-id

The filter name acts as the ID for the IP filter policy expressed as a decimal integer. The filter policy must already exist within the **config>filter>ip** context.

Values 1 to 16384

ipv6-filter-id

The filter name acts as the ID for the IPv6 filter policy expressed as a decimal integer. The filter policy must already exist within the **config>filter>ipv6** context.

Values 1 to 65535

Platforms

7705 SAR Gen 2

filter

Syntax

filter ip *ip-filter-id*

filter ipv6 *ipv6-filter-id*

filter mac *mac-filter-id*

no filter [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]

Context

[Tree] (config>service>pw-template>ingress filter)

[Tree] (config>service>pw-template>egress filter)

Full Context

configure service pw-template ingress filter

configure service pw-template egress filter

Description

This command associates an IP filter policy or MAC filter policy on egress or ingress. Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria. There are two types of filter policies: IP and MAC. Only one type may be applied to a SAP at a time.

The **filter** command is used to associate a filter policy with a specified *filter ID* with an ingress or egress SAP. The *filter ID* must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

This command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **filter-name** command can be used in all configuration modes.

This command is mutually exclusive with the **filter-name** command. Only one or the other can be configured.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Parameters

ip-filter-id

Specifies the IP filter policy.

Values 1 to 65535

ipv6-filter-id

Specifies the IPv6 filter policy.

Values 1 to 65535

mac-filter-id

Specifies the MAC filter policy.

Values 1 to 65535

Platforms

7705 SAR Gen 2

filter

Syntax

filter *filter-id* [**name** *filter-name*]

no filter *filter-id*

Context

[Tree] (config>log filter)

Full Context

configure log filter

Description

This command creates a context for an event filter. An event filter specifies whether to forward or drop an event or trap based on the match criteria.

Filters are configured in the **filter** *filter-id* context and then applied to a log in the **log-id** *log-id* context. Only events for the configured log source streams destined to the log ID where the filter is applied are filtered.

Changes made to an existing filter using any of the sub-commands are immediately applied to the destinations where the filter is applied.

By default, no event filters are defined. Event filters must be explicitly configured.

The **no** form of this command removes the filter association from log IDs, which causes those logs to forward all events.

Parameters

filter-id

Specifies the unique filter ID.

Values 1 to 1500

name filter-name

Configures an optional filter name, up to 64 characters, that can be used to refer to the filter after it is created.

Platforms

7705 SAR Gen 2

filter

Syntax

filter *filter-id*

no filter

Context

[Tree] (config>log>log-id filter)

Full Context

configure log log-id filter

Description

This command adds an event filter policy with the log destination.

The **filter** command is optional. If an event filter is not configured, all events, alarms and traps generated by the source stream will be forwarded to the destination.

An event filter policy defines (limits) the events that are forwarded to the destination configured in the log-id. The event filter policy can also be used to select the alarms and traps to be forwarded to a destination **snmp-trap-group**.

The application of filters for debug messages is limited to application and subject only.

Accounting records cannot be filtered using the **filter** command.

Only one filter ID can be configured per log destination.

The **no** form of this command removes the specified event filter from the *log-id*.

Parameters

filter-id

Specifies the event filter policy ID is used to associate the filter with the *log-id* configuration. The event filter policy ID must already be defined in **config>log>filter** *filter-id*.

Values 1 **to** 1000

Platforms

7705 SAR Gen 2

10.26 filter-id-range

filter-id-range

Syntax

filter-id-range start *filter-id* **end** *filter-id*

no filter-id-range

Context

[Tree] (config>filter>md-auto-id filter-id-range)

Full Context

configure filter md-auto-id filter-id-range

Description

This command specifies the range of IDs used by SR OS to automatically assign an ID to filters that are created in model-driven interfaces without an ID explicitly specified by the user or client.

A filter created with an explicitly-specified ID cannot use an ID in this range. In classic CLI and SNMP, the ID range cannot be changed while objects exist inside the previous or new range. In MD interfaces, the range can be changed, which causes any previously existing objects in the previous ID range to be deleted and re-created using a new ID in the new range.

The **no** form of this command removes the range values.

See the **config>filter md-auto-id** command for further details.

Default

no filter-id-range

Parameters

start *filter-id*

Specifies the lower value of the ID range. The value must be less than or equal to the **end** value.

Values 1 to 2147483647

end *filter-id*

Specifies the upper value of the ID range. The value must be greater than or equal to the **start** value.

Values 1 to 2147483647

Platforms

7705 SAR Gen 2

10.27 filter-name

filter-name

Syntax

[**no**] **filter-name**

Context

[Tree] (config>service>template>vpls-sap-template>ingress filter-name)

[Tree] (config>service>template>vpls-sap-template>egress filter-name)

Full Context

configure service template vpls-sap-template ingress filter-name

configure service template vpls-sap-template egress filter-name

Description

Commands in this context configure filter parameters.

Platforms

7705 SAR Gen 2

filter-name

Syntax

filter-name ip *ip-name*

filter-name ipv6 *ipv6-name*

filter-name mac *mac-name*

no filter-name [**ip**] [**ipv6**] [**mac**]

Context

[Tree] (config>service>pw-template>ingress filter-name)

[Tree] (config>service>pw-template>egress filter-name)

Full Context

configure service pw-template ingress filter-name

configure service pw-template egress filter-name

Description

This command associates an IP filter policy or MAC filter policy on egress or ingress. Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria. There are two types of filter policies: IP and MAC. Only one type may be applied to a SAP at a time.

The **filter-name** command is used to associate a filter policy with a specified *filter name* with an ingress or egress SAP. The *filter name* must already be defined before the **filter-name** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

This command is mutually exclusive with the **filter** command. Only one or the other can be configured.

The **no** form of this command removes any configured filter name association with the SAP or IP interface. The filter name itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter name and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Parameters

ip-name

Specifies the IP filter policy. The filter name must already exist within the created IP filters, up to 64 characters.

ipv6-name

Specifies the IPv6 filter policy. The filter name must already exist within the created IPv6 filters, up to 64 characters.

mac-name

Specifies the MAC filter policy. The specified filter name must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters, up to 64 characters.

Platforms

7705 SAR Gen 2

10.28 filtering

filtering

Syntax

filtering *filtering-mode*

no filtering

Context

[Tree] (config>service>nat>nat-policy filtering)

Full Context

configure service nat nat-policy filtering

Description

This command configures the filtering of the NAT or residential firewall policy.

Default

filtering endpoint-independent

Parameters

filtering-mode

Specifies the method used to filter the inbound traffic.

Values address-and-port-dependent, endpoint-independent

Platforms

7705 SAR Gen 2

10.29 flags-tlv

flags-tlv

Syntax

[no] flags-tlv

Context

[Tree] (config>router>fad>flex-algo flags-tlv)

Full Context

configure router flexible-algorithm-definitions flex-algo flags-tlv

Description

This command advertises the FAD Flags TLV to provide additional context on how the router must run a constrained SPF (cSPF). The IETF definition includes only the M-flag for use in the FAD Flags TLV. When it is set, the M-flag specifies the use of a Flex-Algorithm specific prefix metric. The M-flag is important for inter-area or inter-domain routing support with Flex-Algorithms.

When a router advertises a FAD, it is optional to advertise the FAD Flags TLV. However, when a FAD that includes the FAD Flags TLV is received, then the router must decode the flags before participating in the Flex-Algorithm.

By default, the following considerations apply to the FAD Flags TLV.

- SR OS sets the M-flag and advertises the FAD Flags TLV.
- When a FAD Flags TLV is received, SR OS decodes the flags and modifies the cSPF computation based upon the M-flag status.

The **no** form of this command prevents the advertisement of the FAD Flags TLV within a FAD.

Default

flags-tlv

Platforms

7705 SAR Gen 2

10.30 flex-algo

flex-algo

Syntax

flex-algo *flex-algo*

no flex-algo

Context

[Tree] (config>router>static-route-entry>indirect>tunnel-next-hop flex-algo)

Full Context

configure router static-route-entry indirect tunnel-next-hop flex-algo

Description

This command instructs the tunnel towards the indirect static-route next-hop to use the specified flexible algorithm.

It is assumed that the router using this command is participating in the flexible algorithm. This command instructs the router to lookup the indirect next-hop using flexible algorithm tunnels. If flexible algorithm aware tunnel to the indirect next-hop does not exist, then the static-route is not activated.

The expected outcome of this command is that when the router receives an IP payload packet, that it is steered towards the indirect next-hop using a flexible algorithm aware segment-routing tunnel if such tunnel exists. If such tunnel does not exist, then the route is not active, and the received IP packet will be dropped, if no other Longest Prefix Match (LPM) route exists.

If the *flex-algo* parameter is specified, the resolution filter can only use matching flexible algorithm-aware segment routing tunnels created by flexible algorithm-aware routing protocols (for example, SR IS-IS).

The **no** form of this command disables flexible algorithm-aware indirect next-hop resolution.

Default

no flex-algo

Parameters

flex-algo

Configures or deconfigures tunnel-next-hop flexible algorithm for resolving indirect static-route-entry.

Values 128 to 255

Platforms

7705 SAR Gen 2

flex-algo

Syntax

flex-algo *fad-name* [**create**]

no flex-algo *fad-name*

Context

[Tree] (config>router>fad flex-algo)

Full Context

configure router flexible-algorithm-definitions flex-algo

Description

This command configures the definition context for a Flexible Algorithm Definition (FAD). Parameters, including the FAD priority, metric type, links to construct a flexible algorithm topology graph, and a description of the algorithm. Up to 256 local FADs can be configured on a router.

The FAD configuration parameters are grouped using the *fad-name* as the reference anchor. When an IGP is configured to use and advertise a local configured FAD, the *fad-name* is used as the reference anchor.

The **no** form of this command deletes the configured parameters and removes the defined FAD.

Default

no flex-algo

Parameters

fad-name

Specifies the name of the flexible algorithm, up to 32 characters, that is used as reference anchor for the configured parameters.

create

Specifies the mandatory keyword to create a router instance.

Platforms

7705 SAR Gen 2

flex-algo

Syntax

[**no**] **flex-algo** *flex-algo*

Context

[Tree] (config>router>isis>flex-algos flex-algo)

Full Context

configure router isis flexible-algorithms flex-algo

Description

This command enters the configuration context for an IS-IS flexible algorithm.

A maximum of seven unique flexible algorithms can be configured on a router across all configured IS-IS instances. In each IS-IS flexible algorithm configuration context, the IS-IS instance participation can be either enabled or disabled, and it configures the advertising of a locally-configured flexible algorithm definition.

When flexible algorithm is enabled in an IS-IS instance, it is enabled for all levels (Level 1 and Level 2) within the IS-IS instance.

The **no** form of this command removes the IS-IS flexible algorithm configuration context.

Default

no flex-algo

Parameters

flex-algo

Specifies the number of the IS-IS flexible algorithm.

Values 128 to 255

Platforms

7705 SAR Gen 2

flex-algo

Syntax

[**no**] **flex-algo** *flex-algo-id*

Context

[Tree] (config>router>ospf>flex-algos flex-algo)

Full Context

configure router ospf flexible-algorithms flex-algo

Description

This command enters the configuration context for an OSPFv2 flexible algorithm.

A maximum of seven unique flexible algorithms can be configured on a router across all configured OSPFv2 instances. The supported flexible algorithms are in the range of 128 to 255. In each OSPF flexible algorithm configuration context, the OSPFv2 instance participation can be either enabled or disabled, and it configures the advertising of a locally-configured flexible algorithm definition.

When flexible algorithm is enabled in an OSPF instance, it is enabled for all areas within the OSPF instance.

The **no** form of this command removes the OSPF flexible algorithm configuration context.

Default

no flex-algo

Parameters

flex-algo-id

Specifies the OSPF flexible algorithm number.

Values 128 to 255

Platforms

7705 SAR Gen 2

flex-algo

Syntax

[**no**] **flex-algo** *flex-algo-id*

Context

[Tree] (config>router>ospf>area>if flex-algo)

Full Context

configure router ospf area interface flex-algo

Description

This command enters the OSPFv2 flexible algorithms configuration context on the interface.

The **no** form of this command removes the OSPF flexible algorithm configuration context.

Default

no flex-algo

Parameters

flex-algo-id

Specifies the number of the OSPF flexible algorithm.

Values 128 to 255

Platforms

7705 SAR Gen 2

flex-algo

Syntax

flex-algo *flex-algo-id* | *param-name*

no flex-algo

Context

[Tree] (config>router>policy-options>policy-statement>entry>action flex-algo)

[Tree] (config>router>policy-options>policy-statement>default-action flex-algo)

Full Context

configure router policy-options policy-statement entry action flex-algo

configure router policy-options policy-statement default-action flex-algo

Description

This command configures the Flex-Algorithm for use in the BGP next-hop autobind operation in a BGP import policy. A Flex-Algorithm aware autobind of the BGP next-hop is enabled when the route is matched by the policy statement entry.

Note:

- Flex-Algorithm aware next-hop lookup is supported for unicast BGP, VPRN, and BGP-LU.
- This command is not supported for multicast address families.

The **no** form of this command removes the Flex-Algorithm aware next-hop lookup.

Default

no flex-algo

Parameters

flex-algo-id

Specifies the flexible algorithm forwarding path.

Values 128 to 255

param-name

Specifies the parameter name, up to 32 characters, that starts and ends with an at-sign (@) symbol.

Platforms

7705 SAR Gen 2

flex-algo

Syntax

flex-algo *flex-algo-id*

no flex-algo

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls>sr-isis flex-algo)

[Tree] (config>oam-pm>session>ip>tunnel>mpls>sr-ospf flex-algo)

Full Context

configure oam-pm session ip tunnel mpls sr-isis flex-algo

configure oam-pm session ip tunnel mpls sr-ospf flex-algo

Description

This command configures the flexible algorithm to tunnel IP packets for session tests.

The **no** form of this command disables the flexible algorithm to tunnel IP packets.

Default

no flex-algo

Parameters

flex-algo-id

Specifies the flexible algorithm ID.

Values 128 to 255

Platforms

7705 SAR Gen 2

10.31 flexible-algorithm-definitions

flexible-algorithm-definitions

Syntax

flexible-algorithm-definitions

Context

[Tree] (config>router flexible-algorithm-definitions)

Full Context

configure router flexible-algorithm-definitions

Description

Commands in this context locally configure algorithm definitions.

Platforms

7705 SAR Gen 2

10.32 flexible-algorithms

flexible-algorithms

Syntax

flexible-algorithms

Context

[Tree] (config>router>isis flexible-algorithms)

Full Context

configure router isis flexible-algorithms

Description

Commands in this context configure the IS-IS parameters for flexible algorithm participation.

Platforms

7705 SAR Gen 2

flexible-algorithms

Syntax

flexible-algorithms

Context

[Tree] (config>router>ospf flexible-algorithms)

Full Context

configure router ospf flexible-algorithms

Description

Commands in this context configure the OSPFv2 parameters for flexible algorithm participation.

Platforms

7705 SAR Gen 2

10.33 flood-garp-and-unknown-req

flood-garp-and-unknown-req

Syntax

[**no**] **flood-garp-and-unknown-req**

Context

[Tree] (config>service>ies>if>vpls>evpn>arp flood-garp-and-unknown-req)

[Tree] (config>service>vprn>if>vpls>evpn>arp flood-garp-and-unknown-req)

Full Context

configure service ies interface vpls evpn arp flood-garp-and-unknown-req

configure service vprn interface vpls evpn arp flood-garp-and-unknown-req

Description

This command controls whether CPM-originated ARP frames are flooded in the R-VPLS service. Any frames that are data path flooded, such as the ARP messages received on a SAP, are flooded regardless of the command.

The **no** form of this command disables flooding GARP and unknown requests.

Default

flood-garp-and-unknown-req

Platforms

7705 SAR Gen 2

10.34 flow-label

flow-label

Syntax

flow-label *flow-label* [*mask*]

no flow-label

Context

[Tree] (config>filter>ipv6-filter>entry>match flow-label)

Full Context

configure filter ipv6-filter entry match flow-label

Description

This command configures the flow-label and optional mask match condition.

The **no** form of the command reverts to the default.

Default

no flow-label

Parameters

flow-label

Specifies the flow label to be used as a match criterion. Value can be expressed as a decimal integer, as well as in hexadecimal or binary format. The following value shows decimal integer format only.

Values 0 to 1048575

mask

Specifies the flow label mask value for this policy IPv6 Filter entry. Value can be expressed as a decimal integer, as well as in hexadecimal or binary format. The following value shows decimal integer format only.

Values 0 to 1048575

Platforms

7705 SAR Gen 2

flow-label

Syntax

flow-label *value*

no flow-label

Context

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter>entry flow-label)

Full Context

configure system security management-access-filter ipv6-filter entry flow-label

Description

This command configures flow label match conditions. Flow labeling enables the labeling of packets belonging to particular traffic flows for which the sender requests special handling, such as non-default quality of service or real-time service.

Parameters

value

Specifies the flow identifier in an IPv6 packet header that can be used to discriminate traffic flows (See RFC 3595, *Textual Conventions for IPv6 Flow Label*.)

Values 0 to 1048575

Platforms

7705 SAR Gen 2

10.35 flr-threshold

flr-threshold

Syntax

flr-threshold *percentage*

no flr-threshold

Context

[Tree] (config>oam-pm>session>ip>twamp-light>loss flr-threshold)

Full Context

configure oam-pm session ip twamp-light loss flr-threshold

Description

This command defines the frame loss threshold used to determine whether the delta-t is available or unavailable. An individual delta-t with a frame loss threshold equal to or higher than the configured threshold is marked unavailable. An individual delta-t with a frame loss threshold lower than the configured threshold is marked as available.

The **no** form of this command restores the default value of 50%.

Default

flr-threshold 50

Parameters

percentage

Specifies the percentage of the threshold.

Values 0 to 100

Default 50

Platforms

7705 SAR Gen 2

10.36 force-l2pt-boundary

force-l2pt-boundary

Syntax

force-l2pt-boundary [**cdp**] [**dtp**] [**pagp**] [**stp**] [**udld**] [**vtp**]

no force-l2pt-boundary

Context

[Tree] (config>service>vpls>sap force-l2pt-boundary)

Full Context

configure service vpls sap force-l2pt-boundary

Description

Enabling force-l2pt-boundary will force all SAPs managed by the specified m-vpls instance on the corresponding port to have l2pt-termination enabled. This command is applicable only to SAPs created under m-vpls regardless of the flavor of STP currently active. It is not applicable to spoke-SDPs.

The execution of this command will fail as soon as at least one of the currently managed SAPs (all SAPs falling within the specified managed-vlan-range) does not have l2pt-termination enabled regardless of its admin/operational status.

If force-l2pt-boundary is enabled on a specified m-vpls SAP, all newly created SAPs falling into the specified managed-vlan-range will have l2pt-termination enabled per default.

Extending or adding new range into a managed-vlan-range declaration will fail as soon as there is at least one SAPs falling into the specified vlan-range does not have l2pt-termination enabled.

Disabling l2pt-termination on currently managed SAPs will fail as soon as the force-l2pt-boundary is enabled under corresponding m-vpls SAP.

Parameters

cdp

Specifies the Cisco discovery protocol

dtp

Specifies the dynamic trunking protocol

pagp

Specifies the port aggregation protocol

stp

Specifies all spanning tree protocols: stp, rstp, mstp, pvst (default)

udld

Specifies unidirectional link detection

vtp

Specifies the virtual trunk protocol

Platforms

7705 SAR Gen 2

10.37 force-renews

force-renews

Syntax

[**no**] **force-renews**

Context

[Tree] (config>router>dhcp>server force-renews)

Full Context

configure router dhcp local-dhcp-server force-renews

Description

This command enables the sending of sending FORCERENEW messages for DHCP.

The **no** form of this command disables the sending of FORCERENEW messages.

Platforms

7705 SAR Gen 2

10.38 force-switchover

force-switchover

Syntax

force-switchover [**now**]

Context

[Tree] (admin>redundancy force-switchover)

Full Context

admin redundancy force-switchover

Description

This command forces a switchover to the standby CPM card. The primary CPM reloads its software image and becomes the secondary CPM.

Parameters

now

Forces the switchover to the redundant CPM card immediately.

Platforms

7705 SAR Gen 2

10.39 force-vlan-vc-forwarding

force-vlan-vc-forwarding

Syntax

[**no**] **force-vlan-vc-forwarding**

Context

[Tree] (config>service>vpls>bgp-evpn>mpls force-vlan-vc-forwarding)

[Tree] (config>service>epipe>bgp-evpn>mpls force-vlan-vc-forwarding)

Full Context

configure service vpls bgp-evpn mpls force-vlan-vc-forwarding

configure service epipe bgp-evpn mpls force-vlan-vc-forwarding

Description

This command enables the system to preserve the VLAN ID and 802.1p bits of the service-delimiting qtag in a new tag, which is sent in the customer frame to the EVPN destinations.

If this configuration is used in conjunction with the **sap ingress vlan-translation** command, the configured translated VLAN ID is the VLAN ID sent to the EVPN destinations, instead of the service-delimiting tag VLAN ID. If the ingress SAP or SDP binding is null-encapsulated, the output VLAN ID and p-bits are zero.

The **no** form of this command does not preserve the VLAN ID and 802.1p bits of the service-delimiting qtag.

Default

no force-vlan-vc-forwarding

Platforms

7705 SAR Gen 2

force-vlan-vc-forwarding

Syntax

[**no**] **force-vlan-vc-forwarding**

Context

[Tree] (config>service>epipe>spoke-sdp force-vlan-vc-forwarding)

[Tree] (config>service>vpls>spoke-sdp force-vlan-vc-forwarding)

[Tree] (config>service>vpls>mesh-sdp force-vlan-vc-forwarding)

Full Context

configure service epipe spoke-sdp force-vlan-vc-forwarding

configure service vpls spoke-sdp force-vlan-vc-forwarding

configure service vpls mesh-sdp force-vlan-vc-forwarding

Description

This command forces vc-vlan-type forwarding in the datapath for spoke and mesh SDPs which have either vc-type. This command is not allowed on vlan-vc-type SDPs.

The **no** version of this command sets default behavior.

Default

no force-vlan-vc-forwarding

Platforms

7705 SAR Gen 2

10.40 format

format

Syntax

format [*cflash-id*] [**reliable**]

Context

[Tree] (file format)

Full Context

file format

Description

This command formats the compact flash. The compact flash must be shut down before starting the format.

Parameters

cflash-id

Specifies the compact flash type.

Values cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

reliable

Enables the reliance file system and disables the default DoS file system. This option is valid only on compact flashes 1 and 2.

Platforms

7705 SAR Gen 2

10.41 forward

forward

Syntax

forward

forward bonding-connection *connection-id*

IPv4: forward esi *esi* **sf-ip** *ip-address* **vas-interface** *interface-name* **router** *router-instance*

IPv6: forward esi *esi* **sf-ip** *ipv6-address* **vas-interface** *interface-name* **router** *router-instance*

IPv4: forward esi *esi* **sf-ip** *ip-address* **vas-interface** *interface-name* **router service-name** *service-name*

IPv6: forward esi *esi* **sf-ip** *ipv6-address* **vas-interface** *interface-name* **router service-name** *service-name*

forward esi *esi* **service-id** *vpls-service-id*

forward gre-tunnel *gre-tunnel-name*

forward lsp *lsp-name*

IPv4: forward mpls-policy *ip-address*

IPv6: forward mpls-policy *ipv6-address*

IPv4: forward next-hop *ip-address*

IPv6: forward next-hop *ipv6-address*

IPv4: forward next-hop *ip-address* **router** *router-instance*

IPv6: forward next-hop *ipv6-address* **router** *router-instance*

IPv4: forward next-hop *ip-address* **router service-name** *service-name*

IPv6: forward next-hop *ipv6-address* **router service-name** *service-name*

IPv4: forward next-hop indirect *ip-address*

IPv6: forward next-hop indirect *ipv6-address*

IPv4: forward next-hop indirect *ip-address* **router** *router-instance*

IPv6: forward next-hop indirect *ipv6-address* **router** *router-instance*

IPv4: forward next-hop indirect *ip-address* **router service-name** *service-name*

IPv6: forward next-hop indirect *ipv6-address* **router service-name** *service-name*

forward redirect-policy *policy-name*

forward router *router-instance*

forward router service-name *service-name*

forward sap *sap-id*

forward sdp *sdp-id:vc-id*

IPv4: forward srte-policy *ip-address* **color** *color-id*

IPv6: forward srte-policy *ipv6-address* **color** *color-id*

IPv4: forward srv6-policy *ipv6-address* **color** *color-id* **service-sid** *ipv6-address*

IPv6: forward srv6-policy *ipv6-address* **color** *color-id* **service-sid** *ipv6-address*

IPv4: forward vprn-target bgp-nh *ip-address* **router** *router-instance* [**adv-prefix** *ip-address/mask*] [**lsp** *lsp-name*]

IPv6: forward vprn-target bgp-nh *ip-address* **router** *router-instance* [**adv-prefix** *ipv6-address/prefix-length*] [**lsp** *lsp-name*]

IPv4: forward vprn-target bgp-nh *ip-address* **router service-name** *service-name* [**adv-prefix** *ip-address/mask*] [**lsp** *lsp-name*]

IPv6: forward vprn-target bgp-nh *ip-address* **router service-name** *service-name* [**adv-prefix** *ipv6-address/prefix-length*] [**lsp** *lsp-name*]

Context

[Tree] (config>filter>ip-filter>entry>action forward)

[Tree] (config>filter>ipv6-filter>entry>action forward)

Full Context

configure filter ip-filter entry action forward

configure filter ipv6-filter entry action forward

Description

This command sets the context for specific forward commands to be performed.

Parameters

connection-id

Specifies that the packet should be forwarded over the specified connection (specified by the connection ID under the bonding group interface), if that connect is available. Outside of a bonding egress context, the behavior of this filter is undefined.

Values 1, 2

esi service-id

Specifies that the packet matching the entry is forwarded to an ESI-identified first appliance in Nuage service chain using EVPN-resolved VXLAN tunnel in the specified VPLS service.

esi sf-ip vas-interface router

Specifies that the packet matching the entry is forwarded to ESI/SF-IP identified first appliance in Nuage service chain using EVPN-resolved VXLAN tunnel over the configured VAS interface in the specified VPRN service.

gre-tunnel-name

Specifies the GRE tunnel name up to 32 characters.

lsp

Specifies that the packet matching the entry is forwarded using the specified lsp.

mpls-policy

Specifies the redirection of the traffic to the programed instance of the MPLS FP specified by its endpoint IPv4 or IPv6 address. The behavior results in a simple forward if no policy exists, if no instance is programmed, and if the policy or instance is administratively down.

next-hop

Specifies that the packet matching the entry is forwarded in the routing context of the incoming interface using direct or indirect IPv4 address in the routing lookup.

next-hop router

Specifies that the packet matching the entry is forwarded in the configured routing context using direct or indirect IPv4 address in the routing lookup.

redirect-policy

Specifies that the packet matching the entry is forwarded using forward next-hop or forward next hop router and the IP address of destination selected by the configured redirect policy. If no destination is selected, packets are subject to action forward.

router

Specifies that the packet matching the entry is routed in the configured routing instance and not in the incoming interface routing instance.

sap

Specifies that the packet matching the entry is forwarded using the configured SAP.

sdp

Specifies that the packet matching the entry is forwarded using the configured SDP.

srte-policy

Specifies the redirection of the traffic to the programed instance of the SR-TE policy specified by its endpoint IPv4 address or IPv6 address and color. The behavior results

in a simple forward if no policy exists, if no instance is programmed, and if the policy or instance is administratively down.

color-id

Specifies the color identifier of the specified SR-TE policy.

Values 0 to 4294967295

vprn-target

Specifies that the packet matching the entry is redirected towards a designated BGP next-hop (**bgp-nh**). The user may specify an LSP (**lsp** *lsp-name*) to use towards that next-hop. If no LSP is specified, the system will automatically select one. The user must specify the routing context (**router** {*router-instance* | **service-name** *service-name*}) in which the system will perform the lookups in order to derive the proper VPRN service label. The user may specify an advertised prefix route (**adv-prefix** *ip-address/prefix-length*). This is needed in case label per VRF is not the label allocation method configured at the BGP peer.

esi

Specifies a 10-byte Ethernet Segment Identifier.

ip-address/mask

Specifies an IPv4 advertised route in the CIDR notation. The IPv4 address is in dotted decimal notation.

Values ip-address a.b.c.d (host bits must be 0)

mask: 0 to 32

ipv6-address/prefix-length

Specifies an IPv6 advertised route in the CIDR notation.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d, where "x" is [0..FFFF]H, and "d" is [0..255]

prefix-length: 0 to 128

bgp-nh ip-address

Specifies the IPv4 address (in dotted decimal notation) of the target BGP next-hop.

Values ip-address d.d.d.d

ipv6-address

Specifies the IPv6 address of a direct or indirect next hop to forward matching packets or of the service SID to use with the SRv6 policy.

interface-name

Specifies the (maximum 32-character) name of an egress R-VPLS IP interface used to forward the packets using ESI redirect for VPRN/IES service.

lsp-name

Specifies an existing RSVP-TE, MPLS-TP, or SR-TE LSP that supports LSP redirect.

policy-name

Specifies an IPv4 redirect policy configured in the config>filter>redirect-policy context.

sap-id

Specifies an existing VPLS Ethernet SAP.

sdp-id:vc-id

Specifies an existing VPLS SDP.

router-instance

Specifies "Base" or an existing VPRN service ID. For the **forward vprn-target bgp-nh** command, *router-instance* must specify an existing VPRN service ID.

service-name

Specifies an existing VPRN service name.

vpls-service-id

Specifies an existing VPLS service ID or service name.

Platforms

7705 SAR Gen 2

10.42 forward-6in4

forward-6in4

Syntax

[**no**] **forward-6in4**

Context

[Tree] (config>system>ip forward-6in4)

Full Context

configure system ip forward-6in4

Description

This command enables forwarding of IPv6 traffic encapsulated in an IPv4 transport sent to the system IP address.

The **no** form of this command disables this option and returns the system to the default behavior.

Default

no forward-6in4

Platforms

7705 SAR Gen 2

10.43 forward-delay

forward-delay

Syntax

forward-delay *forward-delay*

no forward-delay [*forward-delay*]

Context

[Tree] (config>service>template>vpls-template>stp forward-delay)

[Tree] (config>service>vpls>stp forward-delay)

Full Context

configure service template vpls-template stp forward-delay

configure service vpls stp forward-delay

Description

RSTP, as defined in the IEEE 802.1D-2004 standards, will normally transition to the forwarding state via a handshaking mechanism (rapid transition), without any waiting times. If handshaking fails (foe example, on shared links, see below), the system falls back to the timer-based mechanism defined in the original STP (802.1D-1998) standard.

A shared link is a link with more than two nodes (for example, a shared 10/100BaseT segment). The port-type command is used to configure a link as point-to-point or shared.

For timer-based transitions, the 802.1D-2004 standard defines an internal variable forward-delay, which is used in calculating the default number of seconds that a SAP or spoke-SDP spends in the discarding and learning states when transitioning to the forwarding state.

The value of the forward-delay variable depends on the STP operating mode of the VPLS instance:

- in rstp or mstp mode, but only when the SAP or spoke-SDP has not fallen back to legacy STP operation, the value configured by the hello-time command is used;
- in all other situations, the value configured by the forward-delay command is used.

Default

forward-delay 15

Parameters

seconds

The forward delay timer for the STP instance in seconds

Values 4 to 30

Platforms
7705 SAR Gen 2

10.44 forward-ip-over-gre

forward-ip-over-gre

Syntax
[**no**] **forward-ip-over-gre**

Context
[Tree] (config>system>ip forward-ip-over-gre)

Full Context
configure system ip forward-ip-over-gre

Description
This command enables forwarding of IP traffic encapsulated in a GRE over IPv4 transport sent to the system IP address.

The **no** form of this command disables this option and returns the system to the default behavior.

Default
no forward-ip-over-gre

Platforms
7705 SAR Gen 2

10.45 forward-ipv4-packets

forward-ipv4-packets

Syntax
[**no**] **forward-ipv4-packets**

Context
[Tree] (config>service>vprn>if>ipv6 forward-ipv4-packets)

Full Context
configure service vprn interface ipv6 forward-ipv4-packets

Description

This command allows an IPv6-only interface (with no configured IPv4 addresses) to be used for forwarding transit and locally originating and terminating IPv4 packets.

The interface will report that its IPv4 oper-state is up if its IPv6 oper-state is up. Be aware that not all protocols will observe the interface as up from an IPv4 perspective. This command is mostly intended to support BGP routing use cases. Refer to RFC 5549, Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop, for further information.

The **no** form of this command restores the default behavior and prevents the interface from forwarding IPv4 packets if it has no configured IPv4 subnets.

Platforms

7705 SAR Gen 2

forward-ipv4-packets

Syntax

[**no**] **forward-ipv4-packets**

Context

[Tree] (config>router>if>ipv6 forward-ipv4-packets)

Full Context

configure router interface ipv6 forward-ipv4-packets

Description

This command allows an IPv6-only interface (with no configured IPv4 addresses) to be used for forwarding transit and locally originating and terminating IPv4 packets.

The interface reports that its IPv4 operational state is up if its IPv6 operational state is up. Be aware that not all protocols observe the interface as up from an IPv4 perspective. This command is mostly intended to support BGP routing use cases. Refer to RFC 5549, Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop, for further information.

The **no** form of this command restores the default behavior and prevents the interface from forwarding IPv4 packets if it has no configured IPv4 subnets.

Default

no forward-ipv4-packets

Platforms

7705 SAR Gen 2

10.46 forwarding

forwarding

Syntax

forwarding *limit*

no forwarding

Context

[Tree] (config>service>nat>nat-policy>port-limits forwarding)

Full Context

configure service nat nat-policy port-limits forwarding

Description

This command configures the maximum number of port forwarding entries.

Default

no forwarding

Parameters

limit

Specifies the maximum number of port forwarding entries per subscriber.

Values 1 to 64

Platforms

7705 SAR Gen 2

forwarding

Syntax

forwarding {**next-hop** *ip-address* | **interface** *interface-name* | **bypass-routing**}

no forwarding

Context

[Tree] (config>oam-pm>session>ip forwarding)

Full Context

configure oam-pm session ip forwarding

Description

This command influences the forwarding decision of the TWAMP Light packet. When this command is used, only one of the forwarding options can be enabled at any time.

The **no** form of this command removes the options and enables the default forwarding logic.

Parameters

ip-address

Specifies the IP address of the next hop on the path.

Values

ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

interface-name

Specifies the name, up to 32 characters, to refer to the interface from which the packet is sent. The name must already exist in the **config>router>interface** context or within the appropriate **config>service** context.

bypass-routing

Specifies to send the packet to a host on a directly attached network, bypassing the routing table.

Platforms

7705 SAR Gen 2

10.47 forwarding-bits-set

forwarding-bits-set

Syntax

forwarding-bits-set {**all** | **non-fwd**}

no forwarding-bits-set

Context

[Tree] (config>service>vprn>bgp>group>graceful-restart>long-lived forwarding-bits-set)

[Tree] (config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived forwarding-bits-set)

[Tree] (config>service>vprn>bgp>graceful-restart>long-lived forwarding-bits-set)

Full Context

configure service vprn bgp group graceful-restart long-lived forwarding-bits-set

configure service vprn bgp group neighbor graceful-restart long-lived forwarding-bits-set

configure service vprn bgp graceful-restart long-lived forwarding-bits-set

Description

This command determines the setting of the F bits in the GR and LLGR capabilities advertised by the router. When the F bit is set for an AFI/SAFI, it indicates that the advertising router was able to preserve forwarding state for the routes of that AFI/SAFI across the last restart. If a router restarts and does not set F=1, then when the session with a peer is re-established, the peer immediately deletes all LLGR stale routes it was preserving on behalf of the restarting router for the corresponding AFI/SAFI.

This command allows the F bits for all advertised AFI/SAFI to be set to 1, or only the F bits for non-forwarding AFI/SAFI to be set to 1. Non-forwarding AFI/SAFI are the following configuration-related address families: L2-VPN, route-target, flow-IPv4, and flow-IPv6.

Default

no forwarding-bits-set

Parameters

all

Specifies that the F bit for all AFI/SAFI should be set to 1.

non-fwd

Specifies that the F bit for only non-forwarding AFI/SAFI should be set to 1. These AFI/SAFI correspond to the following families: L2-VPN, route-target, flow-IPv4, and flow-IPv6.

Platforms

7705 SAR Gen 2

forwarding-bits-set

Syntax

forwarding-bits-set {**all** | **non-fwd**}

no forwarding-bits-set

Context

[Tree] (config>router>bgp>graceful-restart>long-lived forwarding-bits-set)

[Tree] (config>router>bgp>group>graceful-restart>long-lived forwarding-bits-set)

[Tree] (config>router>bgp>group>neighbor>graceful-restart>long-lived forwarding-bits-set)

Full Context

configure router bgp graceful-restart long-lived forwarding-bits-set

configure router bgp group graceful-restart long-lived forwarding-bits-set

configure router bgp group neighbor graceful-restart long-lived forwarding-bits-set

Description

This command determines the setting of the F bits in the GR and LLGR capabilities advertised by the router. When the F bit is set for an AFI/SAFI, it indicates that the advertising router was able to preserve forwarding state for the routes of that AFI/SAFI across the last restart. If a router restarts and does not set F=1, then when the session with a peer re-establishes the peer immediately deletes all LLGR stale routes it was preserving on behalf of the restarting router for the corresponding AFI/SAFI.

This command allows the F bits for all advertised AFI/SAFI to be set to 1, or only the F bits for non-forwarding AFI/SAFI to be set to 1. Non-forwarding AFI/SAFI are the following configuration-related address families: L2-VPN, route-target, flow-IPv4, and flow-IPv6.

Default

no forwarding-bits-set

Parameters

all

Specifies that the F bit for all AFI/SAFI should be set to 1.

non-fwd

Specifies that the F bit for only non-forwarding AFI/SAFI should be set to 1. These AFI/SAFI correspond to the following families: L2-VPN, route-target, flow-IPv4, and flow-IPv6.

Platforms

7705 SAR Gen 2

10.48 forwarding-policies

forwarding-policies

Syntax

[**no**] **forwarding-policies**

Context

[Tree] (config>router>mpls forwarding-policies)

Full Context

configure router mpls forwarding-policies

Description

Commands in this context configure an MPLS forwarding policy.

The **no** form of this command deletes all policies from the forwarding policy database.

Platforms

7705 SAR Gen 2

10.49 forwarding-policy

forwarding-policy

Syntax

[**no**] **forwarding-policy** *name*

Context

[Tree] (config>router>mpls>fwd-policies forwarding-policy)

Full Context

configure router mpls forwarding-policies forwarding-policy

Description

This command creates an MPLS forwarding policy.

There are two types of MPLS forwarding policy:

- endpoint policy
- label-binding policy

The endpoint policy allows the user to forward unlabeled packets over a set of user-defined direct (with option to push a label stack) or indirect next hops. Routes are bound to an endpoint policy when their next hop matches the endpoint address of the policy.

The label-binding policy provides the same capability for labeled packets. In this case, labeled packets matching the ILM of the policy binding label are forwarded over the set of next hops of the policy.

The data model of a forwarding policy represents each pair of {primary next hop, backup next hop} as a group and models the ECMP set as the set of Next-Hop Groups (NHGs). Flows of prefixes can be switched on a per-NHG basis from the primary next hop, when it fails, to the backup next hop without disturbing the flows forwarded over the other NHGs of the policy. The same can be performed when reverting back from a backup next hop to the restored primary next hop of the same NHG.

The MPLS forwarding policy supports two types of NHGs on a per policy basis:

- An NHG of resolution type indirect supported with the label-binding policy and in which forwarding over the primary/backup next hop is modeled as a swap operation from the binding label to an implicit-null label over multiple outgoing interfaces (multiple NHLFEs) corresponding to the resolved next hops of the indirect route.

  Within a given NHG, the primary next hop is the preferred active path in the absence of any failure of the NHG of resolution type indirect.

  The forwarding database tracks the primary or backup next hop in the routing table. A **route delete** of the primary indirect next hop causes the CPM to program the backup indirect next hop in the data path.

A **route modify** to the indirect primary or backup next hop causes the CPM to update the resolved next hops and to update the data path if it is the active indirect next hop.

When the primary indirect next hop is restored and is added back into the routing table, CPM waits for an amount of time equal to the user-programmed revert timer before updating the data path. However, if the backup indirect next hop fails while the timer is running, CPM updates the data path immediately.

- An NHG of resolution type direct is modeled as follows:
  - For a label-binding policy, forwarding over the primary or backup next hop is modeled as a swap operation from the binding label to the configured label stack or to an implicit-null label (if the **pushed-labels** command not configured) over a single outgoing interface to the next hop.
  - For an endpoint policy, forwarding over the primary or backup next hop is modeled as a push operation from the binding label to the configured label stack or to an implicit-null label (if the **pushed-labels** command not configured) over a single outgoing interface to the next hop.
  - The labels configured by the **pushed-labels** command are not validated.

  Within a given NHG, the primary next hop is the preferred active path in the absence of any failure of the NHG of resolution type direct.

  The NHG supports uniform failover. The forwarding policy database assigns a Protect-Group ID (PG-ID) to each of the primary next hop and the backup next hop and programs both of them in data path. A failure of the active path switches traffic to the other path following the uniform failover procedures.

  The forwarding database tracks the primary or backup next hop in the routing table. A **route delete** of the primary/backup direct next hop causes CPM to send the corresponding PG-ID switch to the data path.

  A **route modify** to the direct primary or backup next hop causes CPM to update the MPLS forwarding database and to update the data path since both next hops are programmed.

  When the primary direct next hop is restored and is added back into the routing table, CPM waits for an amount of time equal to the user programmed revert timer before activating it and updating the data path. However, if the backup direct next hop fails while the timer is running, CPM activates it and updates the data path immediately. The latter failover to the restored primary next hop is performed using the uniform failover procedure.

The forwarding policy database activates the best endpoint policy among the named policies sharing the same value of the endpoint parameter by selecting the lowest preference value policy. This policy is then programmed into the TTM and into the tunnel table in data path. If this policy goes down, then the forwarding policy database performs a re-evaluation and activates the named policy with the next lowest preference value for the same endpoint value. If a more preferred policy comes back up, the forwarding policy database reverts to it and activates it.

The forwarding policy database similarly activates the best label-binding policy among the named policies sharing the same binding label by selecting the lowest preference value policy. This policy is then programmed into the label FIB table in data path. If this policy goes down, then the forwarding policy database performs a re-evaluation and activates the names policy with the next lowest preference value for the same binding label value. If a more preferred policy comes back up, the forwarding policy database reverts to it and activates it.

Ingress statistics can be enabled as is associated with binding label, that is the ILM of the forwarding policy, and provides aggregate packet and byte counters for packets matching the binding label.

The **no** form of the command deletes the named MPLS forwarding policy.

Parameters

name

Specifies the name of the MPLS forwarding policy, up to 64 characters.

Platforms

7705 SAR Gen 2

10.50 fp

fp

Syntax

fp [*fp-number*]

Context

[Tree] (config>card fp)

Full Context

configure card fp

Description

This command enables access to the configuration of the forwarding planes on a card.

The default forwarding plane is 1. When entering the FP node, if the forwarding plane number is omitted, the system will assume forwarding plane number 1.

Commands can only be configured under **card>fp** if the hardware that the FP resides on (either a card or an XMA) is provisioned. Conversely, all commands under **card>fp** of the corresponding FPs are automatically removed when that hardware is unprovisioned.

Parameters

fp-number

Specifies that the FP number parameter is optional following the **fp** command.

Values 1 to 8

Default fp 1

Platforms

7705 SAR Gen 2

10.51 fp-redirect-group

fp-redirect-group

Syntax

fp-redirect-group *policer-type policer-id*

no fp-redirect-group *policer-type*

Context

[Tree] (config>qos>network>ingress>fc fp-redirect-group)

Full Context

configure qos network ingress fc fp-redirect-group

Description

This command is used to redirect the FC of a broadcast packet received in a VPLS service over a PW or network IP interface to an ingress forwarding plane queue-group.

It defines the mapping of an FC to a policer-id and redirects the lookup of the policer of the same ID in some ingress forwarding plane queue-group instance. However, the queue-group name and instance are explicitly provided only at the time the network QoS policy is applied to the ingress context of a spoke or mesh SDP or a network IP interface.

The broadcast-policer statement is ignored when the network QoS policy is applied to any object other than a VPLS spoke or mesh SDP or a network IP interface.

The **no** form of this command removes the redirection of the FC.

Parameters

policer-type

The policer type to be used. The *policer-type* is ignored when the network QoS policy is applied to any object other than a VPLS spoke or mesh SDP or a network IP interface.

Values broadcast-policer | mcast-policer | policer | unknown-policer

policer-id

The specified *policer-id* must exist within the queue-group template applied to the ingress context of the forwarding plane.

Values 1 to 32

Platforms

7705 SAR Gen 2

10.52 frag-required

frag-required

Syntax

[**no**] **frag-required**

Context

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>icmp-generation frag-required)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel>icmp-generation frag-required)

[Tree] (config>service>ies>if>sap>ip-tunnel>icmp-generation frag-required)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>icmp-generation frag-required)

[Tree] (config>router>if>ipsec>ipsec-tunnel>icmp-generation frag-required)

[Tree] (config>ipsec>tnl-temp>icmp-generation frag-required)

[Tree] (config>service>vprn>if>sap>ip-tunnel>icmp-generation frag-required)

Full Context

configure service vprn interface ipsec ipsec-tunnel icmp-generation frag-required

configure service vprn interface sap ipsec-tunnel icmp-generation frag-required

configure service ies interface sap ip-tunnel icmp-generation frag-required

configure service ies interface ipsec ipsec-tunnel icmp-generation frag-required

configure router interface ipsec ipsec-tunnel icmp-generation frag-required

configure ipsec tunnel-template icmp-generation frag-required

configure service vprn interface sap ip-tunnel icmp-generation frag-required

Description

Commands in this context configure ICMP Fragmentation Required parameters.

The **no** form of this command disables sending the ICMP messages.

Platforms

7705 SAR Gen 2

10.53 fragment

fragment

Syntax

fragment {**true** | **false**}

no fragment

Context

[Tree] (config>qos>sap-ingress>ip-criteria>entry>match fragment)

[Tree] (config>qos>sap-egress>ip-criteria>entry>match fragment)

Full Context

configure qos sap-ingress ip-criteria entry match fragment

configure qos sap-egress ip-criteria entry match fragment

Description

This command configures fragmented or non-fragmented IP packets as a SAP QoS policy match criterion.

The **no** form of this command removes the match criterion and matches all packets regardless of whether they are fragmented or not.

Default

no fragment

Parameters

true

Configures a match on all fragmented IP packets. A match will occur for all packets that have either the MF (more fragment) bit set or have the Fragment Offset field of the IP header set to a non-zero value.

false

Configures a match on all non-fragmented IP packets. Non-fragmented IP packets are packets that have the MF bit set to zero and have the Fragment Offset field also set to zero.

Platforms

7705 SAR Gen 2

fragment

Syntax

fragment {**true** | **false** | **first-only** | **non-first-only**}

no fragment

Context

[Tree] (config>qos>sap-ingress>ipv6-criteria>entry>match fragment)

Full Context

configure qos sap-ingress ipv6-criteria entry match fragment

Description

This command configures fragmented or non-fragmented IPv6 packets as a SAP ingress QoS policy match criterion.

The **no** form of this command removes the match criterion and matches all packets regardless of whether they are fragmented or not.

Default

no fragment

Parameters

true

Specifies to match on all fragmented IPv6 packets. A match will occur for all packets that contain an IPv6 Fragmentation Extension Header.

false

Specifies to match on all non-fragmented IP packets. Non-fragmented IPv6 packets are packets that do not contain an IPv6 Fragmentation Extension Header.

first-only

Matches if a packet is an initial fragment of the fragmented IPv6 packet.

non-first-only

Matches if a packet is a non-initial fragment of the fragmented IPv6 packet.

Platforms

7705 SAR Gen 2

fragment

Syntax

fragment {**true** | **false**}

no fragment

Context

[Tree] (config>qos>network>ingress>ip-criteria>entry>match fragment)

[Tree] (config>qos>network>egress>ip-criteria>entry>match fragment)

Full Context

configure qos network ingress ip-criteria entry match fragment

configure qos network egress ip-criteria entry match fragment

Description

This command configures fragmented or non-fragmented IP packets as a network QoS policy match criterion.

The **no** form of this command removes the match criterion and matches all packets regardless of whether they are fragmented or not.

Parameters

true

Configures a match on all fragmented IP packets. A match will occur for all packets that have either the MF (more fragment) bit set or have the Fragment Offset field of the IP header set to a non-zero value.

false

Configures a match on all non-fragmented IP packets. Non-fragmented IP packets are packets that have the MF bit set to zero and have the Fragment Offset field also set to zero.

Platforms

7705 SAR Gen 2

fragment

Syntax

fragment {**true** | **false** | **first-only** | **non-first-only**}

no fragment

Context

[Tree] (config>qos>network>egress>ipv6-criteria>entry>match fragment)

[Tree] (config>qos>network>ingress>ipv6-criteria>entry>match fragment)

Full Context

configure qos network egress ipv6-criteria entry match fragment

configure qos network ingress ipv6-criteria entry match fragment

Description

This command configures fragmented or non-fragmented IPv6 packets as a network QoS policy match criterion.

The **no** form of this command removes the match criterion and matches all packets regardless of whether they are fragmented or not.

Parameters

true

Specifies to match on all fragmented IPv6 packets. A match will occur for all packets that contain an IPv6 Fragmentation Extension Header.

false

Specifies to match on all non-fragmented IP packets. Non-fragmented IPv6 packets are packets that do not contain an IPv6 Fragmentation Extension Header.

first-only

Matches if a packet is an initial fragment of the fragmented IPv6 packet.

non-first-only

Matches if a packet is a non-initial fragment of the fragmented IPv6 packet.

Platforms

7705 SAR Gen 2

fragment

Syntax

fragment {**true** | **false** | **first-only** | **non-first-only**}

no fragment

Context

[Tree] (config>filter>ipv6-filter>entry>match fragment)

[Tree] (config>filter>ip-filter>entry>match fragment)

Full Context

configure filter ipv6-filter entry match fragment

configure filter ip-filter entry match fragment

Description

This command specifies match criterion for fragmented packets.

Matches can be based on the presence of a fragmented packet (or otherwise) on the ingress or egress interface.

Matches can also be based on the presence of the first fragment of a packet, or on the presence of a fragment that is not the first fragment on the ingress interface.

The **no** form of the command removes the match criterion.

Default

no fragment

Parameters

true

Specifies to match on all fragmented packets.

false

Specifies to match on all non-fragmented packets.

first-only

Matches if a packet is an initial fragment of a fragmented packet.

non-first-only

Matches if a packet is a non-initial fragment of a fragmented packet.

Platforms

7705 SAR Gen 2

10.54 frame-based-accounting

frame-based-accounting

Syntax

[**no**] **frame-based-accounting**

Context

[Tree] (config>qos>scheduler-policy frame-based-accounting)

Full Context

configure qos scheduler-policy frame-based-accounting

Description

The frame-based-accounting command is used to enable frame-based accounting for both the children queues parented to the scheduling policy and for the schedulers within the scheduler policy.

When frame-based accounting is enabled on the policy, all queues associated with the scheduler (through the parent command on each queue) will have their rate and CIR values interpreted as frame-based values. When shaping, the queues will include the 12-byte Inter-Frame Gap (IFG) and 8 byte preamble for each packet scheduled out the queue. The profiling CIR threshold will also include the 20-byte frame encapsulation overhead. Statistics associated with the queue do not include the frame encapsulation overhead.

The scheduler policy’s scheduler rate and CIR values will be interpreted as frame-based values.

The configuration of **parent-location** and **frame-based-accounting** in a scheduler policy is mutually exclusive to ensure consistency between the different scheduling levels. Packet byte offset settings are not included in the applied rate when frame-based accounting is configured; however, the offsets are applied to the statistics.

The **no** form of this command is used to return all schedulers within the policy and queues associated with the policy to the default packet-based accounting mode. If **frame-based-accounting** is not currently enabled for the scheduling policy, the **no frame-based-accounting** command has no effect.

Platforms

7705 SAR Gen 2

10.55 framed-ip-addr

framed-ip-addr

Syntax

[**no**] **framed-ip-addr**

Context

[Tree] (config>ipsec>rad-acct-plcy>include framed-ip-addr)

Full Context

configure ipsec radius-accounting-policy include-radius-attribute framed-ip-addr

Description

This command enables the inclusion of the **framed-ip-addr** attribute.

Default

no framed-ip-addr

Platforms

7705 SAR Gen 2

10.56 framed-ipv6-prefix

framed-ipv6-prefix

Syntax

[**no**] **framed-ipv6-prefix**

Context

[Tree] (config>ipsec>rad-acct-plcy>include framed-ipv6-prefix)

Full Context

configure ipsec radius-accounting-policy include-radius-attribute framed-ipv6-prefix

Description

This command enables the inclusion of the **framed-ipv6-prefix** attribute.

Default

no framed-ipv6-prefix

Platforms

7705 SAR Gen 2

10.57 frequency

frequency

Syntax

frequency *frequency*

no frequency

Context

[Tree] (config>port>dwdm frequency)

Full Context

configure port dwdm frequency

Description

This command configures the center frequency to use for a tunable DWDM optical interface. It replaces the **configure**>**port**>**dwdm**>**channel** command (used prior to Release 22.2.R1). The **frequency** command supports any frequency in the C band, but the actual operating frequency is dependent on the installed optic module.

Provisioning rules

The provisioned MDA type must have DWDM tunable optics (for example, p1-100g-tun-b) or the MDA must support the option of tunable DWDM optic modules. The following provisioning rules apply:

- The DWDM frequency must set to a non-zero value before the port is set to **no shutdown**.
- The port must be **shutdown** before changing the DWDM frequency.
- The port must be a physical port to set the DWDM frequency.

Default

frequency 0

Parameters

frequency

Specifies the frequency in MHz.

Values 0, 191100000 to 196150000

Platforms

7705 SAR Gen 2

10.58 from

from

Syntax

from [**main**] [**security**] [**change**] [**debug-trace**]

no from

Context

[Tree] (config>service>vprn>log>log-id from)

Full Context

configure service vprn log log-id from

Description

This command selects the source stream to be sent to a log destination.

One or more source streams must be specified. The source of the data stream must be identified using the **from** command before you can configure the destination using the **to** command. The **from** command can identify multiple source streams in a single statement (for example: **from main change debug-trace**).

Only one **from** command may be entered for a single *log-id*. If multiple **from** commands are configured, then the last command entered overwrites the previous **from** command.

The **no** form of this command removes all previously configured source streams.

Default

No source stream is configured.

Parameters

main

Instructs all events in the main event stream to be sent to the destination defined in the **to** command for this destination *log-id*. The main event stream contains the events that are not explicitly directed to any other event stream. To limit the events forwarded to the destination, configure filters using the **filter** command.

security

Instructs all events in the security event stream to be sent to the destination defined in the **to** command for this destination *log-id*. The security event stream contains all events that pertain to attempts to breach system security. To limit the events forwarded to the destination, configure filters using the **filter** command.

change

Instructs all events in the user activity stream to be sent to the destination configured in the **to** command for this destination *log-id*. The change event stream contains all events that directly affect the configuration or operation of this node. To limit the events forwarded to the change stream destination, configure filters using the filter command.

debug-trace

Instructs all events in the debug-trace event stream to be sent to the destination defined in the **to** command for this destination *log-id*. The debug-trace event stream contains all events that pertain to trace or other debugging information. To limit the events forwarded to the destination, configure filters using the **filter** command.

Platforms

7705 SAR Gen 2

from

Syntax

from *ip-address*

Context

[Tree] (config>router>mpls>lsp-template from)

[Tree] (config>router>mpls>lsp from)

Full Context

configure router mpls lsp-template from

configure router mpls lsp from

Description

This optional command specifies the IP address of the ingress router for the LSP. When this command is not specified, the system IP address is used. IP addresses that are not defined in the system are allowed. If an invalid IP address is entered, LSP bring-up fails and an error is logged.

If an interface IP address is specified as the **from** address, and the egress interface of the LSP nexthop IP address is a different interface, the LSP is not signaled. As the egress interface changes due to changes in the routing topology, it is recommended to set the **from** IP address to the system IP address or to the address of a loopback interface to ensure the LSP recovers.

Only one **from** address can be configured.

Default

The system IP address

Parameters

ip-address

Specifies the IP address of the ingress router. This can be either the interface, the system or a loopback interface IP address. If the IP address is local, the LSP must egress through that local interface which ensures local strictness. When the LSP type is **sr-te**, then an IPv6 address can be used.

Values ipv4-address — a.b.c.d

ipv6-address — x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x — 0 to FFFF (hexadecimal)

d — 0 to 255 (decimal)

Platforms

7705 SAR Gen 2

from

Syntax

from {[**main**] [**security**] [**change**] [**debug-trace**]}

no from

Context

[Tree] (config>log>log-id from)

Full Context

configure log log-id from

Description

This command selects the source stream to be sent to a log destination.

One or more source streams must be specified. The source of the data stream must be identified using the **from** command before you can configure the destination using the **to** command. The **from** command can identify multiple source streams in a single statement (for example: **from main change debug-trace**).

Only one **from** command may be entered for a single *log-id*. If multiple **from** commands are configured, then the last command entered overwrites the previous **from** command.

The **no** form of this command removes all previously configured source streams.

Parameters

main

Instructs all events in the main event stream to be sent to the destination defined in the **to** command for this destination *log-id*. The main event stream contains the events that are not explicitly directed to any other event stream. To limit the events forwarded to the destination, configure filters using the **filter** command.

security

Instructs all events in the security event stream to be sent to the destination defined in the **to** command for this destination *log-id*. The security stream contains all events that affect attempts to breach system security such as failed login attempts, attempts to access MIB tables to which the user is not granted access or attempts to enter a branch of the CLI to which access has not been granted. To limit the events forwarded to the destination, configure filters using the **filter** command.

change

Instructs all events in the user activity stream to be sent to the destination configured in the **to** command for this destination *log-id*. The change event stream contains all events that directly affect the configuration or operation of this node. To limit the events forwarded to the change stream destination, configure filters using the **filter** command.

debug-trace

Instructs all debug-trace messages in the debug stream to be sent to the destination configured in the **to** command for this destination *log-id*. Filters applied to debug messages are limited to application and subject.

Platforms

7705 SAR Gen 2

from

Syntax

[**no**] **from**

Context

[Tree] (config>router>policy-options>policy-statement>entry from)

Full Context

configure router policy-options policy-statement entry from

Description

This command creates the context to configure policy match criteria based on a route’s source or the protocol from which the route is received.

If no condition is specified, all route sources are considered to match.

The **no** form of this command deletes the source match criteria for the route policy statement entry.

Platforms

7705 SAR Gen 2

from

Syntax

from *ipv4-address*

no from

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls>rsvp-te-auto from)

Full Context

configure oam-pm session ip tunnel mpls rsvp-te-auto from

Description

This command configures the headend of the RSVP LSP. Configure the following three commands to identify an RSVP-TE Auto LSP: **from**, **to**, and **lsp-template**. When all three of these values are configured, the specific RSVP LSP can be identified and the test packets can be carried across the tunnel

The **no** form of this command removes the IPv4 address.

Parameters

ipv4-address

Specifies an IPv4 address.

Values ipv4-address: a.b.c.d (host bits must be 0)

Platforms

7705 SAR Gen 2

10.59 frr

frr

Syntax

frr [**detail**]

no frr

Context

[Tree] (debug>router>mpls>event frr)

Full Context

debug router mpls event frr

Description

This command debugs fast re-route events.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about re-route events.

Platforms

7705 SAR Gen 2

10.60 frr-object

frr-object

Syntax

[**no**] **frr-object**

Context

[Tree] (config>router>mpls frr-object)

Full Context

configure router mpls frr-object

Description

This command specifies whether fast reroute for LSPs using the **facility** bypass method is signaled with or without the fast reroute object using the **one-to-one** keyword. The value is ignored if fast reroute is disabled for the LSP or if the LSP is using one-to-one Backup.

Default

frr-object — Specifies the value is by default inherited by all LSPs.

Platforms

7705 SAR Gen 2

10.61 fsm-state-changes

fsm-state-changes

Syntax

[**no**] **fsm-state-changes**

Context

[Tree] (debug>service>id>stp fsm-state-changes)

Full Context

debug service id stp fsm-state-changes

Description

This command enables STP debugging for FSM state changes.

The **no** form of the command disables debugging.

Platforms

7705 SAR Gen 2

10.62 fsm-timers

fsm-timers

Syntax

[**no**] **fsm-timers**

Context

[Tree] (debug>service>id>stp fsm-timers)

Full Context

debug service id stp fsm-timers

Description

This command enables STP debugging for FSM timer changes.

The **no** form of the command disables debugging.

Platforms
7705 SAR Gen 2

10.63 ftp

ftp

Syntax
ftp

Context
[Tree] (config>system>login-control ftp)

Full Context
configure system login-control ftp

Description
This command creates the context to configure FTP login control parameters.

Platforms
7705 SAR Gen 2

10.64 ftp-server

ftp-server

Syntax
[**no**] **ftp-server**

Context
[Tree] (config>system>security ftp-server)

Full Context
configure system security ftp-server

Description
This command enables FTP servers running on the system.

FTP servers are disabled by default. At system startup, only SSH servers are enabled.

The **no** form of this command disables FTP servers running on the system.

Platforms

7705 SAR Gen 2

11 g Commands

11.1 garp-flood-evpn

garp-flood-evpn

Syntax

[**no**] **garp-flood-evpn**

Context

[Tree] (config>service>vpls>proxy-arp garp-flood-evpn)

Full Context

configure service vpls proxy-arp garp-flood-evpn

Description

This command controls whether the system floods GARP-requests and GARP-replies to the EVPN. The GARPs impacted by this command are identified by the sender's IP being equal to the target's IP and the MAC DA being broadcast.

The **no** form of the command only floods to local SAPs or binds but not to EVPN destinations.

Disabling this command is only recommended in networks where CEs are routers that are directly connected to the PEs. Networks using aggregation switches between the host/routers and the PEs should flood GARP messages in the EVPN to ensure that the remote caches are updated and the BGP does not miss the advertisement of these entries.

Default

garp-flood-evpn

Platforms

7705 SAR Gen 2

11.2 gateway

gateway

Syntax

gateway name *name* **tunnel** *ip-address*[**:***port*] [**nat-ip** *nat-ip*[**:***port*]] [**detail**] [**no-dpd-debug**] [**display-keys**]

no gateway name *name* **tunnel** *ip-address*[**:***port*] [**nat-ip** *nat-ip*[**:***port*]

gateway name *name* **tunnel-subnet** *ip-prefix/ip-prefix-length* [**port** *port*] [**detail**] [**no-dpd-debug**] [**display-keys**]

no gateway name *name* **tunnel-subnet** *ip-prefix/ip-prefix-length*

Context

[Tree] (debug>ipsec gateway)

Full Context

debug ipsec gateway

Description

This command enables debugging for dynamic IPsec tunnels that terminate on the specified IPsec gateway.

The tunnel to be debugged can be specified by either its source address or source subnet. If a subnet is specified, the system will enable debugging for all tunnels with source addresses in the specified subnet.

Parameters

name

Specifies the name of the IPsec gateway up to 32 characters.

ip-address:port

Specifies the tunnel IP address of the remote peer and, optionally, the remote UDP port of IKE.

nat-ip:port

Specifies the inside IP address of the NAT tunnel and, optionally, the port.

detail

Specifies to display detailed debug information.

no-dpd-debug

Specifies to stop logging IKEv1 and IKEv2 DPD events during debug in order to produce less noise.

ip-prefix/ip-prefix-length

Specifies the subnet of the peer’s tunnel address.

display-keys

Specifies the IKE-SA and CHILD-SA keys for inclusion in the debug output.

Platforms

7705 SAR Gen 2

11.3 gen-keypair

gen-keypair

Syntax

gen-keypair *url-string* **curve** {**secp256r1** | **secp384r1** | **secp521r1**}

gen-keypair *url-string* [**size** *key-size*] [**type** {**rsa** | **dsa**}]

Context

[Tree] (admin>certificate gen-keypair)

Full Context

admin certificate gen-keypair

Description

This command generates RSA, DSA, or ECDSA private key or public key pairs at the specified location.

Parameters

url-string

Specifies the path of the key file.

Values		
	url-string	<local-url> [up to 99 characters]
	local-url	<cflash-id>/<file-path>
	cflash-id	cf1: | cf2: | cf3:

curve

Generates an ECDSA key with a specified curve.

Values secp256r1, secp384r1, secp521r1

key-size

Specifies the key size in bits.

Values 512 to 8192

Default 2048

type

Specifies the type of key.

Values rsa, dsa

Default rsa

Platforms

7705 SAR Gen 2

11.4 gen-local-cert-req

gen-local-cert-req

Syntax

gen-local-cert-req keypair *url-string* **subject-dn** *subject-dn* [**domain-name** *name*] [**ip-addr** *ip-address*] **file** *cert-req-file-url* [**hash-alg hash-algorithm**]

Context

[Tree] (admin>certificate gen-local-cert-req)

Full Context

admin certificate gen-local-cert-req

Description

This command generates a PKCS#10 formatted certificate request by using a local existing key pair file.

Parameters

url-string

Specifies the name of the keyfile in cf3:\system-pki\key that is used to generate a certificate request.

Values

url-string	<local-url> [up to 99 characters]
local-url	<cflash-id>/<file-path>
cflash-id	cf1: | cf2: | cf3:

subject-dn

Specifies the distinguish name that is used as the subject in a certificate request, including:

- C-Country

- ST-State
- O-Organization name
- OU-Organization Unit name
- CN-common name

This parameter is formatted as a text string including any of the above attributes. The attribute and its value is linked by using "=", and "," is used to separate different attributes.

For example: C=US,ST=CA,O=ALU,CN=SR12

Values attr1=val1,attr2=val2... where: attrN={C| ST| O| OU| CN}, 256 chars max

domain-name

Specifies a domain name string can be specified and included as the dNSName in the Subject Alternative Name extension of the certificate request.

ip-address

Specifies an IPv4 address string can be specified and included as the ipAddress in the Subject Alternative Name extension of the certificate request.

cert-req-file-url

Specifies the certificate URL. This URL could be either a local CF card path and filename to save the certificate request; or an FTP URL to upload the certificate request.

hash-algorithm

Specifies the hash algorithm to be used in a certificate request.

Values sha1, sha224, sha256, sha384, sha512

Platforms

7705 SAR Gen 2

11.5 general-port

general-port

Syntax

general-port *port-number*

no general-port

Context

[Tree] (config>system>snmp general-port)

Full Context

configure system snmp general-port

Description

This command configures the port number used to receive SNMP request messages and send replies.

For the port used for SNMP notifications, configure the **configure log snmp-trap-group trap-target port** command.

The **no** form of the command reverts to the default value.

Default

general-port 161

Parameters

port-number

Specifies the port number used to send SNMP traffic other than traps.

Values 1 to 65535

Platforms

7705 SAR Gen 2

11.6 generate-icmp

generate-icmp

Syntax

[**no**] **generate-icmp**

Context

[Tree] (config>service>vprn>static-route-entry>black-hole generate-icmp)

Full Context

configure service vprn static-route-entry black-hole generate-icmp

Description

This optional command causes the ICMP unreachable messages to be sent when received packets match the associated static route. By default, the ICMP unreachable messages for those types of static routes are not generated.

This command can only be associated with a static route that has a black-hole next-hop

The **no** form of this command removes the black-hole next-hop from static route configuration.

Default

no generate-icmp

Platforms

7705 SAR Gen 2

generate-icmp

Syntax

[**no**] **generate-icmp**

Context

[Tree] (config>router>static-route-entry>black-hole generate-icmp)

Full Context

configure router static-route-entry black-hole generate-icmp

Description

This optional command causes the ICMP unreachable messages to be sent when received packets match the associated static route. By default, the ICMP unreachable messages for those types of static routes are not generated.

This command can only be associated with a static route that has a blackhole next-hop

The **no** form of this command removes the black-hole nexthop from the static route configuration.

Default

no generate-icmp

Platforms

7705 SAR Gen 2

11.7 generate-traps

generate-traps

Syntax

[**no**] **generate-traps**

Context

[Tree] (config>system>network-element-discovery generate-traps)

Full Context

configure system network-element-discovery generate-traps

Description

This command configures whether traps are generated every time a node is updated, added, or removed from the OSPF opaque database (using LSA type 10 opaque update).

The **no** form of causes traps to not be generated for database changes.

Platforms

7705 SAR Gen 2

11.8 get

get

Syntax

[**no**] **get**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization get)

Full Context

configure system security profile netconf base-op-authorization get

Description

This command enables the NETCONF <get> RPC.

The **no** form of this command disables the RPC.

Default

no get

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

11.9 get-config

get-config

Syntax

[**no**] **get-config**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization get-config)

Full Context

configure system security profile netconf base-op-authorization get-config

Description

This command enables the NETCONF <get-config> RPC.

The **no** form of this command disables the RPC.

Default

no get-config

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

11.10 get-data

get-data

Syntax

[**no**] **get-data**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization get-data)

Full Context

configure system security profile netconf base-op-authorization get-data

Description

This command enables the NETCONF <get-data> RPC.

The **no** form of this command disables the RPC.

Default

no get-data

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

11.11 get-schema

get-schema

Syntax

[**no**] **get-schema**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization get-schema)

Full Context

configure system security profile netconf base-op-authorization get-schema

Description

This command enables the NETCONF <get-schema> RPC.

The **no** form of this command disables the RPC.

Default

no get-schema

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

11.12 gi-address

gi-address

Syntax

gi-address *ip-address*

no gi-address

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host gi-address)

Full Context

configure subscriber-mgmt local-user-db ipoe host gi-address

Description

This command allows selection of GI addresses based on the host entry in LUDB.

The gi-address must be a valid address (associated with an interface) within the routing context that received the DHCP message on the access side.

The **no** form of this command reverts to the default.

Parameters

ip-address

Specifies the IPv4 gi-address.

Values a.b.c.d

Platforms

7705 SAR Gen 2

gi-address

Syntax

gi-address *ip-address* [**src-ip-address**]

no gi-address

Context

[Tree] (config>service>vprn>if>dhcp gi-address)

[Tree] (config>service>ies>if>dhcp gi-address)

Full Context

configure service vprn interface dhcp gi-address

configure service ies interface dhcp gi-address

Description

This command configures the gateway interface address for the DHCP relay. A subscriber interface can include multiple group interfaces with multiple SAPs. The GI address is needed, when the router functions as a DHCP relay, to distinguish between the different subscriber interfaces and potentially between the group interfaces defined.

By default, the GI address used in the relayed DHCP packet is the primary IP address of a normal IES interface. Specifying the GI address allows the user to choose a secondary address. For group interfaces a GI address must be specified under the group interface DHCP context or subscriber-interface DHCP context in order for DHCP to function.

The **no** form of this command reverts to the default.

Parameters

ip-address

Specifies the host IP address to be used for DHCP relay packets.

Values a.b.c.d

src-ip-address

Specifies that this GI address is to be the source IP address for DHCP relay packets. This parameter is not applicable for PPPoE DHCP client messages (**dhcp client-applications ppp**).

Platforms

7705 SAR Gen 2

gi-address

Syntax

gi-address *ip-address*

no gi-address

Context

[Tree] (config>service>ies>if>sap>ipsec-gw>dhcp gi-address)

[Tree] (config>service>vprn>if>sap>ipsec-gw>dhcp gi-address)

Full Context

configure service ies interface sap ipsec-gw dhcp gi-address

configure service vprn interface sap ipsec-gw dhcp gi-address

Description

This command specifies the gateway IP address of the DHCPv4 packets sent by the system. IPsec DHCP Relay uses only the **gi-address** configuration found under the IPsec gateway and does not take into account **gi-address** with **src-ip-addr** configuration below other interfaces.

Default

no gi-address

Parameters

ip-address

Specifies the host IP address to be used for DHCP relay packets.

Platforms

7705 SAR Gen 2

gi-address

Syntax

gi-address *ip-address* [**src-ip-addr**]

no gi-address

Context

[Tree] (config>router>if>dhcp gi-address)

Full Context

configure router interface dhcp gi-address

Description

This command configures the gateway interface address for the DHCP relay. The GI address is needed, when the router functions as a DHCP relay, to distinguish between the different subscriber interfaces and potentially between the group interfaces defined.

Default

no gi-address

Parameters

ip-address

Specifies the host IP address to be used for DHCP relay packets.

src-ip-addr

Uses the GI address as the source IP.

Platforms

7705 SAR Gen 2

11.13 global

global

Syntax

global *file-url*

no global

Context

[Tree] (config>system>login-control>login-scripts global)

Full Context

configure system login-control login-scripts global

Description

This command enables an operator to define a common CLI script that executes when any user logs into a CLI session. This login exec script is executed when any user (authenticated by any means including local user database, TACACS+, or RADIUS) opens a CLI session. This allows a user, for example, to define a common set of CLI aliases that are made available on the router for all users. This global login exec script is executed before any user-specific login exec files that may be configured.

This CLI script executes in the context of the user who opens the CLI session. Any commands in the script that the user is not authorized to execute will fail.

The **no** form of this command disables the execution of a global login-script.

Default

no global

Parameters

file-url

The path or directory name.

Platforms

7705 SAR Gen 2

11.14 global-timeouts

global-timeouts

Syntax

global-timeouts

Context

[Tree] (config>system>management-interface>ops global-timeouts)

Full Context

configure system management-interface operations global-timeouts

Description

Commands in this context configure system timeout parameters for operational commands.

Timeout parameters provide default system-level control for various types of operational commands in model-driven interfaces. The timeout values are used when specific execution and retention timeouts are not requested for a specific operation.

Platforms

7705 SAR Gen 2

11.15 global-variables

global-variables

Syntax

global-variables

no global-variables

Context

[Tree] (config>router>policy-options global-variables)

Full Context

configure router policy-options global-variables

Description

This command enables the **global-variables** configuration context.

The **no** form of this command removes all global variables.

Platforms
7705 SAR Gen 2

11.16 gnmi

gnmi

Syntax
gnmi

Context
[Tree] (config>system>grpc gnmi)

Full Context
configure system grpc gnmi

Description
Commands in this context configure a gNMI service on gRPC.

Platforms
7705 SAR Gen 2

11.17 gnmi-capabilities

gnmi-capabilities

Syntax
gnmi-capabilities {**permit** | **deny**}

Context
[Tree] (config>system>security>profile>grpc>rpc-authorization gnmi-capabilities)

Full Context
configure system security profile grpc rpc-authorization gnmi-capabilities

Description
This command permits the use of Capability RPC for a user associated with the given format.

The **no** form of this command reverts to the default value.

Default

gnmi-capabilities permit

Parameters

permit

Specifies that the use of the Capability RPC is permitted.

deny

Specifies that the use of the Capability RPC is denied.

Platforms

7705 SAR Gen 2

11.18 gnmi-get

gnmi-get

Syntax

gnmi-get {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnmi-get)

Full Context

configure system security profile grpc rpc-authorization gnmi-get

Description

This command permits the use of Get RPC.

The **no** form of this command reverts to the default value.

Default

gnmi-get permit

Parameters

permit

Specifies that the use of the Get RPC is permitted.

deny

Specifies that the use of the Get RPC is denied.

Platforms
7705 SAR Gen 2

11.19 gnmi-set

gnmi-set

Syntax
gnmi-set {**permit** | **deny**}

Context
[Tree] (config>system>security>profile>grpc>rpc-authorization gnmi-set)

Full Context
configure system security profile grpc rpc-authorization gnmi-set

Description
This command permits the use of Set RPC.

The **no** form of this command reverts to the default value.

Default
gnmi-set permit

Parameters
permit
Specifies that the use of the Set RPC is permitted.

deny
Specifies that the use of the Set RPC is denied.

Platforms
7705 SAR Gen 2

11.20 gnmi-subscribe

gnmi-subscribe

Syntax
gnmi-subscribe {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnmi-subscribe)

Full Context

configure system security profile grpc rpc-authorization gnmi-subscribe

Description

This command permits the use of Subscribe RPC.

The **no** form of this command reverts to the default value.

Default

gnmi-subscribe permit

Parameters

permit

Specifies that the use of the Subscribe RPC is permitted.

deny

Specifies that the use of the Subscribe RPC is denied.

Platforms

7705 SAR Gen 2

11.21 gnoi-cert-mgmt-cangenerate

gnoi-cert-mgmt-cangenerate

Syntax

gnoi-cert-mgmt-cangenerate {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-cert-mgmt-cangenerate)

Full Context

configure system security profile grpc rpc-authorization gnoi-cert-mgmt-cangenerate

Description

This command permits the use of gNOI CanGenerateCSR RPCs for the user profile.

The **no** form of this command reverts to the default value.

Default

gnoi-cert-mgmt-cangenerate deny

Parameters

permit

Specifies that the use of the gNOI CanGenerateCSR RPCs for the user profile is permitted.

deny

Specifies that the use of the gNOI CanGenerateCSR RPCs for the user profile is denied.

Platforms

7705 SAR Gen 2

11.22 gnoi-cert-mgmt-getcert

gnoi-cert-mgmt-getcert

Syntax

gnoi-cert-mgmt-getcert {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-cert-mgmt-getcert)

Full Context

configure system security profile grpc rpc-authorization gnoi-cert-mgmt-getcert

Description

This command permits the use of gNOI GetCertificate RPCs for the user profile.

The **no** form of this command reverts to the default value.

Default

gnoi-cert-mgmt-getcert deny

Parameters

permit

Specifies that the use of the gNOI GetCertificate RPCs for the user profile is permitted.

deny

Specifies that the use of the gNOI GetCertificate RPCs for the user profile is denied.

Platforms
7705 SAR Gen 2

11.23 gnoi-cert-mgmt-install

gnoi-cert-mgmt-install

Syntax
gnoi-cert-mgmt-install {**permit** | **deny**}

Context
[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-cert-mgmt-install)

Full Context
configure system security profile grpc rpc-authorization gnoi-cert-mgmt-install

Description
This command permits the use of gNOI Install RPCs for the user profile.

The **no** form of this command reverts to the default value.

Default
gnoi-cert-mgmt-install deny

Parameters
permit
Specifies that the use of the gNOI Install RPCs for the user profile is permitted.

deny
Specifies that the use of the gNOI Install RPCs for the user profile is denied.

Platforms
7705 SAR Gen 2

11.24 gnoi-cert-mgmt-revoke

gnoi-cert-mgmt-revoke

Syntax
gnoi-cert-mgmt-revoke {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-cert-mgmt-revoke)

Full Context

configure system security profile grpc rpc-authorization gnoi-cert-mgmt-revoke

Description

This command permits or denies the use of gNOI RevokeCertificates RPCs for the user profile.

The **no** form of this command reverts to the default value.

Default

gnoi-cert-mgmt-revoke deny

Parameters

permit

Specifies that the use of gNOI RevokeCertificates RPCs for the user profile is permitted.

deny

Specifies that the use of gNOI RevokeCertificates RPCs for the user profile is denied.

Platforms

7705 SAR Gen 2

11.25 gnoi-cert-mgmt-rotate

gnoi-cert-mgmt-rotate

Syntax

gnoi-cert-mgmt-rotate {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-cert-mgmt-rotate)

Full Context

configure system security profile grpc rpc-authorization gnoi-cert-mgmt-rotate

Description

This command permits the use of gNOI Rotate RPCs for the user profile.

Default

gnoi-cert-mgmt-rotate deny

Parameters

permit

Specifies that the use of the gNOI Rotate RPCs for the user profile is permitted.

deny

Specifies that the use of the gNOI Rotate RPCs for the user profile is denied.

Platforms

7705 SAR Gen 2

11.26 gnoi-file-get

gnoi-file-get

Syntax

gnoi-file-get {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-file-get)

Full Context

configure system security profile grpc rpc-authorization gnoi-file-get

Description

This command permits the use of gNOI File Get RPC for a file from a target location.

Default

gnoi-file-get permit

Parameters

permit

Specifies that the use of the gNOI File Get RPC is permitted.

deny

Specifies that the use of the gNOI File Get RPC is denied.

Platforms

7705 SAR Gen 2

11.27 gnoi-file-put

gnoi-file-put

Syntax

gnoi-file-put {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-file-put)

Full Context

configure system security profile grpc rpc-authorization gnoi-file-put

Description

This command permits the use of gNOI File Put RPC to write to a file on a target location.

Default

gnoi-file-put permit

Parameters

permit

Specifies that the use of the gNOI File Put RPC is permitted.

deny

Specifies that the use of the gNOI File Put RPC is denied.

Platforms

7705 SAR Gen 2

11.28 gnoi-file-remove

gnoi-file-remove

Syntax

gnoi-file-remove {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-file-remove)

Full Context

configure system security profile grpc rpc-authorization gnoi-file-remove

Description

This command permits the use of gNOI File Remove RPC to remove a file from the specified target location.

Default

gnoi-file-remove permit

Parameters

permit

Specifies that the use of the gNOI File Remove RPC is permitted.

deny

Specifies that the use of the gNOI File Remove RPC is denied.

Platforms

7705 SAR Gen 2

11.29 gnoi-file-stat

gnoi-file-stat

Syntax

gnoi-file-stat {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-file-stat)

Full Context

configure system security profile grpc rpc-authorization gnoi-file-stat

Description

This command permits the use of gNOI File Stat RPC to retrieve metadata for a file from the specified target location.

Default

gnoi-file-stat permit

Parameters

permit

Specifies that the use of the gNOI File Stat RPC is permitted.

deny

Specifies that the use of the gNOI File Stat RPC is denied.

Platforms

7705 SAR Gen 2

11.30 gnoi-file-transfertoremote

gnoi-file-transfertoremote

Syntax

gnoi-file-transfertoremote {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-file-transfertoremote)

Full Context

configure system security profile grpc rpc-authorization gnoi-file-transfertoremote

Description

This command permits the use of the gNOI File TransferToRemote RPC to transfer the file from the target node to a specified remote location.

Default

gnoi-file-transfertoremote permit

Parameters

permit

Specifies that the use of the gNOI File TransferToRemote RPC is permitted.

deny

Specifies that the use of the gNOI File TransferToRemote RPC is denied.

Platforms

7705 SAR Gen 2

11.31 gnoi-system-cancelreboot

gnoi-system-cancelreboot

Syntax

gnoi-system-cancelreboot {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-system-cancelreboot)

Full Context

configure system security profile grpc rpc-authorization gnoi-system-cancelreboot

Description

This command permits the use of gNOI System CancelReboot RPC for a user-given profile.

Default

gnoi-system-cancelreboot deny

Parameters

permit

Specifies that the use of gNOI System CancelReboot RPC is permitted.

deny

Specifies that the use of gNOI System CancelReboot RPC is denied.

Platforms

7705 SAR Gen 2

11.32 gnoi-system-ping

gnoi-system-ping

Syntax

gnoi-system-ping {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-system-ping)

Full Context

configure system security profile grpc rpc-authorization gnoi-system-ping

Description

This command permits the use of the gNOI Ping RPC to execute the ping command on the target node and stream back the results.

Default

gnoi-system-ping permit

Parameters

permit

Specifies that the use of the gNOI Ping RPC is permitted.

deny

Specifies that the use of the gNOI Ping RPC is denied.

Platforms

7705 SAR Gen 2

11.33 gnoi-system-reboot

gnoi-system-reboot

Syntax

gnoi-system-reboot {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-system-reboot)

Full Context

configure system security profile grpc rpc-authorization gnoi-system-reboot

Description

This command permits the use of gNOI System Reboot RPC for a user-given profile.

The **no** form of this command reverts to the default value.

Default

gnoi-system-reboot deny

Parameters

permit

Specifies that the use of gNOI System Reboot RPC is permitted.

deny

Specifies that the use of gNOI System Reboot RPC is denied.

Platforms

7705 SAR Gen 2

11.34 gnoi-system-rebootstatus

gnoi-system-rebootstatus

Syntax

gnoi-system-rebootstatus {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-system-rebootstatus)

Full Context

configure system security profile grpc rpc-authorization gnoi-system-rebootstatus

Description

This command permits the use of gNOI System RebootStatus RPC for a user-given profile.

The **no** form of this command reverts to the default value.

Default

gnoi-system-rebootstatus deny

Parameters

permit

Specifies that the use of gNOI System RebootStatus RPC is permitted for a user-given profile.

deny

Specifies that the use of gNOI System RebootStatus RPC is denied.

Platforms

7705 SAR Gen 2

11.35 gnoi-system-setpackage

gnoi-system-setpackage

Syntax

gnoi-system-setpackage {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-system-setpackage)

Full Context

configure system security profile grpc rpc-authorization gnoi-system-setpackage

Description

This command permits the use of gNOI System SetPackage RPC for a user-given profile.

The **no** form of this command reverts to the default value.

Default

gnoi-system-setpackage deny

Parameters

deny

Specifies that the use of gNOI System SetPackage RPC is denied.

permit

Specifies that the use of gNOI System SetPackage RPC is permitted.

Platforms

7705 SAR Gen 2

11.36 gnoi-system-switchcontrolprocessor

gnoi-system-switchcontrolprocessor

Syntax

gnoi-system-switchcontrolprocessor {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-system-switchcontrolprocessor)

Full Context

configure system security profile grpc rpc-authorization gnoi-system-switchcontrolprocessor

Description

This command permits the use of gNOI System SwitchControlProcessor RPC for a user-given profile.

The **no** form of this command reverts to the default value.

Default

gnoi-system-switchcontrolprocessor deny

Parameters

deny

Specifies that the use of gNOI System SwitchControlProcessor RPC is denied.

permit

Specifies that the use of gNOI System SwitchControlProcessor RPC is permitted.

Platforms

7705 SAR Gen 2

11.37 gnoi-system-time

gnoi-system-time

Syntax

gnoi-system-time {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-system-time)

Full Context

configure system security profile grpc rpc-authorization gnoi-system-time

Description

This command permits the use of the gNOI Time RPC to return the current time on the target node.

Default

gnoi-system-time permit

Parameters

permit

Specifies that the use of the gNOI Time RPC is permitted.

deny

Specifies that the use of the gNOI Time RPC is denied.

Platforms

7705 SAR Gen 2

11.38 gnoi-system-traceroute

gnoi-system-traceroute

Syntax

gnoi-system-traceroute {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization gnoi-system-traceroute)

Full Context

configure system security profile grpc rpc-authorization gnoi-system-traceroute

Description

This command permits the use of the gNOI Traceroute RPC to execute the traceroute command on the target node and stream back the results.

Default

gnoi-system-traceroute permit

Parameters

permit

Specifies that the use of the gNOI Traceroute RPC is permitted.

deny

Specifies that the use of the gNOI Traceroute RPC is denied.

Platforms

7705 SAR Gen 2

11.39 goto

goto

Syntax

goto *line*

Context

[Tree] (candidate goto)

Full Context

candidate goto

Description

This command changes the edit point of the candidate configuration. The edit point is the point after which new commands are inserted into the candidate configuration as an operator navigates the CLI and issues commands in edit-cfg mode.

Parameters

line

Indicates which line to change starting at the point indicated by the following options.

Values line, offset, **first**, **edit-point**, **last**

line	absolute line number
offset	relative line number to current edit point. Prefixed with '+' or '-'
first	keyword - first line
edit-point	keyword - current edit point
last	keyword - last line that is not 'exit'

Platforms

7705 SAR Gen 2

11.40 gr-helper

gr-helper

Syntax

gr-helper [**enable** | **disable**]

Context

[Tree] (config>router>rsvp>if gr-helper)

Full Context

configure router rsvp interface gr-helper

Description

This command enables the RSVP Graceful Restart Helper feature.

The RSVP-TE Graceful Restart helper mode allows the SR OS based system (the helper node) to provide another router that has requested it (the restarting node) a grace period, during which the system will continue to use RSVP sessions to neighbors requesting the grace period. This is typically used when another router is rebooting its control plane but its forwarding plane is expected to continue to forward traffic based on the previously available Path and Resv states.

The user can enable Graceful Restart helper on each RSVP interface separately. When the GR helper feature is enabled on an RSVP interface, the node starts inserting a new Restart_Cap Object in the Hello packets to its neighbor. The restarting node does the same and indicates to the helper node the desired Restart Time and Recovery Time.

The GR Restart helper consists of a couple of phases. Once it loses Hello communication with its neighbor, the helper node enters the Restart phase. During this phase, it preserves the state of all RSVP sessions to its neighbor and waits for a new Hello message.

Once the Hello message is received indicating the restarting node preserved state, the helper node enters the recovery phase in which it starts refreshing all the sessions that were preserved. The restarting node will activate all the stale sessions that are refreshed by the helper node. Any Path state which did not get a Resv message from the restarting node once the Recovery Phase time is over is considered to have expired and is deleted by the helper node causing the proper Path Tear generation downstream.

The duration of the restart phase (recovery phase) is equal to the minimum of the neighbor’s advertised Restart Time (Recovery Time) in its last Hello message and the locally configured value of the max-restart (max-recovery) parameter.

When GR helper is enabled on an RSVP interface, its procedures apply to the state of both P2P and P2MP RSVP LSP to a neighbor over this interface.

Default

disable

Platforms

7705 SAR Gen 2

11.41 gr-helper-time

gr-helper-time

Syntax

gr-helper-time max-recovery *recovery-interval* **max-restart** *restart-interval*

no gr-helper-time

Context

[Tree] (config>router>rsvp gr-helper-time)

Full Context

configure router rsvp gr-helper-time

Description

This command configures the local values for the max-recovery and the max-restart intervals used in the RSVP Graceful Restart Helper feature.

The values are configured globally in RSVP but separate instances of the timers are applied to each RSVP interface that has the RSVP Graceful Restart Helper enabled.

The **no** version of this command re-instates the default value for the delay timer.

Default

gr-helper-time max-recovery 300 max-restart 120

Parameters

recovery-interval

Specifies the max recovery interval value in seconds.

Values 1 to 1800

restart-interval

Specifies the max restart interval value in seconds.

Values 1 to 300

Platforms

7705 SAR Gen 2

11.42 graceful-restart

graceful-restart

Syntax

[**no**] **graceful-restart**

Context

[Tree] (config>service>vprn>bgp>group graceful-restart)

[Tree] (config>service>vprn>bgp graceful-restart)

[Tree] (config>service>vprn>bgp>group>neighbor graceful-restart)

Full Context

configure service vprn bgp group graceful-restart

configure service vprn bgp graceful-restart

configure service vprn bgp group neighbor graceful-restart

Description

This command enables BGP graceful restart helper procedures (the "receiving router" role defined in the standard) for address families included in the GR capabilities of both peers. In a VPRN, SR OS can support GR helper functionality for IPv4, IPv6, label-ipv4, flow-ipv4 (IPv4 FlowSpec) and flow-ipv6 (IPv6 FlowSpec) routes.

When a neighbor covered by the GR helper mode restarts its control plane, forwarding can continue uninterrupted while the session is re-established and routes are re-learned.

The **no** form of this command disables graceful restart.

Platforms

7705 SAR Gen 2

graceful-restart

Syntax

[**no**] **graceful-restart**

Context

[Tree] (config>service>vprn>isis graceful-restart)

Full Context

configure service vprn isis graceful-restart

Description

This command enables IS-IS graceful restart (GR) to minimize service interruption. When the control plane of a GR-capable router fails or restarts, the neighboring routers (GR helpers) temporarily preserve IS-IS forwarding information. Traffic continues to be forwarded to the restarting router using the last known forwarding tables. If the control plane of the restarting router becomes operationally and administratively up within the grace period, the restarting router resumes normal IS-IS operation. If the grace period expires, then the restarting router is presumed inactive and the IS-IS topology is recalculated to route traffic around the failure.

The **no** form of this command disables graceful restart and removes the graceful restart configuration from the IS-IS instance.

Default

no graceful-restart

Platforms

7705 SAR Gen 2

graceful-restart

Syntax

[**no**] **graceful-restart**

Context

[Tree] (config>service>vprn>ospf3 graceful-restart)

[Tree] (config>service>vprn>ospf graceful-restart)

Full Context

configure service vprn ospf3 graceful-restart

configure service vprn ospf graceful-restart

Description

This command enables OSPF graceful restart (GR) to minimize service interruption.

When the control plane of a GR-capable router fails or restarts, the neighboring routers (GR helpers) temporarily preserve OSPF forwarding information. Traffic continues to be forwarded to the restarting router using the last known forwarding tables. If the control plane of the restarting router becomes operationally and administratively up within the grace period, the restarting router resumes normal OSPF operation. If the grace period expires, the restarting router is presumed inactive and the OSPF topology is recalculated to route traffic around the failure.

The **no** form of this command disables GR and removes the GR configuration from the OSPF instance.

Default

no graceful-restart

Platforms

7705 SAR Gen 2

graceful-restart

Syntax

[**no**] **graceful-restart**

Context

[Tree] (config>router>ldp graceful-restart)

Full Context

configure router ldp graceful-restart

Description

This command enables graceful restart helper.

The **no** form of this command disables graceful restart.

Graceful restart helper configuration changes, enable/disable, or change of a parameter will cause the LDP session to bounce.

Default

no graceful-restart (disabled) — Graceful-restart must be explicitly enabled.

Platforms

7705 SAR Gen 2

graceful-restart

Syntax

[**no**] **graceful-restart**

Context

[Tree] (config>router>bgp>group graceful-restart)

[Tree] (config>router>bgp>group>neighbor graceful-restart)

[Tree] (config>router>bgp graceful-restart)

Full Context

configure router bgp group graceful-restart

configure router bgp group neighbor graceful-restart

configure router bgp graceful-restart

Description

This command enables BGP graceful restart helper procedures (the "receiving router" role defined in the standard) for address families included in the GR capabilities of both peers. SR OS can support GR helper functionality for IPv4, IPv6, VPN-IPv4, VPN-IPv6, Label-IPv4, Label-IPv6, L2-VPN, Route-Target (RTC), Flow-IPv4 (IPv4 FlowSpec) and Flow-IPv6 (IPv6 FlowSpec) routes.

If a neighbor covered by the GR helper mode restarts its control plane, forwarding can continue uninterrupted while the session is re-established and routes are re-learned.

The **no** form of this command disables graceful restart.

Default

no graceful-restart

Platforms

7705 SAR Gen 2

graceful-restart

Syntax

graceful-restart [**neighbor** *ip-address* | **group** *name*]

no graceful-restart

Context

[Tree] (debug>router>bgp graceful-restart)

Full Context

debug router bgp graceful-restart

Description

This command enables debugging for BGP graceful restart.

The **no** form of this command disables the debugging.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x [-interface] (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H
- d: [0 to 255]D

- interface: up to 32 characters for link local addresses

group *name*

Debugs only events affecting the specified peer group name, up to 64 characters, and associated neighbors.

Platforms

7705 SAR Gen 2

graceful-restart

Syntax

[**no**] **graceful-restart**

Context

[Tree] (config>router>isis graceful-restart)

Full Context

configure router isis graceful-restart

Description

This command enables IS-IS graceful restart (GR) to minimize service interruption. When the control plane of a GR-capable router fails or restarts, the neighboring routers (GR helpers) temporarily preserve IS-IS forwarding information. Traffic continues to be forwarded to the restarting router using the last known forwarding tables. If the control plane of the restarting router becomes operationally and administratively up within the grace period, the restarting router resumes normal IS-IS operation. If the grace period expires, then the restarting router is presumed inactive and the IS-IS topology is recalculated to route traffic around the failure.

The **no** form of this command disables graceful restart and removes the graceful restart configuration from the IS-IS instance.

Default

no graceful-restart

Platforms

7705 SAR Gen 2

graceful-restart

Syntax

[**no**] **graceful-restart**

Context

[Tree] (debug>router>isis graceful-restart)

Full Context

debug router isis graceful-restart

Description

This command enables debugging for IS-IS graceful-restart.

The **no** form of the command disables debugging.

Platforms

7705 SAR Gen 2

graceful-restart

Syntax

[**no**] **graceful-restart**

Context

[Tree] (config>router>ospf graceful-restart)

[Tree] (config>router>ospf3 graceful-restart)

Full Context

configure router ospf graceful-restart

configure router ospf3 graceful-restart

Description

This command enables OSPF graceful restart (GR) to minimize service disruption. When the control plane of a GR-capable router fails or restarts, the neighboring routers (GR helpers) temporarily preserve OSPF forwarding information. Traffic continues to be forwarded to the restarting router using the last known forwarding tables. If the control plane of the restarting router comes back up within the grace period, the restarting router resumes normal OSPF operation. If the grace period expires, then the restarting router is presumed inactive and the OSPF topology is recalculated to route traffic around the failure.

The **no** form of this command disables graceful restart and removes the graceful restart configuration from the OSPF instance.

Default

no graceful-restart

Platforms

7705 SAR Gen 2

graceful-restart

Syntax

[**no**] **graceful-restart**

Context

[Tree] (debug>router>ospf graceful-restart)

[Tree] (debug>router>ospf3 graceful-restart)

Full Context

debug router ospf graceful-restart

debug router ospf3 graceful-restart

Description

This command enables debugging for OSPF and OSPF3 graceful restart.

Platforms

7705 SAR Gen 2

11.43 graceful-shutdown

graceful-shutdown

Syntax

[**no**] **graceful-shutdown**

Context

[Tree] (config>router>rsvp graceful-shutdown)

[Tree] (config>router>rsvp>interface graceful-shutdown)

Full Context

configure router rsvp graceful-shutdown

configure router rsvp interface graceful-shutdown

Description

This command initiates a graceful shutdown of the specified RSVP interface or all RSVP interfaces on the node if applied at the RSVP level. These are referred to as maintenance interface and maintenance node, respectively.

To initiate a graceful shutdown the maintenance node generates a PathErr message with a specific error sub-code of Local Maintenance on TE Link required for each LSP that is exiting the maintenance interface.

The node performs a single make-before-break attempt for all adaptive CSPF LSPs it originates and LSP paths using the maintenance interfaces. If an alternative path for an affected LSP is not found, then the LSP is maintained on its current path. The maintenance node also tears down and re-signals any detour LSP path using listed maintenance interfaces as soon as they are not active.

The maintenance node floods an IGP TE LSA/LSP containing Link TLV for the links under graceful shutdown with TE metric set to 0xffffffff and Unreserved Bandwidth parameter set to zero (0).

A head-end LER node, upon receipt of the PathErr message performs a single make-before-break attempt on the affected adaptive CSPF LSP. If an alternative path is not found, then the LSP is maintained on its current path.

A node does not take any action on the paths of the following originating LSPs after receiving the PathErr message:

a. An adaptive CSPF LSP for which the PathErr indicates a node address in the address list and the node corresponds to the destination of the LSP. In this case, there are no alternative paths which can be found.

b. An adaptive CSPF LSP whose path has explicit hops defined using the listed maintenance interface(s)/ node(s).

c. A CSPF LSP with the adaptive option disabled and which current path is over the listed maintenance interfaces in the PathErr message. These are not subject to make-before-break.

d. A non CSPF LSP which current path is over the listed maintenance interfaces in the PathErr message.

The head-end LER node upon receipt of the updates IPG TE LSA/LSP for the maintenance interfaces updates the TE database. This information will be used at the next scheduled CSPF computation for any LSP which path may traverse any of the maintenance interfaces.

The **no** form of this command disables the graceful shutdown operation at the RSVP interface level or at the RSVP level. The configured TE parameters of the maintenance links are restored and the maintenance node floods the links.

Platforms

7705 SAR Gen 2

11.44 graft

graft

Syntax

graft [**source** *ip-address*] [**group** *grp-ip-address*] [**detail**]

no graft

Context

[Tree] (debug>router>pim graft)

Full Context

debug router pim graft

Description

This command enables debugging for PIM grafts.

The **no** form of this command disables PIM graft debugging.

Parameters

ip-address

Debugs graft information associated with the specified source.

Values source address (ipv4, ipv6)

grp-ip-address

Debugs graft information associated with the specified group.

Values multicast group address (ipv4, ipv6)

detail

Debugs detailed graft information.

Platforms

7705 SAR Gen 2

11.45 gre

gre

Syntax

[**no**] **gre**

Context

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel>resolution-filter gre)

Full Context

configure service vprn bgp-ipvpn mpls auto-bind-tunnel resolution-filter gre

Description

This command enables setting the tunnel type for the auto bind tunnel.

The **gre** encapsulation of the MPLS service packet uses the base 4-byte header as per RFC 2890. The optional fields Checksum (plus Reserved field), Key, and Sequence Number are not inserted.

The **no** form of this command disables the setting the tunnel type for the auto bind tunnel.

Default

no gre

Platforms

7705 SAR Gen 2

gre

Syntax

gre

Context

[Tree] (config>service>vprn>auto-bind-tunnel>res-filter gre)

Full Context

configure service vprn auto-bind-tunnel resolution-filter gre

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

11.46 gre-header

gre-header

Syntax

gre-header send-key *send-key* **receive-key** *receive-key*

no gre-header

Context

[Tree] (config>service>ies>if>sap>ip-tunnel gre-header)

[Tree] (config>service>vprn>if>sap>ip-tunnel gre-header)

Full Context

configure service ies interface sap ip-tunnel gre-header

configure service vprn interface sap ip-tunnel gre-header

Description

This command configures the type of the IP tunnel. If the **gre-header** command is configured then the tunnel is a GRE tunnel with a GRE header inserted between the outer and inner IP headers. If the **no** form of this command is configured then the tunnel is a simple IP-IP tunnel.

Default

no gre-header

Parameters

send-key *send-key*

Specifies a 32-bit unsigned integer.

Values 0 to 4294967295

receive-key *receive-key*

Specifies a 32-bit unsigned integer.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

11.47 group

group

Syntax

[**no**] **group** *group-name*

Context

[Tree] (config>service>vprn>rip group)

Full Context

configure service vprn rip group

Description

This command creates a context for configuring a RIP group of neighbors. RIP groups are a way of logically associating RIP neighbor interfaces to facilitate a common configuration for RIP interfaces.

The **no** form of this command deletes the RIP neighbor interface group. Deleting the group also removes the RIP configuration of all the neighbor interfaces currently assigned to this group.

Default

no group

Parameters

group-name

The RIP group name. Allowed values are any string, up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

group

Syntax

[**no**] **group** *name*

Context

[Tree] (config>router>bgp group)

Full Context

configure router bgp group

Description

Commands in this context configure a BGP peer group.

The **no** form of this command deletes the specified peer group and all configurations associated with the peer group. The group must be shut down before it can be deleted.

Default

no group

Parameters

name

Specifies the peer group name. Allowed values are any string, up to 64 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

group

Syntax

[**no**] **group** *grp-ip-address*

[**no**] **group** *grp-ipv6-address*

Context

[Tree] (config>service>vpls>mesh-sdp>mld-snooping>static group)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping>static group)

[Tree] (config>service>vpls>sap>mld-snooping>static group)

[Tree] (config>service>vpls>spoke-sdp>mld-snooping>static group)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping>static group)

[Tree] (config>service>vpls>sap>igmp-snooping>static group)

Full Context

configure service vpls mesh-sdp mld-snooping static group

configure service vpls spoke-sdp igmp-snooping static group

configure service vpls sap mld-snooping static group

configure service vpls spoke-sdp mld-snooping static group

configure service vpls mesh-sdp igmp-snooping static group

configure service vpls sap igmp-snooping static group

Description

Commands in this context add a static multicast group as a (*, G) or as one or more (S,G) records. When a static MLD or IGMP group is added, multicast data for that (*,G) or (S,G) is forwarded to the specific SAP or SDP without receiving any membership report from a host.

Parameters

grp-ip-address

Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

grp-ipv6-address

Specifies an MLD multicast group address that receives data on an interface. The IP address must be unique for each static group.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

group

Syntax

group *name* [**esm-dynamic-peer**]

no group *name*

Context

[Tree] (config>service>vprn>bgp group)

Full Context

configure service vprn bgp group

Description

This command creates a context to configure a BGP peer group.

The **no** form of this command deletes the specified peer group and all configurations associated with the peer group. The group must be shut down before it can be deleted.

Parameters

name

Specifies the peer group name. Allowed values is a string up to 64 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

esm-dynamic-peer

Specifies that the given BGP group is used by BGP peers created dynamically based on subscriber-hosts pointing to corresponding BGP peering policy. There can be only one BGP group with this flag set in any given VPRN. No BGP neighbors can be manually configured in a BGP group with this flag set.

Default disabled

Platforms

7705 SAR Gen 2

group

Syntax

[**no**] **group** *grp-ip-address*

[**no**] **group start** *grp-ip-address* **end** *grp-ip-address* [**step** *ip-address*]

Context

[Tree] (config>service>vprn>igmp>if>static group)

Full Context

configure service vprn igmp interface static group

Description

This command adds a static multicast group either as a (*,G) or one or more (S,G) records. Use IGMP static group memberships to test multicast forwarding without a receiver host. When IGMP static groups are enabled, data is forwarded to an interface without receiving membership reports from host members.

When static IGMP group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static IGMP group entries do not generate join messages toward the RP.

Parameters

grp-ip-address

Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group. The address must be in dotted decimal notation.

start *grp-ip-address*

Specifies the start multicast group address.

end *grp-ip-address*

Specifies the end multicast group address.

step *ip-address*

Specifies the step increment.

Platforms

7705 SAR Gen 2

group

Syntax

[**no**] **group** *grp-ipv6-address*

[**no**] **group start** *grp-ipv6-address* **end** *grp-ipv6-address* [**step** *ipv6-address*]

Context

[Tree] (config>service>vprn>mld>if>static group)

Full Context

configure service vprn mld interface static group

Description

Commands in this context add a static multicast group either as a (*,G) or one or more (S,G) records. Use MLD static group memberships to test multicast forwarding without a receiver host. When MLD static groups are enabled, data is forwarded to an interface without receiving membership reports from host members.

When static MLD group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static MLD group entries do not generate join messages toward the RP.

The **no** form of this command removes the IPv6 address from the configuration.

Parameters

grp-ipv6-address

Specifies an MLD multicast group address that receives data on an interface. The IP address must be unique for each static group.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

start *grp-ipv6-address*

Specifies the start multicast group address.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

end *grp-ipv6-address*

Specifies the end multicast group address.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

step *ipv6-address*

Specifies the step increment.

Platforms

7705 SAR Gen 2

group

Syntax

[**no**] **group** *grp-ip-address*

[**no**] **group start** *grp-ip-address* **end** *grp-ip-address* [**step** *ip-address*]

Context

[Tree] (config>router>igmp>if>static group)

Full Context

configure router igmp interface static group

Description

Commands in this context add a static multicast group either as a (*,G) or one or more (S,G) records. Use IGMP static group memberships to test multicast forwarding without a receiver host. When IGMP static groups are enabled, data is forwarded to an interface without receiving membership reports from host members.

When static IGMP group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static IGMP group entries do not generate join messages toward the RP.

Parameters

ip-address

Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

start *grp-ip-address*

Specifies the start multicast group address.

end *grp-ip-address*

Specifies the end multicast group address.

step *ip-address*

Specifies the step increment.

Platforms

7705 SAR Gen 2

group

Syntax

[**no**] **group** *grp-ipv6-address*

[**no**] **group start** *grp-ipv6-address* **end** *grp-ipv6-address* [**step** *ipv6-address*]

Context

[Tree] (config>router>mld>if>static group)

Full Context

configure router mld interface static group

Description

Commands in this context add a static multicast group either as a (*,G) or one or more (S,G) records. Use MLD static group memberships to test multicast forwarding without a receiver host. When MLD static groups are enabled, data is forwarded to an interface without receiving membership reports from host members.

When static MLD group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static MLD group entries do not generate join messages toward the RP.

The **no** form of this command removes the IPv6 address from the configuration.

Parameters

grp-ipv6-address

Specifies an MLD multicast group address that receives data on an interface. The IP address must be unique for each static group.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

start *grp-ipv6-address*

Specifies the start multicast group address.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

end *grp-ipv6-address*

Specifies the end multicast group address.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

step *ipv6-address*

Specifies the step increment.

Platforms

7705 SAR Gen 2

group

Syntax

group *group-name*

no group

Context

[Tree] (config>system>security>user>snmp group)

Full Context

configure system security user snmp group

Description

This command associates (or links) a user to a group name. The group name must be configured with the **config>system>security>user >snmp>group** command. The **config>system>security>user access** command links the group with one or more views, security model (s), security level (s), and read, write, and notify permissions.

Parameters

group-name

Enter the group name (between 1 and 32 alphanumeric characters) that is associated with this user. A user can be associated with one group-name per security model.

Platforms

7705 SAR Gen 2

group

Syntax

[**no**] **group** *group-name*

Context

[Tree] (config>router>rip group)

[Tree] (config>router>ripng group)

Full Context

configure router rip group

configure router ripng group

Description

This command creates a context for configuring a RIP group of neighbor interfaces.

RIP groups are a way of logically associating RIP neighbor interfaces to facilitate a common configuration for RIP interfaces.

The **no** form of the command deletes the RIP neighbor interface group. Deleting the group will also remove the RIP configuration of all the neighbor interfaces currently assigned to this group.

Default

no group

Parameters

group-name

Specifies the RIP group name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

group

Syntax

group down *time* | **no group down**

group up *time* | **no group up**

Context

[Tree] (config>service>oper-group>hold-time group)

Full Context

configure service oper-group hold-time group

Description

The **group down** form of the command configures the number of seconds to wait before notifying clients monitoring this group when its operational status transitions from up to down.

The **group up** form of the command configures the number of seconds to wait before notifying clients monitoring this group when its operational status transitions from down to up. A value of zero indicates that transitions are reported immediately to monitoring clients. The up time option is a must to achieve fast convergence: when the group comes up, the monitoring MH site that tracks the group status may wait without impacting the overall convergence; there is usually a pair MH site that is already handling the traffic.

The **no** form of the command sets the values back to the default.

Default

group down 0

group up 4

Parameters

time

Specifies the group up or group down time value.

Values 0 to 3600

Platforms

7705 SAR Gen 2

11.48 group-address

group-address

Syntax

group-address *prefix-list-name*

no group-address

Context

[Tree] (config>router>policy-options>policy-statement>entry>from group-address)

Full Context

configure router policy-options policy-statement entry from group-address

Description

This command specifies the multicast group-address prefix list containing multicast group-addresses that are embedded in the join or prune packet as a filter criterion. The prefix list must be configured prior to entering this command. Prefix lists are configured in the **config>router>policy-options>prefix-list** context.

The **no** form of this command removes the criterion from the configuration.

Default

no group-address

Parameters

prefix-list-name

Specifies the prefix-list name. Allowed values are any string up to 64 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

The *prefix-list-name* is defined in the **config>router>policy-options>prefix-list** context.

Platforms

7705 SAR Gen 2

11.49 group-encryption

group-encryption

Syntax

[**no**] **group-encryption**

Context

[Tree] (config>router>interface group-encryption)

Full Context

configure router interface group-encryption

Description

This command enables NGE on the router interface. When NGE is enabled on the interface, all received Layer 3 packets that have the protocol ID configured as ESP are considered to be NGE packets and must be encrypted using a valid set of keys from any preconfigured key group on the system.

The **no** form of this command disables NGE on the interface. NGE cannot be disabled unless all key groups and IP exception filters are removed.

Default

no group-encryption

Platforms

7705 SAR Gen 2

group-encryption

Syntax

group-encryption

Context

[Tree] (config group-encryption)

Full Context

configure group-encryption

Description

Commands in this context configure group encryption parameters.

Platforms

7705 SAR Gen 2

11.50 group-encryption-label

group-encryption-label

Syntax

group-encryption-label *encryption-label*

no group-encryption-label

Context

[Tree] (config>grp-encryp group-encryption-label)

Full Context

configure group-encryption group-encryption-label

Description

This command configures the group encryption label used to identify when an MPLS payload is encrypted. This label must be unique network-wide and must be configured consistently on all nodes participating in a network group encryption domain. The label cannot be changed or deleted when there are any key groups configured on the node.

The **no** form of the command reverts to the default setting.

Parameters

encryption-label

The network-wide, unique reserved MPLS label for group encryption.

Values 32 to 2047

Platforms

7705 SAR Gen 2

11.51 group-interface

group-interface

Syntax

[**no**] **group-interface** [**fwd-service** *service-id*] [*ip-int-name*]

Context

[Tree] (debug>router>igmp group-interface)

Full Context

debug router igmp group-interface

Description

This command enables debugging for IGMP group-interface.

The **no** form of the command disables debugging.

Parameters

service-id

Debugs information associated with the service ID.

Values service-id: 1 to 2148278386

svc-name: up to 64 characters.

ip-int-name

Debugs information associated with the specified IP interface name.

Values IP interface address

Platforms

7705 SAR Gen 2

11.52 group-list

group-list

Syntax

group-list *name*

no group-list

Context

[Tree] (config>system>security>tls>client-tls-profile group-list)

Full Context

configure system security tls client-tls-profile group-list

Description

This command assigns an existing TLS 1.3 group list to the TLS client profile.

The **no** form of this command removes the group list from the client profile.

Default

no group-list

Parameters

name

Specifies the name of the group list, up to 32 characters.

Platforms

7705 SAR Gen 2

group-list

Syntax

group-list *name*

no group-list

Context

[Tree] (config>system>security>tls>server-tls-profile group-list)

Full Context

configure system security tls server-tls-profile group-list

Description

This command assigns an existing TLS 1.3 group list to the TLS server profile.

The **no** form of this command removes the group list from the server profile.

Default

no group-list

Parameters

name

Specifies the name of the group list, up to 32 characters.

Platforms

7705 SAR Gen 2

11.53 group-name

group-name

Syntax

group-name *group-name* **value** *group-value*

no group-name *group-name*

Context

[Tree] (config>service>sdp-group group-name)

Full Context

configure service sdp-group group-name

Description

This command defines SDP administrative groups, referred to as SDP admin groups.

SDP admin groups provide a way for services using a pseudowire template to automatically include or exclude specific provisioned SDPs. SDPs sharing a specific characteristic or attribute can be made members of the same admin group. When users configure a pseudowire template, they can include and/or exclude one or more admin groups. When the service is bound to the pseudowire template, the SDP selection rules will enforce the admin group constraints specified in the **sdp-include** and **sdp-exclude** commands.

A maximum of 32 admin groups can be created. The group value ranges from zero (0) to 31. It is uniquely associated with the group name at creation time. If the user attempts to configure another group name for a group value that is already assigned to an existing group name, the SDP admin group creation is failed. The same happens if the user attempts to configure an SDP admin group with a new name but associates it to a group value already assigned to an existing group name.

The **no** option of this command deletes the SDP admin group but is only allowed if the group-name is not referenced in a PW template or SDP.

Parameters

group-name

Specifies the name of the SDP admin group. A maximum of 32 characters can be entered.

group-value

Specifies the group value associated with this SDP admin group. This value is unique within the system.

Values 0 to 31

Platforms

7705 SAR Gen 2

11.54 group-prefix

group-prefix

Syntax

[**no**] **group-prefix** *grp-ipv6-address*|*prefix-length*

Context

[Tree] (config>router>pim>rp>ipv6>static>address group-prefix)

[Tree] (config>router>pim>rp>static>address group-prefix)

Full Context

configure router pim rp ipv6 static address group-prefix

configure router pim rp static address group-prefix

Description

This command specifies the range of multicast group addresses which should be used by the router as the Rendezvous Point (RP). The **config**>**router**>**pim**>**rp**>**static**> **address** *a.b.c.d* implicitly defaults to deny all for all multicast groups (224.0.0.0/4). A group-prefix must be specified for that static address. This command does not apply to the whole group range.

The **no** form of this command removes the group-prefix from the configuration.

Parameters

grp-ipv6-address

Specifies the multicast group IPv6 address expressed in dotted decimal notation.

Values grp-ipv6-address x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x [0..FFFF]H
d [0..255]D

prefix-length

Specifies the prefix length of the IPv6 address.

Values 8 to 128

Platforms

7705 SAR Gen 2

11.55 group-range

group-range

Syntax

[**no**] **group-range** {*ipv6-address*/*prefix-length*}

Context

[Tree] (config>service>vprn>pim>rp>ipv6>rp-candidate group-range)

[Tree] (config>service>vprn>pim>rp>ipv6>embedded-rp group-range)

Full Context

configure service vprn pim rp ipv6 rp-candidate group-range

configure service vprn pim rp ipv6 embedded-rp group-range

Description

This command configures the group address or range of group addresses for which this router can be the rendezvous point (RP).

The **no** form of this command removes the group address or range of group addresses for which this router can be the RP from the configuration.

Parameters

ipv6-address

Specifies the addresses or address ranges that this router can be an RP.

prefix-length

Specifies the address prefix length.

Values

ipv6-address	: x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x [0 to FFFF]H
	d [0 to 255]D
prefix-length	[8 to 128] // for embedded-rp
prefix-length	[16 to 128] // for rp-candidate

Platforms

7705 SAR Gen 2

group-range

Syntax

[**no**] **group-range** {*ip-prefixImask* | *ip-prefix netmask*}

Context

[Tree] (config>service>vprn>pim>ssm group-range)

[Tree] (config>service>vprn>pim>rp>rp-candidate group-range)

Full Context

configure service vprn pim ssm-groups group-range

configure service vprn pim rp rp-candidate group-range

Description

This command configures the group address or range of group addresses for which this router can be the rendezvous point (RP).

Use the **no** form of this command to remove the group address or range of group addresses for which this router can be the RP from the configuration.

Parameters

ip-prefix

Specifies the addresses or address ranges that this router can be an RP.

Values ipv4-prefix - a.b.c.d ipv4-prefix-le - [0 to 32] ipv6-prefix - x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x - [0 to FFFF]H d - [0 to 255]D ipv6-prefix-le - [0 to 128]

mask

Specifies the address mask with the address to define a range of addresses.

netmask

Specifies the subnet mask in dotted decimal notation.

Values :a.b.c.d (network bits all 1 and host bits all 0)

Platforms

7705 SAR Gen 2

group-range

Syntax

[**no**] **group-range** *ipv6-address/prefix-length*

Context

[Tree] (config>router>pim>rp>ipv6>embedded-rp group-range)

[Tree] (config>router>pim>rp>ipv6>rp-candidate group-range)

Full Context

configure router pim rp ipv6 embedded-rp group-range

configure router pim rp ipv6 rp-candidate group-range

Description

This command defines which multicast groups can embed RP address information besides FF70::/12. Embedded RP information is only used when the multicast group is in FF70::/12 or the configured group range.

The **no** form of this command removes the parameter from the

Parameters

ipv6-address/prefix-length

Specifies the group range for embedded RP.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

prefix-length: 16 to 128

Platforms

7705 SAR Gen 2

group-range

Syntax

[**no**] **group-range** {*grp-ip-address/mask* | *grp-ip-address netmask*}

Context

[Tree] (config>router>pim>rp>rp-candidate group-range)

Full Context

configure router pim rp rp-candidate group-range

Description

This command configures the address ranges of the multicast groups for which this router can be an RP.

The **no** form of this commands removes the parameter from the configuration.

Parameters

grp-ip-address

Specifies the multicast group IP address expressed in dotted decimal notation.

Values 224.0.0.0 to 239.255.255.255

mask

Specifies the mask associated with the IP prefix expressed as a mask length or in dotted decimal notation; for example, /16 for a sixteen-bit mask. The mask can also be entered in dotted decimal notation (255.255.0.0).

Values 4 to 32

netmask

Specifies the subnet mask in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

Platforms

7705 SAR Gen 2

group-range

Syntax

[**no**] **group-range** {*ip-prefix/mask* | *ip-prefix netmask*}

Context

[Tree] (config>router>pim>ssm-groups group-range)

Full Context

configure router pim ssm-groups group-range

Description

This command configures the address ranges of the multicast groups for this router. When there are parameters present, the command configures the SSM group ranges for IPv6 addresses and netmasks.

The **no** form of this command removes the parameter from the configuration.

Parameters

ip-prefix/mask

Specifies the IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area ipv6-prefix.

Values ipv4-prefix:

- a.b.c.d

ipv4-prefix-le: 0 to 32

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

ipv6-prefix-le: 0 to 128

Values 0 to 32 (mask length), 0.0.0.0 to 255.255.255.255 (dotted decimal)

netmask

Specifies the subnet mask in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

Platforms

7705 SAR Gen 2

11.56 grp-if-query-src-ip

grp-if-query-src-ip

Syntax

grp-if-query-src-ip *ip-address*

no grp-if-query-src-ip

Context

[Tree] (config>service>vprn>igmp grp-if-query-src-ip)

Full Context

configure service vprn igmp grp-if-query-src-ip

Description

This command configures the query source IP address for all group interfaces.

The **no** form of this command removes the IP address.

Platforms

7705 SAR Gen 2

grp-if-query-src-ip

Syntax

grp-if-query-src-ip *ip-address*

no grp-if-query-src-ip

Context

[Tree] (config>router>igmp grp-if-query-src-ip)

Full Context

configure router igmp grp-if-query-src-ip

Description

This command configures the query source IP address for all group interfaces.

The **no** form of the command removes the IP address.

Parameters

ip-address

Sets the query source IP address.

Platforms

7705 SAR Gen 2

grp-if-query-src-ip

Syntax

grp-if-query-src-ip *ipv6-address*

no grp-if-query-src-ip

Context

[Tree] (config>router>mld grp-if-query-src-ip)

Full Context

configure router mld grp-if-query-src-ip

Description

This command configures the query source IPv6 address for all group interfaces.

The **no** form of this command removes the IP address.

Parameters

ipv6-address

Sets the source IPv6 address for all group interfaces. The address can be up to 64 characters. The source address should be link local.

Platforms

7705 SAR Gen 2

11.57 grp-range

grp-range

Syntax

[**no**] **grp-range** *start end*

Context

[Tree] (config>service>vprn>igmp>ssm-translate grp-range)

Full Context

configure service vprn igmp ssm-translate grp-range

Description

This command is used to configure group ranges which are translated to SSM (S,G) entries.

Parameters

start

An IP address that specifies the start of the group range.

end

An IP address that specifies the end of the group range. This value should always be greater than or equal to the value of the *start* value.

Platforms

7705 SAR Gen 2

grp-range

Syntax

[**no**] **grp-range** *start end*

Context

[Tree] (config>service>vprn>mld>ssm-translate grp-range)

Full Context

configure service vprn mld ssm-translate grp-range

Description

This command is used to configure group ranges which are translated to SSM (S,G) entries.

Parameters

start

An IP address that specifies the start of the group range.

end

An IP address that specifies the end of the group range. This value should always be greater than or equal to the value of the *start* value.

Platforms

7705 SAR Gen 2

grp-range

Syntax

[**no**] **grp-range** *start end*

Context

[Tree] (config>router>igmp>ssm-translate grp-range)

[Tree] (config>router>igmp>if>ssm-translate grp-range)

Full Context

configure router igmp ssm-translate grp-range

configure router igmp interface ssm-translate grp-range

Description

This command is used to configure group ranges which are translated to SSM (S,G) entries.

Parameters

start

An IP address that specifies the start of the group range.

end

An IP address that specifies the end of the group range. This value should always be greater than or equal to the value of the *start* value.

Platforms

7705 SAR Gen 2

grp-range

Syntax

[**no**] **grp-range** *start end*

Context

[Tree] (config>router>mld>if>ssm-translate grp-range)

[Tree] (config>router>mld>ssm-translate grp-range)

Full Context

configure router mld interface ssm-translate grp-range

configure router mld ssm-translate grp-range

Description

This command is used to configure group ranges which are translated to SSM (S,G) entries.

The **no** form of this command removes the start and end ranges from the configuration.

Parameters

start

Specifies an IP address for the start of the group range.

end

Specifies an IP address for the end of the group range. This value should always be greater than or equal to the value of the *start* value.

Platforms

7705 SAR Gen 2

11.58 grpc

grpc

Syntax

[**no**] **grpc**

Context

[Tree] (debug>system grpc)

Full Context
debug system grpc

Description
This command enables the debug context for gRPC.

The **no** form of this command removes any debug activation within the gRPC context.

Platforms
7705 SAR Gen 2

grpc

Syntax
grpc

Context
[Tree] (config>system>security>management-interface grpc)

Full Context
configure system security management-interface grpc

Description
Commands in this context configure hash-control for the gRPC interface.

Platforms
7705 SAR Gen 2

grpc

Syntax
grpc

Context
[Tree] (config>system>security>profile grpc)

Full Context
configure system security profile grpc

Description
Commands in this context configure a specific gRPC security profile.

Platforms
7705 SAR Gen 2

grpc

Syntax
grpc

Context
[Tree] (admin>system>telemetry grpc)
[Tree] (config>system grpc)

Full Context
admin system telemetry grpc
configure system grpc

Description
Commands in this context configure gRPC parameters.

Platforms
7705 SAR Gen 2

11.59 grpc-tunnel

grpc-tunnel

Syntax
grpc-tunnel

Context
[Tree] (config>system grpc-tunnel)

Full Context
configure system grpc-tunnel

Description
Commands in this context configure the GRPC tunnel.

Platforms
7705 SAR Gen 2

11.60 grt

grt

Syntax

[**no**] **grt**

Context

[Tree] (config>service>vprn>static-route-entry grt)

Full Context

configure service vprn static-route-entry grt

Description

This command creates a static route in a VPRN service context that points to the global routing context (base router). This is primarily used to allow traffic that ingress through a VPRN service to be routed out of the global routing context.

This next-hop type cannot be used in conjunction with any other next-hop types.

Default

no grt

Platforms

7705 SAR Gen 2

11.61 grt-lookup

grt-lookup

Syntax

grt-lookup

Context

[Tree] (config>service>vprn grt-lookup)

Full Context

configure service vprn grt-lookup

Description

Commands in this context configure all Global Route Table (GRT) leaking commands. If all the supporting commands in the context are removed, this command is also removed.

Platforms

7705 SAR Gen 2

12 h Commands

12.1 half-life

half-life

Syntax

half-life *half-life* **max-suppress-time** *max-time*

Context

[Tree] (config>port>ethernet>dampening half-life)

Full Context

configure port ethernet dampening half-life

Description

This command configures the half-life decay time and the maximum period of time for which the port up state can be suppressed.

The *half-life* and *max-time* values must be set at the same time; the ratio of *max-time*/ *half-life* must be less than or equal to 49 and greater than or equal to 1.

Parameters

half-life

Specifies the required elapsed time, in seconds, before penalties decay to one-half the initial amount.

Values 1 to 2000

Default 5

max-time

Specifies the maximum suppression time, in seconds, which is the time it can take after the physical link comes up before the worst case accumulated penalties have decayed to the reuse threshold. The maximum penalty is derived from the maximum suppression time, half-life, and reuse threshold, using the following equation:

maximum penalty = (reuse threshold) X 2 expo:(max-time/half-life)

Values 1 to 43200

Default 20

Platforms

7705 SAR Gen 2

half-life

Syntax

half-life *minutes*

no half-life

Context

[Tree] (config>router>policy-options>damping half-life)

Full Context

configure router policy-options damping half-life

Description

This command configures the **half-life** parameter for the route damping profile.

The half-life value is the time, expressed in minutes, required for a route to remain stable in order for the Figure of Merit (FoM) value to be reduced by one half; for example, if the half-life value is 6 (minutes) and the route remains stable for 6 minutes, then the new FoM value is 3 (minutes). After another 3 minutes pass and the route remains stable, the new FoM value is 1.5 (minutes).

When the FoM value falls below the **config>router>policy-options>damping reuse** threshold, the route is once again considered valid and can be reused or included in route advertisements.

The **no** form of this command removes the half life parameter from the damping profile.

Default

no half-life

Parameters

minutes

Specifies the half-life in minutes expressed as a decimal integer.

Values 1 to 45

Platforms

7705 SAR Gen 2

12.2 handler

handler

Syntax

[**no**] **handler** *event-handler-name*

Context

[Tree] (config>log>event-handling handler)

Full Context

configure log event-handling handler

Description

This command configures an EHS handler.

The **no** form of this command removes the specified EHS handler.

Parameters

event-handler-name

Specifies the name of the EHS handler, up to 32 characters maximum.

Platforms

7705 SAR Gen 2

handler

Syntax

handler *name* **[create]**

no handler *name*

Context

[Tree] (config>system>grpc-tunnel>tunnel handler)

Full Context

configure system grpc-tunnel tunnel handler

Description

Commands in this context configure tunnel handler parameters. There can be multiple handlers created for any tunnel.

The **no** form of this command removes the specified tunnel handler.

Parameters

name

Specifies the handler name, up to 32 characters.

create

Keyword used to create a tunnel.

Platforms

7705 SAR Gen 2

12.3 hash-algorithm

hash-algorithm

Syntax

hash-algorithm {**hash** | **hash2** | **custom**| **cleartext**}

no hash-algorithm

Context

[Tree] (config>system>security>management-interface>md-cli hash-algorithm)

[Tree] (config>system>security>management-interface>grpc hash-algorithm)

[Tree] (config>system>security>management-interface>netconf hash-algorithm)

Full Context

configure system security management-interface md-cli hash-algorithm

configure system security management-interface grpc hash-algorithm

configure system security management-interface netconf hash-algorithm

Description

This command specifies the format of the input and output for encrypted configuration secrets.

The **no** form of this command reverts to the default value.

Default

hash-algorithm hash2

Parameters

hash

Specifies hash. Use this option to transport a phrase between modules and nodes.

hash2

Specifies hash2 which is module-specific.

custom

Specifies the custom encryption to management interface.

cleartext

Specifies that the phrase is displayed as cleartext everywhere.

Platforms

7705 SAR Gen 2

hash-algorithm

Syntax

hash-algorithm *algorithm*

Context

[Tree] (config>system>security>pki>cert-upd-prof hash-algorithm)

Full Context

configure system security pki certificate-update-profile hash-algorithm

Description

This command configures the hash algorithm used to generate a certificate request.

Default

hash-algorithm sha256

Parameters

algorithm

Specifies the hash option.

Values md5, sha1, sha224, sha256, sha384, sha512

Platforms

7705 SAR Gen 2

12.4 hash-label

hash-label

Syntax

hash-label

hash-label [**signal-capability**]

no hash-label

Context

[Tree] (config>service>pw-template hash-label)

[Tree] (config>service>epipe>spoke-sdp hash-label)

Full Context

configure service pw-template hash-label

configure service epipe spoke-sdp hash-label

Description

This command enables the use of the hash label on a VLL, VPRN or VPLS service bound to any MPLS type encapsulated SDP, as well as to a VPRN service that is using the **auto-bind-tunnel** with the **resolution-filter** set to any MPLS tunnel type. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the gre option. This feature is also not supported on multicast packets forwarded using RSVP P2MP LSP or mLDP LSP in both the base router instance and in the multicast VPN (mVPN) instance. It is, however, supported when forwarding multicast packets using an IES/VPRN spoke-interface.

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to one (1).

To allow applications where the egress LER infers the presence of the hash label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the Hash Label. This means that the value of the hash label will always be in the range [524,288 - 1,048,575] and will not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. Note, however, that for VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

Packets generated in CPM and that are forwarded labeled within the context of a service (for example, OAM packets) must also include a Hash Label at the BoS and set the S-bit accordingly.

The TTL of the hash label is set to a value of 0.

The user enables the signaling of the hash-label capability under a VLL spoke-sdp, a VPLS spoke-sdp or mesh SDP, or an IES/VPRN spoke interface by adding the **signal-capability** option. In this case, the decision whether to insert the hash label on the user and control plane packets by the local PE is solely determined by the outcome of the signaling process and can override the local PE configuration. The following are the procedures:

- The local PE will insert the flow label interface parameters sub-TLV with F=1 in the PW ID FEC element in the label mapping message for that spoke SDP or mesh SDP.
- If the remote PE includes this sub-TLV with F=1 or F=0, then local PE must insert the hash label in the user and control plane packets.

- If remote PE does not include this sub-TLV (for example, it does not support it, or it is supported but the user did not enable the **hash-label** option or the **signal-capability** option), then the local PE establishes the PW but must not insert the hash label in the user and control packets over that spoke SDP or mesh SDP. If the remote PE does not support the **signal-capability** option, then there are a couple of possible outcomes:
  - If the **hash-label** option was enabled on the local configuration of the spoke SDP or mesh SDP at the remote PE, the PW packets received by the local PE will have the hash label included. These packets must be dropped. The only way to solve this is to disable the signaling capability option on the local node which will result in the insertion of the hash label by both PE nodes.
  - If the **hash-label** option is not supported or was not enabled on the local configuration of the spoke SDP or mesh SDP at the remote PE, the PW received by the local PE will not have the hash label included.
- The user can enable or disable the signal-capability option in CLI as needed. When doing so, the 7705 SAR Gen 2 must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the PW ID FEC element.

The **no** form of this command disables the use of the hash label.

Default

no hash-label

Parameters

signal-capability

Enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes. The **signal-capability** option is not supported on a VPRN spoke-sdp.

Platforms

7705 SAR Gen 2

hash-label

Syntax

hash-label signal-capability

hash-label

no hash-label

Context

[Tree] (config>service>vpls>mesh-sdp hash-label)

[Tree] (config>service>vpls>spoke-sdp hash-label)

Full Context

configure service vpls mesh-sdp hash-label

configure service vpls spoke-sdp hash-label

Description

This command enables the use of the hash label on a VLL, VPRN, or VPLS service bound to any MPLS type encapsulated SDP, as well as to a VPRN service using the **auto-bind-tunnel** with the **resolution-filter** set to any MPLS tunnel type. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the gre option. This feature is also not supported on multicast packets forwarded using RSVP P2MP LSP or mLDP LSP in both the base router instance and in the multicast VPN (mVPN) instance. It is, however, supported when forwarding multicast packets using an IES/VPRN spoke-interface.

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to one (1).

To allow applications where the egress LER infers the presence of the hash label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the Hash Label. This means that the value of the hash label will always be in the range [524,288 - 1,048,575] and will not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. Note, however, that for VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

Packets generated in CPM and that are forwarded labeled within the context of a service (for example, OAM packets) must also include a Hash Label at the BoS and set the S-bit accordingly.

The TTL of the hash label is set to a value of 0.

The user enables the signaling of the hash-label capability under a VLL spoke-sdp, a VPLS spoke-sdp or mesh-sdp, or an IES/VPRN spoke interface by adding the **signal-capability** option. In this case, the decision whether to insert the hash label on the user and control plane packets by the local PE is solely determined by the outcome of the signaling process and can override the local PE configuration. The following are the procedures:

- The 7705 SAR Gen 2 local PE will insert the flow label interface parameters sub-TLV with F=1 in the pseudowire ID FEC element in the label mapping message for that spoke-sdp or mesh-sdp.
- If the remote PE includes this sub-TLV with F=1 or F=0, then local PE must insert the hash label in the user and control plane packets.
- If remote PE does not include this sub-TLV (for example, it does not support it, or it is supported but the user did not enable the **hash-label** option or the **signal-capability** option), then the local PE establishes the pseudowire but must not insert the hash label in the user and control packets over that spoke-sdp or mesh-sdp. If the remote PE does not support the **signal-capability** option, then there are a couple of possible outcomes:
  - If the **hash-label** option was enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the pseudowire packets received by the local PE will have the hash label included. These packets must be dropped. The only way to solve this is to disable the signaling capability option on the local node which will result in the insertion of the hash label by both PE nodes.
  - If the **hash-label** option is not supported or was not enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the pseudowire received by the local PE will not have the hash label included.

- The user can enable or disable the signal-capability option in CLI as needed. When doing so, the 7705 SAR Gen 2 must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the pseudowire ID FEC element.

The **no** form of this command disables the use of the hash label.

Default

no hash-label

Parameters

signal-capability

Enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes. The **signal-capability** option is not supported on a VPRN spoke-sdp.

Platforms

7705 SAR Gen 2

hash-label

Syntax

hash-label [**signal-capability**]

no hash-label

Context

[Tree] (config>service>ies>if>spoke-sdp hash-label)

Full Context

configure service ies interface spoke-sdp hash-label

Description

This command enables the use of the hash label on a VLL, VPLS, or VPRN service bound to any MPLS-type encapsulated SDP, as well as to a VPRN service using **auto-bind-tunnel** with the **resolution-filter** configures as any MPLS tunnel type. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the gre option.

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to 1 to indicate that.

In order to allow for applications whereby the egress LER infers the presence of the hash label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the hash label. This means that the value of the hash label will always be in the range [524,288 to 1,048,575] and will not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. For VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

Packets that are generated in CPM and forwarded labeled within the context of a service (for example, OAM packets) must also include a hash label at the BoS and set the S-bit accordingly.

The TTL of the hash label is set to a value of 0.

The user enables the signaling of the hash-label capability under a VLL spoke-sdp, a VPLS spoke-sdp or mesh-sdp, or an IES/VPRN spoke interface by adding the **signal-capability** option. In this case, the decision whether to insert the hash label on the user and control plane packets by the local PE is solely determined by the outcome of the signaling process and can override the local PE configuration. The following are the procedures:

- The local PE will insert the flow label interface parameters sub-TLV with F=1 in the PW ID FEC element in the label mapping message for that spoke-sdp or mesh-sdp.
- If the remote PE includes this sub-TLV with F=1 or F=0, then local PE must insert the hash label in the user and control plane packets.
- If remote PE does not include this sub-TLV (for example, it does not support it, or it is supported but the user did not enable the **hash-label** option or the **signal-capability** option), then the local PE establishes the PW but must not insert the hash label in the user and control packets over that spoke-sdp or mesh-sdp. If the remote PE does not support the **signal-capability** option, then there are a couple of possible outcomes:
  - If the **hash-label** option was enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the PW packets received by the local PE will have the hash label included. These packets must be dropped. The only way to solve this is to disable the signaling capability option on the local node which will result in the insertion of the hash label by both PE nodes.
  - If the **hash-label** option is not supported or was not enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the PW received by the local PE will not have the hash label included.
- The user can enable or disable the signal-capability option in CLI as needed. When doing so, the router must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the PW ID FEC element.

The **no** form of this command disables the use of the hash label.

Default

no hash-label

Parameters

signal-capability

Enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes. The **signal-capability** option is not supported on a VPRN spoke-sdp.

Platforms

7705 SAR Gen 2

hash-label

Syntax

hash-label

hash-label signal-capability

no hash-label

Context

[Tree] (config>service>vprn>if>spoke-sdp hash-label)

[Tree] (config>service>vprn hash-label)

Full Context

configure service vprn interface spoke-sdp hash-label

configure service vprn hash-label

Description

This command enables the use of the hash label on a VLL, VPLS, or VPRN service bound to any MPLS-type encapsulated SDP as well as to a VPRN service using **auto-bind-tunnel** with the **resolution-filter** configured as any MPLS tunnel type. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the gre option.

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to 1 to indicate that.

In order to allow for applications whereby the egress LER infers the presence of the Hash Label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the Hash Label. This means that the value of the hash label will always be in the range [524,288 - 1,048,575] and will not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. For VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

Packets that are generated in CPM and forwarded labeled within the context of a service (for example, OAM packets) must also include a Hash Label at the BoS and set the S-bit accordingly.

The TTL of the hash label is set to a value of 0.

The **no** form of this command disables the use of the hash label.

Default

no hash-label

Parameters

signal-capability

Specifies whether the service should send the Stack Capability and check whether the capability is received from the peer via LDP interface parameters.

Platforms

7705 SAR Gen 2

hash-label

Syntax

[**no**] **hash-label**

Context

[Tree] (config>service>vpls>bgp-evpn>mpls hash-label)

Full Context

configure service vpls bgp-evpn mpls hash-label

Description

This command pushes the hash label based on the following:

- If the **no incl-mcast-l2-attributes-advertisement** command is configured, the hash label is pushed to a unicast EVPN destination.
- If the **incl-mcast-l2-attributes-advertisement** command is configured, the F bit is set to 1 in the Layer 2 Attributes Extended Community of the EVPN IMET route for the service. The hash label is pushed only if the remote PE signaled support for hash label (received F bit is set to 1).

The hash label is never used for BUM packets.

The **no** form of this command disables the push of the hash label.

Default

no hash-label

Platforms

7705 SAR Gen 2

12.5 hash-mask-len

hash-mask-len

Syntax

hash-mask-len *hash-mask-length*

no hash-mask-len

Context

[Tree] (config>service>vprn>pim>rp>bsr-candidate hash-mask-len)

Full Context

configure service vprn pim rp bsr-candidate hash-mask-len

Description

This command is used to configure the length of a mask that is to be combined with the group address before the hash function is called. All groups with the same hash map to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This mechanism is used to map one group or multiple groups to an RP.

Default

hash-mask-len 30

Parameters

hash-mask-length

The hash mask length.

Values 0 to 32

Platforms

7705 SAR Gen 2

hash-mask-len

Syntax

hash-mask-len *hash-mask-length*

no hash-mask-len

Context

[Tree] (config>service>vprn>pim>rp>ipv6>bsr-candidate hash-mask-len)

Full Context

configure service vprn pim rp ipv6 bsr-candidate hash-mask-len

Description

This command is used to configure the length of a mask that is to be combined with the group address before the hash function is called. All groups with the same hash map to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This mechanism is used to map one group or multiple groups to an RP.

Default

hash-mask-len 126

Parameters

hash-mask-length

The hash mask length.

Values 0 to 128

Platforms

7705 SAR Gen 2

hash-mask-len

Syntax

hash-mask-len *hash-mask-length*

no hash-mask-len

Context

[Tree] (config>router>pim>rp>bsr-candidate hash-mask-len)

[Tree] (config>router>pim>rp>ipv6>bsr-candidate hash-mask-len)

Full Context

configure router pim rp bsr-candidate hash-mask-len

configure router pim rp ipv6 bsr-candidate hash-mask-len

Description

This command configures the length of a mask that is to be combined with the group address before the hash function is called. All groups with the same hash map to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This mechanism is used to map one group or multiple groups to an RP.

The **no** form of this command reverts to the default value.

Default

hash-mask-len 30 — for **config**>**router**>**pim**>**rp**>**bsr-candidate**

hash-mask-len 126 — for **config**>**router**>**pim**>**rp**>**ipv6**> **bsr-candidate**

Parameters

hash-mask-length

Specifies the hash mask length.

Values 0 to 32 (v4)
0 to 128 (v6)

Platforms

7705 SAR Gen 2

12.6 hash-weight-threshold

hash-weight-threshold

Syntax

hash-weight-threshold *weight* [**action** *action*] [**cost** *static-cost*]

no hash-weight-threshold

Context

[Tree] (config>lag hash-weight-threshold)

Full Context

configure lag hash-weight-threshold

Description

This command controls the operational status of the LAG or the IGP cost based on the sum of the **hash-weight** values for the active links in the LAG.

The **no** form of this command disables the hash weight threshold.

Parameters

weight

Specifies the value for the sum of all the active LAG ports **hash-weight** at or below which the configured action is invoked. If the sum of **hash-weight** for operational LAG links exceeds the **hash-weight-threshold** value, then no action is taken.

Values 1 to 6400000

action

Specifies the action to take if the sum of the **hash-weight** for active links in the LAG is equal or below the threshold value.

Values **down —** Specifies that the LAG is operationally DOWN. The LAG is only considered as UP once the number of **hash-weight** for the active links exceeds the configured threshold value.

dynamic-cost — Specifies that dynamic cost is activated. The LAG remains operationally UP with a link cost relative to the number of operational links. The link is only considered as operationally DOWN when all links in the LAG are down.

static-cost — Specifies that static cost is activated. The LAG remains operationally UP with the configured cost, regardless of the number of operational links. The link is only considered as operationally DOWN when all links in the LAG are down. If this parameter is used with an IGP, its **reference-bandwidth** must also be configured.

static-cost

Specifies the decimal integer static cost of the LAG.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

12.7 hashing

hashing

Syntax

hashing {**bcrypt** | **sha2-pbkdf2**| **sha3-pbkdf2**}

Context

[Tree] (config>system>security>password hashing)

Full Context

configure system security password hashing

Description

This command configures the password hashing algorithm.

Default

hashing bcrypt

Parameters

bcrypt

Keyword to indicate that the command configures the bcrypt algorithm.

sha2-pbkdf2

Keyword to indicate that the command configures the PBKDF2 algorithm hashed via SHA2.

sha3-pbkdf2

Keyword to indicate that the command configures the PBKDF2 algorithm hashed via SHA3.

Platforms

7705 SAR Gen 2

12.8 head-end

head-end

Syntax

head-end local

head-end *ipv4-address*

no head-end

Context

[Tree] (conf>router>segment-routing>sr-policies>policy head-end)

Full Context

configure router segment-routing sr-policies static-policy head-end

Description

This command associates a head-end location with a statically-defined segment-routing policy. The head-end identifies the router that is the target to install the policy. This is a mandatory parameter and configuration command for enabling the segment-routing policy; if the head-end parameter value is not configured, the execution of the **no shutdown** command on the static segment routing policy fails.

To associate a static policy with the local router as head-end, the keyword **local** must be specified. The static policy is associated with another (non-local) router, if the head-end parameter is set to any IPv4 address. When a non-local, static segment routing policy that originates as a BGP route is imported into BGP, the configured head-end address is converted to an IPv4-address specific route-target extended community that is automatically added to the route.

The **no** form of this command removes the head-end association.

Default

no head-end

Parameters

local

Keyword indicating that the policy is intended to be used by the local router and not advertised to other BGP routers.

ipv4-address

Specifies the IP address of the target head-end router.

Values ipv4-address: a.b.c.d

Platforms

7705 SAR Gen 2

12.9 health-check

health-check

Syntax

health-check

Context

[Tree] (config>aaa>radius-server-policy>servers health-check)

Full Context

configure aaa radius-server-policy servers health-check

Description

Commands in this context configure health check parameters for the RADIUS server.

Platforms

7705 SAR Gen 2

health-check

Syntax

[**no**] **health-check** [**interval** *interval*]

Context

[Tree] (config>system>security>password health-check)

Full Context

configure system security password health-check

Description

This command enables health check monitoring of the RADIUS, TACACS+, and LDAP servers by sending authentication requests for an unknown user at regular intervals. If a response is not received, the operational status of the server is changed to down. The operational status is changed to up when responses are received.

When RADIUS over TLS is configured, Status-Server packets are sent at 30-second intervals as specified in *RFC 3539*, regardless of whether health checks are enabled.

The **no** form of this command disables health monitoring of RADIUS, TACACS+, and LDAP servers. In this case, the operational status for the server is up if a response was received for the last user request.

Default

health-check interval 30

Parameters

interval

Specifies the polling interval for RADIUS, TACACS+, and LDAP servers.

Values 6 to 1500

Default 30

Platforms

7705 SAR Gen 2

12.10 hello

hello

Syntax

hello *timeout factor*

no hello

Context

[Tree] (config>router>ldp>if-params>ipv4 hello)

[Tree] (config>router>ldp>targ-session>ipv6 hello)

[Tree] (config>router>ldp>if-params>if>ipv6 hello)

[Tree] (config>router>ldp>if-params>ipv6 hello)

[Tree] (config>router>ldp>targ-session>ipv4 hello)

[Tree] (config>router>ldp>if-params>if>ipv4 hello)

[Tree] (config>router>ldp>targ-session>peer hello)

[Tree] (config>router>ldp>targ-session>peer-template hello)

Full Context

configure router ldp interface-parameters ipv4 hello

configure router ldp targeted-session ipv6 hello

configure router ldp interface-parameters interface ipv6 hello

configure router ldp interface-parameters ipv6 hello

configure router ldp targeted-session ipv4 hello

configure router ldp interface-parameters interface ipv4 hello

configure router ldp targeted-session peer hello

configure router ldp targeted-session peer-template hello

Description

This command configures the time interval to wait before declaring a neighbor down. The **factor** parameter derives the Hello interval.

Hold time is local to the system and sent in the Hello messages to the neighbor. Hold time cannot be less than three times the Hello interval. The hold time can be configured globally (applies to all LDP interfaces) or per interface. The most specific value is used.

When LDP session is being set up, the hold down time is negotiated to the lower of the two peers. Once an operational value is agreed upon, the Hello factor is used to derive the value of the Hello interval.

The **no** form of the command at the interface-parameters and targeted-session level sets the **hello timeout** and the **hello factor** to the default values.

The **no** form of the command, at the interface level, sets the **hello timeout** and the **hello factor** to the value defined under the interface-parameters level.

The **no** form of this command, at the peer level, sets the **hello timeout** and the **hello factor** to the value defined under the targeted-session level.

The session must be flapped for the new settings to operate.

Default

Table 41: Hello Timeout Factors lists the default values.

Table 41: Hello Timeout Factors

Context	Timeout	Factor
config>router>ldp>if-params	15	3

Context	Timeout	Factor
config>router>ldp>targ-session	45	3
config>router>ldp>if-params>if	Inherits values from interface-parameters context.	
config>router>ldp>targ-session>peer	Inherits values from targeted-session context.	

Parameters

timeout

Configures the time interval, in seconds, that LDP waits before a neighbor down.

Values 1 to 65535

factor

Specifies the number of keepalive messages that should be sent on an idle LDP session in the Hello timeout interval.

Values 1 to 255

Platforms

7705 SAR Gen 2

hello

Syntax

hello [**detail**]

no hello

Context

[Tree] (debug>router>ldp>if>packet hello)

[Tree] (debug>router>ldp>peer>packet hello)

Full Context

debug router ldp interface packet hello

debug router ldp peer packet hello

Description

This command enables debugging for LDP Hello packets.

The **no** form of the command disables the debugging output.

Parameters

detail

Displays detailed information.

Platforms

7705 SAR Gen 2

hello

Syntax

hello [**detail**]

no hello

Context

[Tree] (debug>router>rsvp>packet hello)

Full Context

debug router rsvp packet hello

Description

This command debugs Hello packets.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about Hello packets.

Platforms

7705 SAR Gen 2

12.11 hello-auth-keychain

hello-auth-keychain

Syntax

hello-auth-keychain *name*

Context

[Tree] (config>service>vprn>isis>interface hello-auth-keychain)

[Tree] (config>service>vprn>isis>interface>level hello-auth-keychain)

Full Context

configure service vprn isis interface hello-auth-keychain

configure service vprn isis interface level hello-auth-keychain

Description

This command configures an authentication keychain to use for the protocol interface. The keychain allows the rollover of authentication keys during the lifetime of a session.

Default

no hello-auth-keychain

Parameters

name

Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

Platforms

7705 SAR Gen 2

12.12 hello-authentication

hello-authentication

Syntax

[**no**] **hello-authentication**

Context

[Tree] (config>service>vprn>isis>level hello-authentication)

[Tree] (config>service>vprn>isis>if hello-authentication)

[Tree] (config>service>vprn>isis hello-authentication)

Full Context

configure service vprn isis level hello-authentication

configure service vprn isis interface hello-authentication

configure service vprn isis hello-authentication

Description

This command enables authentication of individual IS-IS Hello packets for the VPRN instance.

The **no** form of this command suppresses authentication of Hello packets.

Platforms

7705 SAR Gen 2

hello-authentication

Syntax

[**no**] **hello-authentication**

Context

[Tree] (config>router>isis hello-authentication)

[Tree] (config>router>isis>level hello-authentication)

[Tree] (config>router>isis>interface hello-authentication)

Full Context

configure router isis hello-authentication

configure router isis level hello-authentication

configure router isis interface hello-authentication

Description

This command enables authentication of individual IS-IS packets of HELLO type.

The **no** form of this command suppresses authentication of HELLO packets.

Default

hello-authentication

Platforms

7705 SAR Gen 2

12.13 hello-authentication-key

hello-authentication-key

Syntax

hello-authentication-key {*authentication-key* | *hash-key*} [**hash** | **hash2** | **custom**]

no hello-authentication-key

Context

[Tree] (config>service>vprn>isis>if hello-authentication-key)

[Tree] (config>service>vprn>isis>if>level hello-authentication-key)

Full Context

configure service vprn isis interface hello-authentication-key

configure service vprn isis interface level hello-authentication-key

Description

This command configures the authentication key (password) for Hello PDUs. Neighboring routers use the password to verify the authenticity of Hello PDUs sent from this interface. Both the Hello authentication key and the Hello authentication type on a segment must match. The **hello-authentication-type** must be specified.

To configure the Hello authentication key in the interface context use the **hello-authentication-key** in the **config>router>isis>if** context.

To configure or override the Hello authentication key for a specific level, configure the **hello-authentication-key** in the **config>router>isis>if>level** context.

If both IS-IS and hello-authentication are configured, Hello messages are validated using Hello authentication. If only IS-IS authentication is configured, it will be used to authenticate all IS-IS (including Hello) protocol PDUs.

When the Hello authentication key is configured in the **config>router>isis>if** context, it applies to all levels configured for the interface.

The **no** form of this command removes the authentication-key from the configuration.

Default

no hello-authentication-key — No Hello authentication key is configured.

Parameters

authentication-key

The Hello authentication key (password). The key can be any combination of ASCII characters up to 254 characters in length (un-encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be

in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

hello-authentication-key

Syntax

hello-authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2** | **custom**]

no hello-authentication-key

Context

[Tree] (config>router>isis>interface hello-authentication-key)

[Tree] (config>router>isis>if>level hello-authentication-key)

Full Context

configure router isis interface hello-authentication-key

configure router isis interface level hello-authentication-key

Description

This command configures the authentication key (password) for Hello PDUs. Neighboring routers use the password to verify the authenticity of Hello PDUs sent from this interface. Both the Hello authentication key and the Hello authentication type on a segment must match. The **hello-authentication-type** must be specified.

To configure the Hello authentication key in the interface context, use the **hello-authentication-key** in the **config>router>isis>interface** context.

To configure or override the Hello authentication key for a specific level, configure the **hello-authentication-key** in the **config>router>isis>interface>level** context.

If both IS-IS and hello-authentication are configured, Hello messages are validated using Hello authentication. If only IS-IS authentication is configured, it will be used to authenticate all IS-IS (including Hello) protocol PDUs.

When the Hello authentication key is configured in the **config>router>isis>interface** context, it applies to all levels configured for the interface.

The **no** form of this command removes the authentication-key from the configuration.

Parameters

authentication-key

Specifies the Hello authentication key (password). The key can be any combination of ASCII characters, up to 254 characters (un-encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ”).

hash-key

Specifies the hash key. The key can be any combination of ASCII characters, up to 342 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

12.14 hello-authentication-type

hello-authentication-type

Syntax

hello-authentication-type {**password** | **message-digest**}

no hello-authentication-type

Context

[Tree] (config>service>vprn>isis>if>level hello-authentication-type)

[Tree] (config>service>vprn>isis>if hello-authentication-type)

Full Context

configure service vprn isis interface level hello-authentication-type

configure service vprn isis interface hello-authentication-type

Description

This command enables Hello authentication at either the interface or level context. Both the Hello authentication key and the Hello authentication type on a segment must match. The hello **authentication-key** statement must also be included.

To configure the Hello authentication type at the interface context, use **hello-authentication-type** in the **config>router>isis>if** context.

To configure or override the Hello authentication setting for a given level, configure the **hello-authentication-type** in the **config>router>isis>if>level** context.

The **no** form of this command disables Hello authentication.

Default

no hello-authentication-type — Hello authentication is disabled

Parameters

password

Specifies simple password (plain text) authentication is required.

message-digest

Specifies MD5 authentication in accordance with RFC 2104 (HMAC: Keyed-Hashing for Message Authentication) is required.

Platforms

7705 SAR Gen 2

hello-authentication-type

Syntax

hello-authentication-type {**password** | **message-digest**}

no hello-authentication-type

Context

[Tree] (config>router>isis>interface hello-authentication-type)

[Tree] (config>router>isis>if>level hello-authentication-type)

Full Context

configure router isis interface hello-authentication-type

configure router isis interface level hello-authentication-type

Description

This command enables Hello authentication at either the interface or level context. Both the Hello authentication key and the Hello authentication type on a segment must match. The hello **authentication-key** statement must also be included.

To configure the Hello authentication type at the interface context, use **hello-authentication-type** in the **config>router>isis>interface** context.

To configure or override the Hello authentication setting for a given level, configure the **hello-authentication-type** in the **config>router>isis>interface>level** context.

The **no** form of this command disables Hello authentication.

Parameters

password

Specifies simple password (plain text) authentication is required.

message-digest

Specifies MD5 authentication in accordance with RFC 2104 (HMAC: Keyed-Hashing for Message Authentication) is required.

Platforms

7705 SAR Gen 2

12.15 hello-interval

hello-interval

Syntax

hello-interval *seconds*

no hello-interval

Context

[Tree] (config>service>vprn>isis>if>level hello-interval)

[Tree] (config>router>isis>if>level hello-interval)

Full Context

configure service vprn isis interface level hello-interval

configure router isis interface level hello-interval

Description

This command configures the interval between IS-IS Hello PDUs issued on the interface at this level. The **hello-interval**, along with the **hello-multiplier**, is used to calculate a hold time, which is communicated to a neighbor in a Hello PDU.

Note:
The neighbor hold time is (hello multiplier X hello interval) on non-designated intermediate system broadcast interfaces and point-to-point interfaces and is (hello multiplier X hello interval / 3) on

designated intermediate system broadcast interfaces. Hello values can be adjusted for faster convergence, but the hold time should always be > 3 to reduce routing instability.

The **no** form of this command reverts to the default value.

Default

3 – for designated intermediate system interfaces

9 – for non-designated intermediate system interfaces and point-to-point interfaces

Parameters

seconds

The Hello interval in seconds expressed as a decimal integer.

Values 1 to 20000

Platforms

7705 SAR Gen 2

hello-interval

Syntax

hello-interval *seconds*

no hello-interval

Context

[Tree] (config>service>vprn>ospf>area>sham-link hello-interval)

[Tree] (config>service>vprn>ospf3>area>virtual-link hello-interval)

[Tree] (config>service>vprn>ospf3>area>if hello-interval)

[Tree] (config>service>vprn>ospf>area>virtual-link hello-interval)

[Tree] (config>service>vprn>ospf>area>if hello-interval)

Full Context

configure service vprn ospf area sham-link hello-interval

configure service vprn ospf3 area virtual-link hello-interval

configure service vprn ospf3 area interface hello-interval

configure service vprn ospf area virtual-link hello-interval

configure service vprn ospf area interface hello-interval

Description

This command configures the interval between OSPF Hello messages issued on the interface, virtual link, or sham-link.

The Hello interval, in combination with the dead-interval, is used to establish and maintain the adjacency. Use this parameter to edit the frequency that Hello packets are sent.

Reducing the interval, in combination with an appropriate reduction in the associated **dead-interval**, allows for faster detection of link and/or router failures at the cost of higher processing costs.

The **no** form of this command reverts to the default value.

Default

hello-interval 10 — a 10-second Hello interval

Parameters

seconds

The Hello interval in seconds expressed as a decimal integer.

Values 1 to 65535

Platforms

7705 SAR Gen 2

hello-interval

Syntax

hello-interval *hello-interval*

no hello-interval

Context

[Tree] (config>service>vprn>pim>if hello-interval)

Full Context

configure service vprn pim interface hello-interval

Description

This command configures the frequency at which PIM Hello messages are transmitted on this interface.

The **no** form of this command resets the configuration to the default value.

Default

hello-interval 30

Parameters

hello-interval

Specifies the Hello interval in seconds. A 0 (zero) value disables the sending of Hello messages (the PIM neighbor will never timeout the adjacency).

Values 0 to 255 seconds

Platforms

7705 SAR Gen 2

hello-interval

Syntax

hello-interval *milli-seconds*

no hello-interval

Context

[Tree] (config>router>rsvp>interface hello-interval)

Full Context

configure router rsvp interface hello-interval

Description

This command configures the time interval between RSVP Hello messages.

RSVP Hello packets are used to detect loss of RSVP connectivity with the neighboring node. Hello packets detect the loss of neighbor far quicker than it would take for the RSVP session to time out based on the refresh interval. After the loss of the of number keep-multiplier consecutive Hello packets, the neighbor is declared to be in a down state.

The **no** form of this command reverts to the default value of the hello-interval. To disable sending hello messages, set the value to zero.

Default

hello-interval 3000

Parameters

milli-seconds

Specifies the RSVP Hello interval (in ms), in multiples of 1000. A 0 (zero) value disables the sending of RSVP Hello messages.

Values 0 to 60000 ms (in multiples of 1000)

Platforms

7705 SAR Gen 2

hello-interval

Syntax

hello-interval *hello-interval*

no hello-interval

Context

[Tree] (config>router>pim>interface hello-interval)

Full Context

configure router pim interface hello-interval

Description

This command configures the frequency at which PIM Hello messages are transmitted on this interface.

The **no** form of this command resets the configuration to the default value.

Default

hello-interval 30

Parameters

hello-interval

Specifies the Hello interval in seconds. A 0 (zero) value disables the sending of Hello messages (the PIM neighbor will never timeout the adjacency).

Values 0 to 255 seconds

Platforms

7705 SAR Gen 2

hello-interval

Syntax

hello-interval *seconds*

no hello-interval

Context

[Tree] (config>router>ospf>area>virtual-link hello-interval)

[Tree] (config>router>ospf3>area>interface hello-interval)

[Tree] (config>router>ospf>area>interface hello-interval)

[Tree] (config>router>ospf3>area>virtual-link hello-interval)

Full Context

configure router ospf area virtual-link hello-interval

configure router ospf3 area interface hello-interval

configure router ospf area interface hello-interval

configure router ospf3 area virtual-link hello-interval

Description

This command configures the interval between OSPF Hellos issued on the interface or virtual link.

The Hello interval, in combination with the dead-interval, is used to establish and maintain the adjacency. Use this parameter to edit the frequency that Hello packets are sent.

Reducing the interval, in combination with an appropriate reduction in the associated **dead-interval ,** allows for faster detection of link and/or router failures at the cost of higher processing costs.

The **no** form of this command reverts to the default value.

Default

hello-interval 10

Parameters

seconds

Specifies the Hello interval, in seconds, expressed as a decimal integer.

Values 1 to 65535

Platforms

7705 SAR Gen 2

hello-interval

Syntax

hello-interval *number*

no hello-interval

Context

[Tree] (config>system>management-interface>remote-management hello-interval)

Full Context

configure system management-interface remote-management hello-interval

Description

This command configures the time interval between Hello messages sent from the SR OS node to the remote manager.

Default

hello-interval 10

Parameters

number

Specifies the Hello interval, in minutes.

Values 10 to 3600

Platforms

7705 SAR Gen 2

12.16 hello-multiplier

hello-multiplier

Syntax

hello-multiplier *multiplier*

no hello-multiplier

Context

[Tree] (config>service>vprn>isis>if>level hello-multiplier)

Full Context

configure service vprn isis interface level hello-multiplier

Description

This command configures the number of missing Hello messages from a neighbor before the router declares the adjacency down.

Note:
The neighbor hold time is (hello multiplier X hello interval) on point-to-point interfaces, and (hello multiplier X hello interval / 3) on broadcast interfaces. Hello values can be adjusted for faster convergence, but the hold-time should always be > 3 to reduce routing instability.

The **no** form of this command reverts to the default value.

Default

hello-multiplier 3

Parameters

multiplier

The multiplier for the Hello interval expressed as a decimal integer.

Values 2 to 100

Platforms

7705 SAR Gen 2

hello-multiplier

Syntax

hello-multiplier *deci-units*

no hello-multiplier

Context

[Tree] (config>service>vprn>pim>if hello-multiplier)

Full Context

configure service vprn pim interface hello-multiplier

Description

This command configures the multiplier to determine the hold time for a PIM neighbor on this interface.

The **hello-multiplier** in conjunction with the **hello-interval** determines the holdtime for a PIM neighbor.

Default

hello-multiplier 35

Parameters

deci-units

Specify the value, specified in multiples of 0.1, for the formula used to calculate the holdtime based on the **hello-multiplier**:

(**hello-interval** X **hello-multiplier**) / 10

This allows the PIMv2 default **hello-multiplier** of 3.5 and the default timeout of 105 seconds to be supported.

Values 20 to 100

Platforms

7705 SAR Gen 2

hello-multiplier

Syntax

hello-multiplier *deci-units*

no hello-multiplier

Context

[Tree] (config>router>pim>interface hello-multiplier)

Full Context

configure router pim interface hello-multiplier

Description

This command configures the multiplier to determine the holdtime for a PIM neighbor on this interface.

The **hello-multiplier** in conjunction with the **hello-interval** determines the holdtime for a PIM neighbor.

The **no** form of this command reverts to the default value.

Default

hello-multiplier 35

Parameters

deci-units

Specifies the value, in multiples of 0.1, for the formula used to calculate the holdtime based on the **hello-multiplier**:

(**hello-interval** X **hello-multiplier**) / 10

This allows the PIMv2 default **hello-multiplier** of 3.5 and the default timeout of 105 seconds to be supported.

Values 20 to 100

Default 35

Platforms

7705 SAR Gen 2

hello-multiplier

Syntax

hello-multiplier *multiplier*

no hello-multiplier

Context

[Tree] (config>router>isis>if>level hello-multiplier)

Full Context

configure router isis interface level hello-multiplier

Description

This command configures a Hello multiplier. The **hello-multiplier**, along with the **hello-interval**, is used to calculate a hold time, which is communicated to a neighbor in a Hello PDU.

The hold time is the time in which the neighbor expects to receive the next Hello PDU. If the neighbor receives a Hello within this time, the hold time is reset. If the neighbor does not receive a Hello within the hold time, it brings the adjacency down.

Note:
The neighbor hold time is (hello multiplier X hello interval) on non-designated intermediate system broadcast interfaces and point-to-point interfaces and is (hello multiplier X hello interval / 3) on designated intermediate system broadcast interfaces. Hello values can be adjusted for faster convergence, but the hold time should always be > 3 to reduce routing instability.

The **no** form of this command reverts to the default value.

Default

hello-multiplier 3

Parameters

multiplier

Specifies the multiplier for the Hello interval expressed as a decimal integer.

Values 2 to 100

Platforms

7705 SAR Gen 2

12.17 hello-padding

hello-padding

Syntax

hello-padding {**none** | **adaptive** | **loose** | **strict**}

no hello-padding

Context

[Tree] (config>service>vprn>isis hello-padding)

[Tree] (config>service>vprn>isis>if hello-padding)

[Tree] (config>service>vprn>isis>if>level hello-padding)

[Tree] (config>service>vprn>isis>level hello-padding)

Full Context

configure service vprn isis hello-padding

configure service vprn isis interface hello-padding

configure service vprn isis interface level hello-padding

configure service vprn isis level hello-padding

Description

This command enables the IS-IS Hello (IIH) message padding to ensure that IS-IS LSPs can traverse the link. When this option is enabled, IS-IS Hello messages are padded to the maximum LSP MTU value, which can be set with the **lsp-mtu-size** command. If link MTU is greater than the maximum LSP MTU value, padding to the link MTU is applied.

The **no** form of this command disables IS-IS Hello message padding at this level. However, the router may still perform Hello padding if it was set at a higher level in the configuration. To ensure that Hello message padding is disabled, set all levels of configuration to **no hello-padding**.

Default

no hello-padding

Parameters

adaptive

Specifies the adaptive padding option; this option is able to detect MTU asymmetry from one side of the connection but uses more overhead than loose padding.

- point-to-point interface—Hello PDUs are padded until the sender declares an adjacency on the link to be in the state up. If the implementation supports RFC 3373/5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*, then this is when the three-way state is up. If the implementation uses the "classic" algorithm described in ISO 10589, this is when the adjacency state is up. If the neighbor does not support the adjacency state TLV, then padding continues.
- broadcast interface—Padding starts until at least one adjacency is up on the interface.

loose

Specifies the loose padding option; the loose padding may not be able to detect certain conditions such as asymmetrical MTUs between the routing devices.

- point-to-point interface—the Hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the INIT state
- broadcast interface—padding starts until at least one adjacency (broadcast only has up/down) is up on the interface

none

Specifies that the Hello message padding is not enabled at this level, even if it is configured at one of the parent levels.

strict

Specifies the strict padding option.

- point-to-point interface—padding is done for all adjacency states, and is continuous. Strict padding has the most overhead but detects MTU issues on both sides of a link
- broadcast interface—padding is done for all adjacency states, and is continuous. Strict padding has the most overhead but detects MTU issues on both sides of a link

Platforms

7705 SAR Gen 2

hello-padding

Syntax

[**no**] **hello-padding** {**none** | **adaptive** | **loose** | **strict**}

Context

[Tree] (config>router>isis>level hello-padding)

[Tree] (config>router>isis>interface>level hello-padding)

[Tree] (config>router>isis>interface hello-padding)

[Tree] (config>router>isis hello-padding)

Full Context

configure router isis level hello-padding

configure router isis interface level hello-padding

configure router isis interface hello-padding

configure router isis hello-padding

Description

This command enables IS-IS Hello (IIH) message padding to ensure that IS-IS LSPs can traverse the link. When this option is enabled, IS-IS Hello messages are padded to the maximum LSP MTU value, which can be set with the **lsp-mtu-size** command. If link MTU is greater than the maximum LSP MTU value, padding to the link MTU is applied.

The **no** form of this command disables IS-IS Hello padding at this level. However, the router may still perform Hello padding if it was set at a higher level in the configuration. To ensure that Hello message padding is disabled, set all levels of configuration to **no hello-padding**.

Default

no hello-padding

Parameters

none

Specifies that the Hello message padding is not enabled at this level, even if it is configured at one of the parent levels.

adaptive

Specifies the adaptive padding option; this option is able to detect LSP MTU asymmetry from one side of the connection but uses more overhead than loose padding.

1. point-to-point interface—Hello PDUs are padded until the sender declares an adjacency on the link to be in state up. If the implementation supports RFC 3373/5303, "Three-Way Handshake for IS-IS Point-to- Point Adjacencies” then this is when the three-way state is Up. If the implementation use the "classic” algorithm described in ISO 10589, this is when adjacency state is Up. If the neighbor does not support the adjacency state TLV, then padding continues.

2. broadcast interface—Padding starts until at least one adjacency is up on the interface.

loose

Specifies the loose padding option; the loose padding may not be able to detect certain situations such as asymmetrical LSP MTUs between the routing devices.

1. point-to-point interface—The Hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the INIT state.
2. broadcast interface—Padding starts until there is at least one adjacency (broadcast only has up/down) is up on the interface.

strict

Specifies the strict padding option; this option is the most overhead-intensive but detects LSP MTU issues on both sides of a link.

1. point-to-point interface—Padding is done for all adjacency states, and is continuous.
2. broadcast interface—Padding is done for all adjacency states, and is continuous.

Platforms

7705 SAR Gen 2

12.18 hello-reduction

hello-reduction

Syntax

hello-reduction {**enable** *factor* | **disable**}

no hello-reduction

Context

[Tree] (config>router>ldp>targ-session>peer-template hello-reduction)

[Tree] (config>router>ldp>targ-session>ipv6 hello-reduction)

[Tree] (config>router>ldp>targ-session>ipv4 hello-reduction)

[Tree] (config>router>ldp>targ-session>peer hello-reduction)

Full Context

configure router ldp targeted-session peer-template hello-reduction

configure router ldp targeted-session ipv6 hello-reduction

configure router ldp targeted-session ipv4 hello-reduction

configure router ldp targeted-session peer hello-reduction

Description

This command enables the suppression of periodic targeted Hello messages between LDP peers once the targeted LDP session is brought up.

When this feature is enabled, the target Hello adjacency is brought up by advertising the Hold-Time value the user configured in the " **hello** timeout" parameter for the targeted session. The LSR node will then start advertising an exponentially increasing Hold-Time value in the Hello message as soon as the targeted LDP session to the peer is up. Each new incremented Hold-Time value is sent in a number of Hello messages equal to the value of the argument *factor*, which represents the dampening factor, before the next exponential value is advertised. This provides time for the two peers to settle on the new value. When the Hold-Time reaches the maximum value of 0xffff (binary 65535), the two peers will send Hello messages at a frequency of every [(65535-1)/local helloFactor] seconds for the lifetime of the targeted-LDP session (for example, if the local Hello Factor is three (3), then Hello messages will be sent every 21844 seconds.

The LSR node continues to compute the frequency of sending the Hello messages based on the minimum of its local Hold-time value and the one advertised by its peer as in RFC 5036. Thus for the targeted LDP session to suppress the periodic Hello messages, both peers must bring their advertised Hold-Time to the maximum value. If one of the LDP peers does not, the frequency of the Hello messages sent by both peers will continue to be governed by the smaller of the two Hold-Time values.

When the user enables the Hello reduction option on the LSR node while the targeted LDP session to the peer is operationally up, the change will take effect immediately. In other words, the LSR node will start advertising an exponentially increasing Hold-Time value in the Hello message, starting with the current configured Hold-Time value.

When the user disables the Hello reduction option while the targeted LDP session to the peer is operationally up, the change in the Hold-Time from 0xffff (binary 65535) to the user configured value for this peer will take effect immediately. The local LSR will immediately advertise the value of the user configured Hold-Time value and will not wait until the next scheduled time to send a Hello to make sure the peer adjusts its local hold timeout value immediately.

In general, any configuration change to the parameters of the T-LDP Hello adjacency (modifying the Hello adjacency Hello Timeout or factor, enabling/disabling Hello reduction, or modifying Hello reduction factor) will cause the LSR node to trigger immediately an updated Hello message with the updated Hold Time value without waiting for the next scheduled time to send a Hello.

The **no** form of this command disables the Hello reduction feature.

Default

no hello-reduction

Parameters

factor

Specifies the integer that specifies the Hello reduction dampening factor.

Values 3 to20

Platforms

7705 SAR Gen 2

12.19 hello-time

hello-time

Syntax

hello-time *hello-time*

no hello-time [*hello-time*]

Context

[Tree] (config>service>vpls>stp hello-time)

[Tree] (config>service>template>vpls-template>stp hello-time)

Full Context

configure service vpls stp hello-time

configure service template vpls-template stp hello-time

Description

This command configures the Spanning Tree Protocol (STP) Hello time for the Virtual Private LAN Service (VPLS) STP instance.

The Hello time parameter defines the default timer value that controls the sending interval between BPDU configuration messages by this bridge, on ports where this bridge assumes the designated role.

The active Hello time for the spanning tree is determined by the root bridge (except when the STP is running in RSTP mode, then the Hello time is always taken from the locally configured parameter).

The configured **hello-time** can also be used to calculate the forward delay. See **auto-edge** (**config>service>vpls>sap>stp auto-edge, config>service>template>vpls-sap-template>stp auto-edge, config>service>vpls>spoke-sdp>stp auto-edge**).

The **no** form of this command returns the Hello time to the default value.

Default

hello-time 2

Parameters

hello-time

The Hello time for the STP instance in seconds.

Values 1 to 10

Platforms

7705 SAR Gen 2

hello-time

Syntax

[**no**] **hello-time** *seconds*

Context

[Tree] (config>service>sdp>keep-alive hello-time)

Full Context

configure service sdp keep-alive hello-time

Description

This command configures the time period between SDP keepalive messages on the SDP-ID for the SDP connectivity monitoring messages.

The **no** form of this command reverts the **hello-time** *seconds* value to the default setting.

Default

hello-time 10

Parameters

seconds

Specifies the time period in seconds between SDP keepalive messages, expressed as a decimal integer.

Values 1 to 3600

Platforms

7705 SAR Gen 2

12.20 help

help

Syntax

help

help edit

help global

help special-characters

Context

[Tree] (help)

Full Context

help

Description

This command provides a brief description of the help system. The following information is shown:

```
Help may be requested at any point by hitting a question mark '?'.
In case of an executable node, the syntax for that node will be displayed with an
explanation of all parameters.
In case of sub-commands, a brief description is provided.
Global Commands:
Help on global commands can be observed by issuing "help globals" at any time.
Editing Commands:
Help on editing commands can be observed by issuing "help edit" at any time.
```

Parameters

help

Displays a brief description of the help system.

edit

Displays help on editing.

Available editing keystrokes:

```
Delete current character.....................Ctrl-d
Delete text up to cursor.....................Ctrl-u
Delete text after cursor.....................Ctrl-k
Move to beginning of line....................Ctrl-a
Move to end of line..........................Ctrl-e
Get prior command from history...............Ctrl-p
Get next command from history................Ctrl-n
Move cursor left.............................Ctrl-b
Move cursor right............................Ctrl-f
Move back one word...........................Esc-b
Move forward one word........................Esc-f
Convert rest of word to uppercase............Esc-c
Convert rest of word to lowercase............Esc-l
Delete remainder of word.....................Esc-d
Delete word up to cursor.....................Ctrl-w
Transpose current and previous character.....Ctrl-t
Enter command and return to root prompt......Ctrl-z
Refresh input line...........................Ctrl-l
```

global

Displays help on global commands.

Available global commands:

```
back            - Go back a level in the command tree
echo            - Echo the text that is typed in
exec            - Execute a file - use -echo to show the commands and
                  prompts on the screen
exit            - Exit to intermediate mode - use option all to exit to
```

```
                  root prompt
help            - Display help
history         - Show command history
info            - Display configuration for the present node
logout          - Log off this system
oam             + OAM Test Suite
ping            - Verify the reachability of a remote host
pwc             - Show the present working context
sleep           - Sleep for specified number of seconds
ssh             - SSH to a host
telnet          - Telnet to a host
traceroute      - Determine the route to a destination address
tree            - Display command tree structure from the context of
                  execution
write           - Write text to another user
```

special-characters

Displays help on special characters.

Use the following CLI commands to display more information about commands and command syntax:

?

Lists all commands in the current context.

string?

Lists all commands available in the current context that start with the string.

command ?

Displays command's syntax and associated keywords.

string<Tab> or string<Space>

Completes a partial command name (auto-completion) or lists available commands that match the string.

Platforms

7705 SAR Gen 2

12.21 helper-disable

helper-disable

Syntax

[**no**] **helper-disable**

Context

[Tree] (config>service>vprn>isis>graceful-restart helper-disable)

Full Context

configure service vprn isis graceful-restart helper-disable

Description

This command disables helper support for IS-IS graceful restart (GR).

When **graceful-restart** is enabled, the router can be a helper (that is, the router is helping a neighbor to restart), a restarting router, or both. The router only supports the helper mode. It will not act as a restarting router, because the high availability feature set already preserves IS-IS forwarding information such that this functionality is not needed.

Note:
This command is a historical command and should not be disabled. Configuring **helper-disable** has the effect of disabling graceful restart, because the router only supports helper mode.

The **no helper-disable** command enables helper support and is the default when graceful restart is enabled.

Default

no helper-disable

Platforms

7705 SAR Gen 2

helper-disable

Syntax

[**no**] **helper-disable**

Context

[Tree] (config>service>vprn>ospf>graceful-restart helper-disable)

[Tree] (config>service>vprn>ospf3>graceful-restart helper-disable)

Full Context

configure service vprn ospf graceful-restart helper-disable

configure service vprn ospf3 graceful-restart helper-disable

Description

This command disables helper support for OSPF graceful restart (GR).

When **graceful-restart** is enabled, the router can be a helper (that is, the router is helping a neighbor to restart), a restarting router, or both. The router only supports helper mode. It will not act as a restarting router, because the high availability feature set already preserves OSPF forwarding information such that this functionality is not needed.

Note:
This command is a historical command and should not be disabled. Configuring **helper-disable** has the effect of disabling graceful restart, because the router only supports helper mode.

The **no helper-disable** command enables helper support and is the default when graceful restart is enabled.

Default

no helper-disable

Platforms

7705 SAR Gen 2

helper-disable

Syntax

[**no**] **helper-disable**

Context

[Tree] (config>router>isis>graceful-restart helper-disable)

Full Context

configure router isis graceful-restart helper-disable

Description

This command disables helper support for IS-IS graceful restart (GR).

When **graceful-restart** is enabled, the router can be a helper (that is, the router is helping a neighbor to restart), a restarting router, or both. The router only supports the helper mode. It will not act as a restarting router, because the high availability feature set already preserves IS-IS forwarding information so that this functionality is not needed.

Note:
This command is a historical command and should not be disabled. Configuring **helper-disable** has the effect of disabling graceful restart, because the router only supports helper mode.

The **no** form of this command enables helper support and is the default when graceful restart is enabled.

Platforms

7705 SAR Gen 2

helper-disable

Syntax

[**no**] **helper-disable**

Context

[Tree] (config>router>ospf3>graceful-restart helper-disable)

[Tree] (config>router>ospf>graceful-restart helper-disable)

Full Context

configure router ospf3 graceful-restart helper-disable

configure router ospf graceful-restart helper-disable

Description

This command disables helper support for OSPF graceful restart (GR).

When **graceful-restart** is enabled, the router can be a helper (that is, the router is helping a neighbor to restart), a restarting router, or both. The router only supports the helper mode. It will not act as a restarting router because the high availability feature set already preserves OSPF forwarding information so that this functionality is not needed.

Note:
This command is a historical command and should not be disabled. Configuring **helper-disable** has the effect of disabling graceful restart, because the router only supports helper mode.

The **no** form of this command enables helper support and is the default when **graceful-restart** is enabled.

Default

no helper-disable

Platforms

7705 SAR Gen 2

12.22 helper-override-restart-time

helper-override-restart-time

Syntax

helper-override-restart-time *seconds*

no helper-override-restart-time

Context

[Tree] (config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived helper-override-restart-time)

[Tree] (config>service>vprn>bgp>group>graceful-restart>long-lived helper-override-restart-time)

[Tree] (config>service>vprn>bgp>graceful-restart>long-lived helper-override-restart-time)

Full Context

configure service vprn bgp group neighbor graceful-restart long-lived helper-override-restart-time

configure service vprn bgp group graceful-restart long-lived helper-override-restart-time

configure service vprn bgp graceful-restart long-lived helper-override-restart-time

Description

This command overrides the restart-time advertised by a peer (in its GR capability) with a locally-configured value. This override applies only to AFI/SAFI that were included in the GR capability of the peer. The restart-time is always zero for AFI/SAFI not included in the GR capability. This command is useful if the local router wants to force LLGR phase to begin after a set time for all protected AFI/SAFI.

By default, the restart time for all AFI/SAFI in the GR capability is the value signaled by the peer.

Default

no helper-override-restart-time

Parameters

seconds

The locally-imposed restart time for all AFI/SAFI included in the peer's GR capability.

Values 0 to 4095

Platforms

7705 SAR Gen 2

helper-override-restart-time

Syntax

helper-override-restart-time *seconds*

no helper-override-restart-time

Context

[Tree] (config>router>bgp>group>neighbor>graceful-restart>long-lived helper-override-restart-time)

[Tree] (config>router>bgp>graceful-restart>long-lived helper-override-restart-time)

[Tree] (config>router>bgp>group>graceful-restart>long-lived helper-override-restart-time)

Full Context

configure router bgp group neighbor graceful-restart long-lived helper-override-restart-time

configure router bgp graceful-restart long-lived helper-override-restart-time

configure router bgp group graceful-restart long-lived helper-override-restart-time

Description

This command overrides the restart-time advertised by a peer (in its GR capability) with a locally-configured value. This override applies only to AFI/SAFI that were included in the GR capability of the peer. The restart-time is always zero for AFI/SAFI not included in the GR capability. This command is useful if the local router wants to force LLGR phase to begin after a set time for all protected AFI/SAFI.

By default, the restart time for all AFI/SAFI in the GR capability is the value signaled by the peer.

Default

no helper-override-restart-time

Parameters

seconds

The locally-imposed restart time for all AFI/SAFI included in the peer's GR capability.

Values 0 to 4095

Platforms

7705 SAR Gen 2

12.23 helper-override-stale-time

helper-override-stale-time

Syntax

helper-override-stale-time *seconds*

no helper-override-stale-time

Context

[Tree] (config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived>family helper-override-stale-time)

[Tree] (config>service>vprn>bgp>group>graceful-restart>long-lived helper-override-stale-time)

[Tree] (config>service>vprn>bgp>graceful-restart>long-lived helper-override-stale-time)

[Tree] (config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived helper-override-stale-time)

[Tree] (config>service>vprn>bgp>group>graceful-restart>long-lived>family helper-override-stale-time)

[Tree] (config>service>vprn>bgp>graceful-restart>long-lived>family helper-override-stale-time)

Full Context

configure service vprn bgp group neighbor graceful-restart long-lived family helper-override-stale-time

configure service vprn bgp group graceful-restart long-lived helper-override-stale-time

configure service vprn bgp graceful-restart long-lived helper-override-stale-time

configure service vprn bgp group neighbor graceful-restart long-lived helper-override-stale-time

configure service vprn bgp group graceful-restart long-lived family helper-override-stale-time

configure service vprn bgp graceful-restart long-lived family helper-override-stale-time

Description

This command overrides the LLGR stale-time advertised by a peer (in its LLGR capability) with a locally-configured value. When configured in the long-lived configuration context, **helper-override-stale-time** applies to all AFI/SAFI in the advertised LLGR capability except for any AFI/SAFI with a family-specific override. A family-specific override is configured with the **helper-override-stale-time** command in a family context.

By default, the LLGR stale-time for an AFI/SAFI is the value signaled by the peer in the corresponding AFI/SAFI part of the LLGR capability.

Default

no helper-override-stale-time

Parameters

seconds

Specifies the locally imposed LLGR stale time in seconds.

Values 0 to 16777215

Platforms

7705 SAR Gen 2

helper-override-stale-time

Syntax

helper-override-stale-time *seconds*

no helper-override-stale-time

Context

[Tree] (config>router>bgp>group>neighbor>graceful-restart>long-lived>family helper-override-stale-time)

[Tree] (config>router>bgp>graceful-restart>long-lived helper-override-stale-time)

[Tree] (config>router>bgp>group>graceful-restart>long-lived helper-override-stale-time)

[Tree] (config>router>bgp>graceful-restart>long-lived>family helper-override-stale-time)

[Tree] (config>router>bgp>group>graceful-restart>long-lived>family helper-override-stale-time)

[Tree] (config>router>bgp>group>neighbor>graceful-restart>long-lived helper-override-stale-time)

Full Context

configure router bgp group neighbor graceful-restart long-lived family helper-override-stale-time

configure router bgp graceful-restart long-lived helper-override-stale-time

configure router bgp group graceful-restart long-lived helper-override-stale-time

configure router bgp graceful-restart long-lived family helper-override-stale-time

configure router bgp group graceful-restart long-lived family helper-override-stale-time

configure router bgp group neighbor graceful-restart long-lived helper-override-stale-time

Description

This command overrides the LLGR stale-time advertised by a peer (in its LLGR capability) with a locally-configured value. When configured in the **long-lived** configuration context, **helper-override-stale-time** applies to all AFI/SAFI in the advertised LLGR capability except for any AFI/SAFI with a family-specific override. A family-specific override is configured with the **helper-override-stale-time** command in a family context.

By default, the LLGR stale-time for an AFI/SAFI is the value signaled by the peer in the corresponding AFI/SAFI part of the LLGR capability.

Default

no helper-override-stale-time

Parameters

seconds

Specifies the locally imposed LLGR stale time in seconds.

Values 0 to 16777215

Platforms

7705 SAR Gen 2

12.24 high

high

Syntax

high

Context

[Tree] (config>qos>sap-egress>queue>drop-tail high)

Full Context

configure qos sap-egress queue drop-tail high

Description

Commands in this context configure the queue high drop tail parameters. The high drop tail defines the queue depth beyond which in-profile packets will not be accepted into the queue and will be discarded.

Platforms

7705 SAR Gen 2

high

Syntax

high

Context

[Tree] (cfg>qos>qgrps>egr>qgrp>queue>drop-tail high)

Full Context

configure qos queue-group-templates egress queue-group queue drop-tail high

Description

Commands in this context configure the queue high drop-tail parameters. The high drop tail defines the queue depth beyond which in-profile packets will not be accepted into the queue and will be discarded.

Platforms

7705 SAR Gen 2

12.25 high-octets-discarded-count

high-octets-discarded-count

Syntax

[**no**] **high-octets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>queue>i-counters high-octets-discarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters high-octets-discarded-count)

Full Context

configure log accounting-policy custom-record queue i-counters high-octets-discarded-count

configure log accounting-policy custom-record ref-queue i-counters high-octets-discarded-count

Description

This command includes the high octets discarded count.

The **no** form of this command excludes the high octets discarded count.

Default

no high-octets-discarded-count

Platforms

7705 SAR Gen 2

12.26 high-octets-offered-count

high-octets-offered-count

Syntax

[**no**] **high-octets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters high-octets-offered-count)

[Tree] (config>log>acct-policy>cr>queue>i-counters high-octets-offered-count)

Full Context

configure log accounting-policy custom-record ref-queue i-counters high-octets-offered-count

configure log accounting-policy custom-record queue i-counters high-octets-offered-count

Description

This command includes the high octets offered count.

The **no** form of this command excludes the high octets offered count.

Default

no high-octets-offered-count

Platforms

7705 SAR Gen 2

12.27 high-packets-discarded-count

high-packets-discarded-count

Syntax

[**no**] **high-packets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters high-packets-discarded-count)

[Tree] (config>log>acct-policy>cr>queue>i-counters high-packets-discarded-count)

Full Context

configure log accounting-policy custom-record ref-queue i-counters high-packets-discarded-count

configure log accounting-policy custom-record queue i-counters high-packets-discarded-count

Description

This command includes the high packets discarded count.

The **no** form of this command excludes the high packets discarded count.

Default

no high-packets-discarded-count

Platforms

7705 SAR Gen 2

12.28 high-packets-offered-count

high-packets-offered-count

Syntax

[**no**] **high-packets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>queue>i-counters high-packets-offered-count)

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters high-packets-offered-count)

Full Context

configure log accounting-policy custom-record queue i-counters high-packets-offered-count

configure log accounting-policy custom-record ref-queue i-counters high-packets-offered-count

Description

This command includes the high packets offered count.

The **no** form of this command excludes the high packets offered count.

Default

no high-packets-offered-count

Platforms

7705 SAR Gen 2

12.29 high-prio-only

high-prio-only

Syntax

high-prio-only *percent-of-mbs*

no high-prio-only

Context

[Tree] (config>qos>sap-egress>policer high-prio-only)

[Tree] (config>qos>sap-ingress>policer high-prio-only)

Full Context

configure qos sap-egress policer high-prio-only

configure qos sap-ingress policer high-prio-only

Description

This command is used to configure the percentage of the policer’s PIR leaky bucket's MBS (maximum burst size) that is reserved for high-priority traffic. While the **mbs** value defines the policer’s high-priority violate threshold, the percentage value defined is applied to the **mbs** value to derive the bucket’s low-priority violate threshold. See the **mbs** command details for information about which types of traffic are associated with each violate threshold.

Parameters

percent-of-mbs

The *percent-of-mbs* parameter is required when specifying **high-prio-only** and is expressed as a percentage.

Values 0 to 100

Default 10

Platforms

7705 SAR Gen 2

high-prio-only

Syntax

high-prio-only *percent-of-mbs*

no high-prio-only

Context

[Tree] (config>qos>qgrps>egr>qgrp>policer high-prio-only)

[Tree] (config>qos>qgrps>ing>qgrp>policer high-prio-only)

Full Context

configure qos queue-group-templates egress queue-group policer high-prio-only

configure qos queue-group-templates ingress queue-group policer high-prio-only

Description

This command is used to configure the percentage of the policer's PIR leaky bucket's MBS (maximum burst size) that is reserved for high-priority traffic. While the **mbs** value defines the policer's high-priority violate threshold, the percentage value defined is applied to the **mbs** value to derive the bucket's low-priority violate threshold. See the **mbs** command details for information on which types of traffic is associated with each violate threshold.

Parameters

percent-of-mbs

The *percent-of-mbs* parameter is required when specifying **high-prio-only** and is expressed as a percentage.

Values 0 to 100

Default 10

Platforms

7705 SAR Gen 2

12.30 highplus

highplus

Syntax

highplus

Context

[Tree] (config>qos>sap-egress>queue>drop-tail highplus)

Full Context

configure qos sap-egress queue drop-tail highplus

Description

Commands in this context configure the queue highplus drop tail parameters. The highplus drop tail defines the queue depth beyond which inplus-profile packets will not be accepted into the queue and will be discarded.

Platforms

7705 SAR Gen 2

highplus

Syntax

highplus

Context

[Tree] (cfg>qos>qgrps>egr>qgrp>queue>drop-tail highplus)

Full Context

configure qos queue-group-templates egress queue-group queue drop-tail highplus

Description

Commands in this context configure the queue highplus drop-tail parameters. The highplus drop tail defines the queue depth beyond which inplus-profile packets will not be accepted into the queue and will be discarded.

Platforms

7705 SAR Gen 2

12.31 history

history

Syntax

history

Context

[Tree] (history)

Full Context

history

Description

This command lists the last 30 commands entered in this session.

Re-execute a command in the history with the **!n** command, where **n** is the line number associated with the command in the history output.

Example:

```
A:ALA-1# history
  68 info
  69 exit
  70 info
  71 filter
  72 exit all
  73 configure
  74 router
  75 info
  76 interface "test"
  77 exit
  78 reduced-prompt
  79 info
  80 interface "test"
  81 icmp unreachables exit all
  82 exit all
  83 reduced-prompt
  84 configure router
  85 interface
  86 info
  87 interface "test"
  88 info
  89 reduced-prompt
  90 exit all
  91 configure
  92 card 1
  93 card-type
  94 exit
  95 router
  96 exit
  97 history
A:ALA-1# !91
A:ALA-1# configure
A:ALA-1>config#
```

Platforms

7705 SAR Gen 2

history

Syntax

history

Context

[Tree] (config>system>management-interface>cli>md-cli>environment history)

Full Context
configure system management-interface cli md-cli environment history

Description
Commands in this context configure the command history.

Platforms
7705 SAR Gen 2

history

Syntax
[**no**] **history**

Context
[Tree] (debug>system>nsp-proxy history)

Full Context
debug system nsp-proxy history

Description
This command enables the NSP proxy history for debugging purposes.

The **no** form of this command disables the NSP proxy history.

Default
no history

Platforms
7705 SAR Gen 2

12.32 history-size

history-size

Syntax
history-size *size*

no history-size

Context
[Tree] (config>system>security>password history-size)

Full Context

configure system security password history-size

Description

Configure how many previous passwords a new password is matched against.

Default

history-size 0

Parameters

size

Specifies how many previous passwords a new password is matched against.

Values 0 to 20

Platforms

7705 SAR Gen 2

12.33 hli-event

hli-event

Syntax

hli-event {**forward** | **backward** | **aggregate**} **threshold** *raise-threshold* [**clear** *clear-threshold*]

no hli-event {**forward** | **backward** | **aggregate**}

Context

[Tree] (config>oam-pm>session>ip>twamp-light>loss-events hli-event)

Full Context

configure oam-pm session ip twamp-light loss-events hli-event

Description

This command sets the high loss interval (HLI) threshold to be monitored and the associated thresholds using the counter of the specified direction. The aggregate is a function of summing forward and backward. This value is only used as a threshold mechanism and is not part of the stored statistics. If the optional **clear** *clear-threshold* parameter is not specified, the traffic crossing alarm is stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and regardless of any previous window. Each unique event can only be raised once within measurement interval. If the optional **clear** *clear-threshold* parameter is specified, the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to following measurement intervals. If a threshold crossing event is raised another is raised until a measurement

interval completes and the clear threshold has not been exceeded. A clear event is raised under that condition.

The **no** form of this command removes the event threshold for frame loss ratio. The direction must be included with the **no** command.

Default

no hli-event forward

no hli-event backward

no hli-event aggregate

Parameters

forward

Specifies the threshold is applied to the forward direction count.

backward

Specifies the threshold is applied to the backward direction count.

aggregate

Specifies the threshold is applied to the aggregate count (sum of forward and backward).

raise-threshold

Specifies the rising threshold that determines when the event is to be generated, when the percentage of loss value is reached.

Values 1 to 864000

clear-threshold

Specifies an optional value used for stateful behavior that allows the operator to configure a percentage of loss value lower than the rising percentage to indicate when the clear event should be generated.

Values 0 to 863999

A value of zero means that the HLI counter must be 0.

Platforms

7705 SAR Gen 2

12.34 hli-force-count

hli-force-count

Syntax

[**no**] **hli-force-count**

Context

[Tree] (config>oam-pm>session>ip>twamp-light>loss hli-force-count)

Full Context

configure oam-pm session ip twamp-light loss hli-force-count

Description

This command allows High Loss Interval (HLI) and Consecutive High Loss Interval (CHLI) counters to increment regardless of availability. Without this command, HLI and CHLI counters can only increment during times of availability, which includes undetermined availability. During times of complete packet loss, the forward direction HLI is marked as high loss. The backward direction is not marked as high loss during times of complete packet loss.

The **no** form of this command configures HLI and CHLI counters to increment during times of availability only.

Platforms

7705 SAR Gen 2

12.35 hold-clear

hold-clear

Syntax

hold-clear *seconds*

no hold-clear

Context

[Tree] (config>vrrp>policy>priority-event>lag-port-down hold-clear)

[Tree] (config>vrrp>policy>priority-event>route-unknown hold-clear)

[Tree] (config>vrrp>policy>priority-event>port-down hold-clear)

[Tree] (config>vrrp>policy>priority-event>mc-ipsec-non-forwarding hold-clear)

[Tree] (config>vrrp>policy>priority-event>host-unreachable hold-clear)

Full Context

configure vrrp policy priority-event lag-port-down hold-clear

configure vrrp policy priority-event route-unknown hold-clear

configure vrrp policy priority-event port-down hold-clear

configure vrrp policy priority-event mc-ipsec-non-forwarding hold-clear

configure vrrp policy priority-event host-unreachable hold-clear

Description

This command configures the hold clear time for the event. The *seconds* parameter specifies the hold-clear time, the amount of time in seconds by which the effect of a cleared event on the associated virtual router instance is delayed.

The hold-clear time is used to prevent black hole conditions when a virtual router instance advertises itself as a master before other conditions associated with the cleared event have had a chance to enter a forwarding state.

Default

no hold-clear

Parameters

seconds

Specifies the amount of time in seconds by which the effect of a cleared event on the associated virtual router instance is delayed.

Values 0 to 86400

Platforms

7705 SAR Gen 2

12.36 hold-count

hold-count

Syntax

hold-count *BPDU tx hold count*

no hold-count

Context

[Tree] (config>service>vpls>stp hold-count)

[Tree] (config>service>template>vpls-template>stp hold-count)

Full Context

configure service vpls stp hold-count

configure service template vpls-template stp hold-count

Description

This command configures the peak number of BPDUs that can be transmitted in a period of one second.

The **no** form of this command returns the hold count to the default value

Default

hold-count 6

Parameters

BPDU tx hold count

The hold count for the STP instance in seconds

Values 1 to 10

Platforms

7705 SAR Gen 2

12.37 hold-down-time

hold-down-time

Syntax

hold-down-time [**sec** *seconds*] [**min** *minutes*] [**hrs** *hours*] [**days** *days*]

no hold-down-time

Context

[Tree] (config>aaa>radius-srv-plcy>servers hold-down-time)

Full Context

configure aaa radius-server-policy servers hold-down-time

Description

This command determines the interval during which no new communication attempts are made to a RADIUS server that is marked down to prevent immediately overloading the server when it is starting up. The only exception is when all servers in the authentication policy are marked down; in that case, they will all be used again to prevent failures on new client connections.

The **no** form of this command reverts to the default.

Default

hold-down-time sec 30

Parameters

days

Specifies the hold time in days before re-using a RADIUS server that was down.

Values 1 to 1

hours

Specifies the hold time in hours before re-using a RADIUS server that was down.

Values 1 to 23

minutes

Specifies the hold time in minutes before re-using a RADIUS server that was down.

Values 1 to 59

seconds

Specifies the hold time in seconds before re-using a RADIUS server that was down.

Values 1 to 59

Platforms

7705 SAR Gen 2

hold-down-time

Syntax

hold-down-time *seconds*

no hold-down-time

Context

[Tree] (config>service>sdp>keep-alive hold-down-time)

Full Context

configure service sdp keep-alive hold-down-time

Description

This command configures the minimum time period the SDP will remain in the operationally down state in response to SDP keepalive monitoring.

This parameter can be used to prevent the SDP operational state from "flapping" by rapidly transitioning between the operationally up and operationally down states based on keepalive messages.

When an SDP keepalive response is received that indicates an error condition or the **max-drop-count** keepalive messages receive no reply, the *sdp-id* will immediately be brought operationally down. If a keepalive response is received that indicates the error has cleared, the *sdp-id* will be eligible to be put into the operationally up state only after the **hold-down-time** interval has expired.

The **no** form of this command reverts the **hold-down-time seconds** *value* to the default setting.

Default

hold-down-time 10

Parameters

seconds

Specifies time, in seconds, expressed as a decimal integer. The SDP ID will remain in the operationally down state before it is eligible to enter the operationally up state. A value of 0 indicates that no **hold-down-time** will be enforced for SDP ID.

Values 0 to 3600

Platforms

7705 SAR Gen 2

12.38 hold-down-timer

hold-down-timer

Syntax

hold-down-timer *hold-down-timer*

no hold-down-timer

Context

[Tree] (config>router>segment-routing>maintenance-policy hold-down-timer)

Full Context

configure router segment-routing maintenance-policy hold-down-timer

Description

This command configures the hold down timer for SR policy candidate paths.

This command is intended to prevent bouncing of the SR policy path state if one or more S-BFD sessions associated with segment lists flap and therefore cause the threshold to be repeatedly crossed in a short period of time. It is started when the number of up S-BFD sessions drops below the threshold. The SR policy path is not considered to be up again until the hold down timer has expired and the number of up S-BFD sessions equals or exceeds the threshold and the internal hold timer is not running.

Note:
If the revert timer is also configured, the revert timer is not started until after the number of S-BFD sessions that are up ≥ threshold and the hold down timer for the primary candidate path has expired.

The **no** form of this command reverts to the default.

Default

hold-down-timer 0

Parameters

hold-down-timer

Specifies the hold-down timer, in deciseconds, in 10ms steps.

Values 0 to 5000

Platforms

7705 SAR Gen 2

12.39 hold-on-neighbor-failure

hold-on-neighbor-failure

Syntax

hold-on-neighbor-failure *multiplier*

no hold-on-neighbor-failure

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-lag hold-on-neighbor-failure)

Full Context

configure redundancy multi-chassis peer mc-lag hold-on-neighbor-failure

Description

This command specifies the interval that the standby node will wait for packets from the active node before assuming a redundant-neighbor node failure. This delay in switch-over operation is required to accommodate different factors influencing node failure detection rate, such as IGP convergence, or HA switch-over times and to prevent the standby node to act prematurely.

The **no** form of this command reverts to the default.

Default

hold-on-neighbor-failure 3

Parameters

multiplier

Specifies the time interval that the standby node waits for packets from the active node before assuming a redundant-neighbor node failure.

Values 2 to 25

Platforms

7705 SAR Gen 2

hold-on-neighbor-failure

Syntax

hold-on-neighbor-failure *multiplier*

no hold-on-neighbor-failure

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ep hold-on-neighbor-failure)

Full Context

configure redundancy multi-chassis peer mc-endpoint hold-on-neighbor-failure

Description

This command specifies the number of keep-alive intervals that the local node will wait for packets from the MC-EP peer before assuming failure. After this time interval passed the all the mc-endpoints configured under services will revert to single chassis behavior, activating the best local pseudowire.

The **no** form of this command sets the multiplier to default value

Default

no hold-on-neighbor-failure

Parameters

multiplier

Specifies the hold time applied on neighbor failure.

Values 2 to 25

Platforms

7705 SAR Gen 2

hold-on-neighbor-failure

Syntax

hold-on-neighbor-failure *multiplier*

no hold-on-neighbor-failure

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ipsec hold-on-neighbor-failure)

Full Context

configure redundancy multi-chassis peer mc-ipsec hold-on-neighbor-failure

Description

This command specifies the number of keep-alive failures before the peer is considered to be down.

The **no** form of this command reverts to the default.

Default

hold-on-neighbor-failure 3

Parameters

multiplier

Specifies the hold time applied on the neighbor failure.

Values 2 to 25

Platforms

7705 SAR Gen 2

12.40 hold-set

hold-set

Syntax

hold-set *seconds*

no hold-set

Context

[Tree] (config>vrrp>policy>priority-event>host-unreachable hold-set)

[Tree] (config>vrrp>policy>priority-event>mc-ipsec-non-forwarding hold-set)

[Tree] (config>vrrp>policy>priority-event>lag-port-down hold-set)

[Tree] (config>vrrp>policy>priority-event>port-down hold-set)

[Tree] (config>vrrp>policy>priority-event>route-unknown hold-set)

Full Context

configure vrrp policy priority-event host-unreachable hold-set

configure vrrp policy priority-event mc-ipsec-non-forwarding hold-set

configure vrrp policy priority-event lag-port-down hold-set

configure vrrp policy priority-event port-down hold-set

configure vrrp policy priority-event route-unknown hold-set

Description

This command specifies the amount of time that must pass before the set state for a VRRP priority control event can transition to the cleared state to dampen flapping events. A flapping event continually transitions between clear and set.

The **hold-set** command is used to dampen the effect of a flapping event. The **hold-set** value is loaded into a hold-set timer that prevents a set event from transitioning to the cleared state until it expires.

Each time an event transitions between cleared and set, the timer is loaded and begins a countdown to zero. When the timer reaches zero, the event is allowed to enter the cleared state. Entering the cleared state is dependent on the object controlling the event, conforming to the requirements defined in the event itself. It is possible, on some event types, to have another set action reload the hold-set timer. This extends the amount of time that must expire before entering the cleared state.

Once the hold-set timer expires and the event meets the cleared state requirements or is set to a lower threshold, the current set effect on the virtual router instances in-use priority can be removed. As with **lag-port-down** events, this may be a decrease in the set effect if the *clearing* amounts to a lower set threshold.

The **hold-set** command can be executed at any time. If the hold-set timer value is configured larger than the new *seconds* setting, the timer is loaded with the new **hold-set** value.

The **no** form of the command disables the hold timer so that event transitions are processed immediately.

Default

no hold-set

Parameters

seconds

The number of seconds that the hold-set timer waits after an event enters a set state or enters a higher threshold set state, depending on the event type.

The value of 0 disables the hold-set timer, preventing any delay in processing lower set thresholds or cleared events.

Values 0 to 86400

Platforms

7705 SAR Gen 2

12.41 hold-time

hold-time

Syntax

hold-time {[**up** *hold-time-up*] [**down** *hold-time-down*] [**seconds** | **centiseconds**]}

no hold-time

Context

[Tree] (config>port>ethernet hold-time)

Full Context

configure port ethernet hold-time

Description

This command configures port link dampening timers which reduce the number of link transitions reported to upper layer protocols. The **hold-time** value dampens interface transitions.

When an interface transitions from an up state to a down state, it is immediately advertised to the rest of the system if the hold-time down interval is zero, but if the hold-time down interval is greater than zero, interface down transitions are not advertised to upper layers until the hold-time down interval has expired. Likewise, an interface is immediately advertised as up to the rest of the system if the hold-time up interval is zero, but if the hold-time up interval is greater than zero, up transitions are not advertised until the hold-time up interval has expired.

For ESM SRRP setup, MCS synchronizes subscriber information between the two chassis. After a chassis recovers from a power reset/down, MCS immediately synchronizes all subscriber information at once. The longer the host list, the longer it will take to synchronize the chassis. In a fully populated chassis, it is recommended to allow at least 45 minutes for MCS synchronization. It is also recommended to hold the port down, facing the subscriber, on the recovering chassis for 45 minutes before it is allowed to forward traffic again.

The **no** form of this command reverts to the default values.

Default

down 0 seconds — No port link down dampening is enabled; link down transitions are immediately reported to upper layer protocols.

up 0 seconds — No port link up dampening is enabled; link up transitions are immediately reported to upper layer protocols.

Parameters

hold-time-up

The delay, in seconds or centiseconds, after which to notify the upper layers when an interface transitions from a down state to an up state.

Values 0 to 36000 seconds, 0 or 10 to 3600000 centiseconds in 5 centisecond increments

hold-time-down

The delay, in seconds or centiseconds, after which to notify the upper layers when an interface transitions from an up state to a down state.

Values 0 to 36000 seconds, 0 or 10 to 3600000 centiseconds in 5 centisecond increments

seconds | centiseconds

Specifies the hold time units as **seconds** or **centiseconds**.

Platforms

7705 SAR Gen 2

hold-time

Syntax

hold-time down *hold-down-time*

no hold-time

Context

[Tree] (config>lag hold-time)

Full Context

configure lag hold-time

Description

This command specifies the timer, in tenths of seconds, which controls the delay between detecting that a LAG is down (all active ports are down) and reporting it to the higher levels.

A non-zero value can be configured, for example, when active/standby signaling is used in a 1:1 fashion to avoid informing higher levels during the small time interval between detecting that the LAG is down and the time needed to activate the standby link.

Default

no hold-time

Parameters

hold-down-time

Specifies the hold-time for event reporting.

Values 0 to 2000

Platforms

7705 SAR Gen 2

hold-time

Syntax

hold-time

Context

[Tree] (config>service>vpls>interface hold-time)

[Tree] (config>service>ies>interface hold-time)

[Tree] (config>service>vprn>network-interface hold-time)

[Tree] (config>service>vprn>interface hold-time)

[Tree] (config>router>if hold-time)

Full Context

configure service vpls interface hold-time

configure service ies interface hold-time

configure service vprn network-interface hold-time

configure service vprn interface hold-time

configure router interface hold-time

Description

This command creates the CLI context to configure interface level hold-up and hold-down timers for the associated IP interface.

The **up** timer controls a delay for the associated IPv4 or IPv6 interface so that the system will delay the deactivation of the associated interface for the specified amount of time.

The **down** timer controls a delay for the associated IPv4 or IPv6 interface so that the system will delay the activation of the associated interface for the specified amount of time

Platforms

7705 SAR Gen 2

hold-time

Syntax

hold-time *seconds* [**min** *seconds2*]

no hold-time

Context

[Tree] (config>service>vprn>bgp hold-time)

[Tree] (config>service>vprn>bgp>group hold-time)

[Tree] (config>service>vprn>bgp>group>neighbor hold-time)

Full Context

configure service vprn bgp hold-time

configure service vprn bgp group hold-time

configure service vprn bgp group neighbor hold-time

Description

This command configures the BGP hold time, expressed in seconds.

The BGP hold time specifies the maximum time BGP waits between successive messages (either keepalive or update) from its peer, before closing the connection. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

Even though the router OS implementation allows setting the **keepalive** (**config>service>vprn>bgp keepalive**, **config>service>vprn>bgp>group keepalive**, **config>service>vprn>bgp>group>neighbor keepalive**) time separately, the configured **keepalive** timer is overridden by the **hold-time** value under the following circumstances:

- If the specified hold-time is less than the configured **keepalive** time, then the operational **keepalive** time is set to a third of the **hold-time**; the configured **keepalive** time is not changed.
- If the **hold-time** is set to zero, then the operational value of the **keepalive** time is set to zero; the configured **keepalive** time is not changed. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

hold-time 90

Parameters

seconds

Specifies the hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently.

Values 0, 3 to 65535

seconds2

Specifies the minimum hold-time that is accepted for the session. If the peer proposes a hold-time lower than this value the session attempt is rejected.

Platforms

7705 SAR Gen 2

hold-time

Syntax

hold-time

Context

[Tree] (config>service>oper-group hold-time)

Full Context

configure service oper-group hold-time

Description

Commands in this context configure hold time information.

Platforms

7705 SAR Gen 2

hold-time

Syntax

hold-time *seconds* [**min** *seconds*]

no hold-time

Context

[Tree] (config>router>bgp>group hold-time)

[Tree] (config>router>bgp hold-time)

[Tree] (config>router>bgp>group>neighbor hold-time)

Full Context

configure router bgp group hold-time

configure router bgp hold-time

configure router bgp group neighbor hold-time

Description

This command configures the BGP hold time, expressed in seconds.

The BGP hold time specifies the maximum time BGP waits between successive messages (either keepalive or update) from its peer, before closing the connection. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

Even though the implementation allows setting the **keepalive** time separately, the configured **keepalive** timer is overridden by the **hold-time** value under the following circumstances:

- If the specified hold-time is less than the configured **keepalive** time, then the operational **keepalive** time is set to a third of the **hold-time**; the configured **keepalive** time is not changed.
- If the **hold-time** is set to zero, then the operational value of the **keepalive** time is set to zero; the configured **keepalive** time is not changed. This means that the connection with the peer is up permanently and no keepalive packets are sent to the peer.

The **no** form of this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

hold-time 90

Parameters

seconds

Specifies the hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently.

Values 0, 3 to 65535

min *seconds2*

Specifies the minimum hold-time that will be accepted for the session. If the peer proposes a hold-time lower than this value, the session attempt will be rejected.

Platforms

7705 SAR Gen 2

12.42 hold-timer

hold-timer

Syntax

hold-timer *seconds*

no hold-timer

Context

[Tree] (config>router>mpls hold-timer)

Full Context

configure router mpls hold-timer

Description

This command specifies the amount of time that the ingress node holds before programming its data plane and declaring the LSP up to the service module. This occurs anytime the ingress node brings up an LSP path or switches traffic from a working path to another working path of the same LSP.

The **no** form of this command reverts to the default value.

Default

no hold-timer

Parameters

seconds

Specifies the time (in seconds), for which the ingress node holds before programming its data plane and declaring the LSP up to the service module.

Values 0 to 1000

Default 1

Platforms

7705 SAR Gen 2

12.43 holddown

holddown

Syntax

[**no**] **holddown** [**neighbor** *ip-int-name* | *ip-address*]

Context

[Tree] (debug>router>rip holddown)

Full Context

debug router rip holddown

Description

This command enables debugging for RIP holddowns.

Parameters

ip-int-name* | *ip-address

Debugs the RIP holddowns sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

holddown

Syntax

[**no**] **holddown** [**neighbor** *ip-int-name* | *ipv6-address*]

Context

[Tree] (debug>router>ripng holddown)

Full Context

debug router ripng holddown

Description

This command enables debugging for RIPng holddowns.

Parameters

ip-int-name*| *ipv6-address

Debugs the RIPng holddowns sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

12.44 holdtime

holdtime

Syntax

holdtime *holdtime*

no holdtime

Context

[Tree] (config>service>vprn>pim>rp>ipv6>rp-candidate holdtime)

[Tree] (config>service>vprn>pim>rp>rp-candidate holdtime)

Full Context

configure service vprn pim rp ipv6 rp-candidate holdtime

configure service vprn pim rp rp-candidate holdtime

Description

This command specifies the length of time a neighbor considers the sending router to be operationally up.

The **no** form of this command reverts to the default value.

Default

holdtime 150

Parameters

holdtime

Specifies the length of time, in seconds, that a neighbor should consider the sending router to be operational.

Values 5 to 255

Platforms

7705 SAR Gen 2

holdtime

Syntax

holdtime *holdtime*

no holdtime

Context

[Tree] (config>router>pim>rp>rp-candidate holdtime)

[Tree] (config>router>pim>rp>ipv6>rp-candidate holdtime)

Full Context

configure router pim rp rp-candidate holdtime

configure router pim rp ipv6 rp-candidate holdtime

Description

This command configures the length of time, in seconds, that neighbors should consider the sending router to be operationally up. A local RP cannot be configured on a logical router.

The **no** form of this command reverts to the default value.

Default

holdtime 150

Parameters

holdtime

Specifies the hold time, in seconds.

Values 5 to 255

Platforms

7705 SAR Gen 2

12.45 home-directory

home-directory

Syntax

home-directory *url-prefix* [*directory*] [*directory*/*directory* ..]

no home-directory

Context

[Tree] (config>system>security>user home-directory)

[Tree] (config>system>security>user-template home-directory)

Full Context

configure system security user home-directory

configure system security user-template home-directory

Description

This command configures the home directory of the user for file access. Files can be accessed locally by CLI **file** commands and output modifiers such as **>** (file redirect), or remotely via FTP and SCP. If the home directory does not exist, a warning message is displayed when the user logs in.

When **restricted-to-home** is configured, file access is denied unless the **home-directory** is configured and the directory is created by an administrator.

The **no** form of this command removes the configured home directory of the user. The directory must be also removed by the administrator.

Default

no home-directory

Parameters

url-prefix [directory] [directory/directory ..]

Specifies the local home directory URL prefix of the user and directory structure, up to 200 characters.

Platforms

7705 SAR Gen 2

12.46 hop

hop

Syntax

hop *hop-index ip-address* {**strict** | **loose**}

hop *hop-index* **sid-label** *sid-value*

no hop *hop-index*

Context

[Tree] (config>router>mpls>path hop)

Full Context

configure router mpls path hop

Description

This command specifies the hops that the LSP should traverse on its way to the egress router. When specified, the IP address can be the interface IP address, a loopback interface address, or the system IP address. If a loopback interface or the system IP address is specified then the LSP can choose the best available interface.

When an IPv6 hop is specified, the interface IP address must be a global unicast IPv6 address. A link-local address is not allowed and is rejected in the configuration if attempted.

Optionally, the LSP ingress and egress IP address can be included as the first and the last hop. A hop list can include the ingress interface IP address, the system IP address, and the egress IP address of any of the hops being specified.

When the **sid-label** parameter is specified, this command specifies an MPLS label value for a hop in the path of an SR-TE LSP. The label value implied by the SID is only used when the path is used by an SR-TE LSP.

The **no** form of this command deletes hop list entries for the path. All the LSPs currently using this path are affected. Additionally, all services actively using these LSPs are affected. The path must be shutdown first in order to delete the hop from the hop list. The **no hop** *hop-index* command will not result in any action except a warning message on the console indicating that the path is administratively up.

Parameters

hop-index

Specifies the hop index is used to order the hops specified. The LSP always traverses from the lowest hop index to the highest. The hop index does not need to be sequential.

Values 1 to 1024

ip-address

Specifies a loopback interface, the system or network interface IP address of the transit router. An interface IPv6 address must be a global unicast address.

Values ipv4-address — a.b.c.d

ipv6-address — x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x — 0 to FFFF (hexadecimal)

d — 0 to 255 (decimal)

loose

This keyword specifies that the route taken by the LSP from the previous hop to this hop can traverse through other routers. Multiple hop entries with the same IP address are flagged as errors. Either the **loose** *or* **strict** keyword must be specified.

strict

This keyword specifies that the LSP must take a direct path from the previous hop router to this router. No transit routers between the previous router and this router are allowed. If the IP address specified is the interface address, then that is the interface the LSP must use.

If there are direct parallel links between the previous router and this router and if system IP address is specified, then any one of the available interfaces can be used by the LSP. The user must ensure that the previous router and this router have a direct link. Multiple hop entries with the same IP address are flagged as errors. Either the **loose** *or* **strict** keyword must be specified.

sid-value

Specifies the SID value. The *sid-value* can be any valid MPLS/SR label value. It is not restricted by any locally-defined label ranges since these may be different on the remote node or adjacency for which the SID is defined.

Values 32 to 1048575

Platforms

7705 SAR Gen 2

hop

Syntax

hop *hop-index ip-address*

no hop *hop-index*

Context

[Tree] (config>service>pw-routing>path hop)

Full Context

configure service pw-routing path hop

Description

This command configures each hop on an explicit path that can be used by one or more dynamic MS-PWs. It specifies the IP addresses of the hops that the MS-PE should traverse. These IP addresses can correspond to the system IP address of each S-PE, or the IP address on which the T-LDP session to a given S-PE terminates.

The **no** form of this command deletes hop list entries for the path. All the MS-PWs currently using this path are unaffected. Additionally, all services actively using these MS-PWs are unaffected. The path must be shutdown first in order to delete the hop from the hop list. The ' **no hop hop-index**' command will not result in any action, except for a warning message on the console indicating that the path is administratively up.

Default

no hop

Parameters

hop-index

Specifies a locally significant numeric identifier for the hop. The hop index is used to order the hops specified. The LSP always traverses from the lowest hop index to the highest. The hop index does not need to be sequential.

Values 1 to 1024

ip-address

Specifies the system IP address or terminating IP address for the T-LDP session to the S-PE corresponding to this hop. For a given IP address on a hop, the system will choose the appropriate SDP to use.

Platforms

7705 SAR Gen 2

12.47 hop-by-hop-opt

hop-by-hop-opt

Syntax

hop-by-hop-opt {**true** | **false**}

no hop-by-hop-opt

Context

[Tree] (config>filter>ipv6-filter>entry>match hop-by-hop-opt)

Full Context

configure filter ipv6-filter entry match hop-by-hop-opt

Description

This command enables match on existence of Hop-by-Hop Options Extension Header in the IPv6 filter policy.

The **no** form of this command ignores Hop-by-Hop Options Extension Header presence/absence in a packet when evaluating match criteria of a given filter policy entry.

Default

no hop-by-hop-opt

Parameters

true

Matches a packet with a Hop-by-Hop Options Extension header.

false

Matches a packet without a Hop-by-Hop Options Extension header.

Platforms

7705 SAR Gen 2

12.48 hop-limit

hop-limit

Syntax

hop-limit *limit*

no hop-limit

Context

[Tree] (config>router>mpls>lsp>fast-reroute hop-limit)

[Tree] (config>router>mpls>lsp-template>fast-reroute hop-limit)

Full Context

configure router mpls lsp fast-reroute hop-limit

configure router mpls lsp-template fast-reroute hop-limit

Description

For fast reroute, how many more routers a detour is allowed to traverse compared to the LSP itself. For example, if an LSP traverses four routers, any detour for the LSP can be no more than ten router hops, including the ingress and egress routers.

The **no** form of this command reverts to the default value.

Default

hop-limit 16

Parameters

limit

Specify the maximum number of hops.

Values 0 to 255

Platforms

7705 SAR Gen 2

hop-limit

Syntax

hop-limit *number*

no hop-limit

Context

[Tree] (config>router>mpls>lsp hop-limit)

Full Context

configure router mpls lsp hop-limit

Description

This command specifies the maximum number of hops that an LSP can traverse, including the ingress and egress routers. An LSP is not set up if the hop limit is exceeded. This value can be changed dynamically for an LSP that is already set up with the following implications.

If the new value is less than the current number of hops of the established LSP, the LSP is brought down. The software then tries to re-establish the LSP within the new **hop-limit** number. If the new value is equal to or greater than the current number hops of the established LSP, the LSP is not affected.

The **no** form of this command returns the parameter to the default value.

Default

hop-limit 255

Parameters

number

Specifies the number of hops the LSP can traverse, expressed as an integer.

Values 2 to 255

Platforms

7705 SAR Gen 2

hop-limit

Syntax

hop-limit *number*

no hop-limit

Context

[Tree] (config>router>mpls>lsp>primary hop-limit)

[Tree] (config>router>mpls>lsp>secondary hop-limit)

Full Context

configure router mpls lsp primary hop-limit

configure router mpls lsp secondary hop-limit

Description

This optional command overrides the **config>router>mpls>lsp** *lsp-name***>hop-limit** command. This command specifies the total number of hops that an LSP traverses, including the ingress and egress routers.

This value can be changed dynamically for an LSP that is already set up with the following implications:

If the new value is less than the current hops of the established LSP, the LSP is brought down. MPLS then tries to re-establish the LSP within the new hop-limit number. If the new value is equal or more than the current hops of the established LSP then the LSP will be unaffected.

The **no** form of this command reverts the values defined under the LSP definition using the **config>router>mpls>lsp** *lsp-name***>hop-limit** command.

Default

no hop-limit

Parameters

number

Specifies the number of hops the LSP can traverse, expressed as an integer.

Values 2 to 255

Platforms

7705 SAR Gen 2

12.49 host

host

Syntax

host *host-name* [**create**]

no host *host-name*

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe host)

Full Context

configure subscriber-mgmt local-user-db ipoe host

Description

This command creates an IPoE or PPP host entry in the local user database. A host entry in the local user database is matched based on the specified match-list criteria and an optional mask that is applied to the host-identification parameters.

A default host entry can be created without host-identification parameters which is used when no other host entries match. Note that creating a default host entry also requires a match-list to be specified.

The **no** form of this command removes the host entry from the local user database.

Parameters

host-name

Specifies a unique host name, up to 32 characters. The *host-name* **default** creates a special match-all host entry that should not have host-identification parameters and is used when no other host entries match.

create

Keyword used to create the host name. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

host

Syntax

[**no**] **host** [*ip-address*]

[**no**] **host** [**fwd-service** *service-id*] **group-interface** *ip-int-name*

Context

[Tree] (debug>router>igmp host)

Full Context

debug router igmp host

Description

This command enables debugging for the IGMP host.

The **no** form of the command disables debugging.

Parameters

ip-address

Debugs the information associated with the specified IP address.

service-id

Debugs information associated with the service ID.

Values service-id: 1 to 2148278386

svc-name: up to 64 characters.

group-interface ***ip-int-name***

Debugs the information associated with the specified IP interface name.

Values IP interface address

Platforms

7705 SAR Gen 2

12.50 host-identification

host-identification

Syntax

host-identification

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host host-identification)

Full Context

configure subscriber-mgmt local-user-db ipoe host host-identification

Description

Commands in this context configure host identification parameters.

Platforms

7705 SAR Gen 2

12.51 host-ip

host-ip

Syntax

host-ip *prefix-list-name*

Context

[Tree] (config>router>policy-options>policy-statement>entry>from host-ip)

Full Context

configure router policy-options policy-statement entry from host-ip

Description

This command specifies a prefix list host IP address as a match criterion for the route policy-statement entry.

Default

no host-ip

Parameters

prefix-list-name

Specifies the prefix-list name. Allowed values are any string up to 64 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

The *prefix-list-name* is defined in the **config>router>policy-options>prefix-list** context.

Platforms

7705 SAR Gen 2

12.52 host-key

host-key

Syntax

host-key *index* **name** *host-key-name*

no host-key *index*

Context

[Tree] (config>system>security>ssh>client-host-key host-key)

[Tree] (config>system>security>ssh>server-host-key host-key)

Full Context

configure system security ssh client-host-key-list host-key

configure system security ssh server-host-key-list host-key

Description

This command configures a host key. Client host keys are used when the SR OS is acting as an SSH client. Server host keys are used when the SR OS is acting as an SSH server.

The **no** form of this command removes the index and host-key name from the configuration.

Default

no host-key *index*

Parameters

index

Specifies the index of the host key in the list.

Values 1 to 255

host-key

Specifies the host-key algorithm.

Values ssh-dss, ssh-rsa, rsa-sha2-256, rsa-sha2-512, ecdsa-sha2-nistp256, ecdsa-sha2-nistp384, ecdsa-sha2-nistp521, ssh-ed25519

Platforms

7705 SAR Gen 2

12.53 host-unreachable

host-unreachable

Syntax

[**no**] **host-unreachable** *ip-address*

[**no**] **host-unreachable** *ipv6-address*

Context

[Tree] (config>vrrp>policy>priority-event host-unreachable)

Full Context

configure vrrp policy priority-event host-unreachable

Description

This command creates the context to configure a host unreachable priority control event to monitor the ability to receive ICMP echo reply packets from an IP host address.

A host unreachable priority event creates a continuous ICMP echo request (ping) probe to the specified *ip-address*. If a ping fails, the event is considered to be set. If a ping is successful, the event is considered to be cleared.

Multiple unique (different *ip-address*) **host-unreachable** event nodes can be configured within the **priority-event** node to a maximum of 32 events.

The **host-unreachable** command can reference any valid local or remote IP address. The ability to ARP a local IP address or find a remote IP address within a route prefix in the route table is considered part of the monitoring procedure. The **host-unreachable** priority event operational state tracks ARP or route table entries dynamically appearing and disappearing from the system. The operational state of the **host-unreachable** event are listed in Table 42: Host Unreachable Operational States.

Table 42: Host Unreachable Operational States

Host Unreachable Operational State	Description
Set – no ARP	No ARP address found for *ip-addr* for **drop-count** consecutive attempts; only applies when IP address is considered local
Set – no route	No route exists for *ip-addr* for **drop-count** consecutive attempts; only when IP address is considered remote
Set – host unreachable	ICMP host unreachable message received for **drop-count** consecutive attempts
Set – no reply	ICMP echo request timed out for **drop-count** consecutive attempts
Set – reply received	Last ICMP echo request attempt received an echo reply but historically not able to clear the event
Cleared – no ARP	No ARP address found for *ip-addr* - not enough failed attempts to set the event
Cleared – no route	No route exists for *ip-addr* - not enough failed attempts to set the event
Cleared – host unreachable	ICMP host unreachable message received - not enough failed attempts to set the event
Cleared – no reply	ICMP echo request timed out - not enough failed attempts to set the event
Cleared – reply received	Event is cleared - last ICMP echo request received an echo reply

Unlike other priority event types, the **host-unreachable** priority event monitors a repetitive task. A historical evaluation is performed on the success rate of receiving ICMP echo reply messages. The operational state takes its cleared and set orientation from the historical success rate. The informational portion of the operational state is derived from the last attempt's result. It is possible for the previous attempt to fail while the operational state is still cleared due to an insufficient number of failures to cause it to become set. It is also possible for the state to be set while the previous attempt was successful.

When an event transitions from clear to set, the set is processed immediately and must be reflected in the associated virtual router instances in-use priority value. As the event transitions from clear to set, a hold-set timer is loaded with the value configured by the events **hold-set** command. This timer prevents the event from clearing until it expires, damping the effect of event flapping. If the event clears and becomes set again before the hold-set timer expires, the timer is reset to the **hold-set** value, extending the time before another clear can take effect.

The hold-set timer be expired and the historical success rate must be met prior to the event operational state becoming cleared.

The **no** form of the command deletes the specific IP host monitoring event. The event may be deleted at any time. When the event is deleted, the in-use priority of all associated virtual router instances must be reevaluated. The event's **hold-set** timer has no effect on the removal procedure.

Default

no host-unreachable — No host unreachable priority events are created.

Parameters

ip-address

The IP address of the host for which the specific event will monitor connectivity. The *ip-addr* can only be monitored by a single event in this policy. The IP address can be monitored by multiple VRRP priority control policies. The IP address can be used in one or multiple **ping** requests. Each VRRP priority control **host-unreachable** and **ping** destined to the same *ip-addr* is uniquely identified on a per message basis. Each session originates a unique identifier value for the ICMP echo request messages it generates. This allows received ICMP echo reply messages to be directed to the appropriate sending application.

Values

ipv4-address:	a.b.c.d	
ipv6-address:	x:x:x:x:x:x:x:x[-interface]	
	x:	[0..FFFF]H
	interface:	32 chars maximum, mandatory for link local addresses

The link-local IPv6 address must have an interface name specified. The global IPv6 address must not have an interface name specified.

Platforms

7705 SAR Gen 2

12.54 host-unsolicited-na-flood-evpn

host-unsolicited-na-flood-evpn

Syntax

[**no**] **host-unsolicited-na-flood-evpn**

Context

[Tree] (config>service>vpls>proxy-nd host-unsolicited-na-flood-evpn)

Full Context

configure service vpls proxy-nd host-unsolicited-na-flood-evpn

Description

This command controls whether the system floods host unsolicited Neighbor Advertisements to the EVPN. The NA messages impacted by this command are NA messages with the following flags: S=0 and R=0.

The **no** form of the command will only flood to local SAPs/binds but not to the EVPN destinations. This is only recommended in networks where CEs are routers that are directly connected to the PEs. Networks using aggregation switches between the host/routers and the PEs should flood unsolicited NA messages in the EVPN to ensure that the remote caches are updated and the BGP does not miss the advertisement of these entries.

Default

host-unsolicited-na-flood-evpn

Platforms

7705 SAR Gen 2

12.55 hostname

hostname

Syntax

hostname {**use-system-name** | **value** *value-string*}

no hostname

Context

[Tree] (config>log>syslog hostname)

Full Context

configure log syslog hostname

Description

This command controls how the HOSTNAME field of syslog messages is populated.

The **no** form of this command causes the HOSTNAME to be populated with an IP address.

Default

no hostname

Parameters

use-system-name

Keyword used to specify the HOSTNAME uses the system name as configured by the **configure system name** command. Do not use any spaces in the system name if it is used for the syslog HOSTNAME.

value-string

Specifies a string, up to 255 characters with no spaces, that is used as the HOSTNAME of syslog messages.

Platforms

7705 SAR Gen 2

hostname

Syntax

hostname {**use-system-name** | **use-vprn-name** | **value** *value-string*}

no hostname

Context

[Tree] (config>service>vprn>log>syslog hostname)

Full Context

configure service vprn log syslog hostname

Description

This command controls how the HOSTNAME field of syslog messages is populated.

The **no** form of this command causes the HOSTNAME to be populated with an IP address.

Default

no hostname

Parameters

use-system-name

Keyword used to specify the HOSTNAME uses the system name as configured by the **configure system name** command. Do not use any spaces in the system name if it is used for the syslog HOSTNAME.

use-vprn-name

Keyword used to specify the HOSTNAME uses the VPRN name as configured by the **configure service vprn name** command. Do not use any spaces in the VPRN name if it is used for the syslog HOSTNAME.

value-string

Specifies a string, up to 255 characters with no spaces, that is used as the HOSTNAME of syslog messages.

Platforms

7705 SAR Gen 2

12.56 hour

hour

Syntax

hour *hour-number* [*..hour-number*] | **all**}

no hour

Context

[Tree] (config>system>cron>sched hour)

Full Context

configure system cron schedule hour

Description

This command specifies which hour to schedule a command. Multiple hours of the day can be specified. When multiple hours are configured, each of them will cause the schedule to trigger. **Day-of-month** or **weekday** must also be specified. All days of the month or weekdays can be specified. If an hour is configured without configuring month, weekday, day-of-month, and minute, the event will not execute.

The **no** form of this command removes the specified hour from the configuration.

Default

no hour

Parameters

hour-number

Specifies the hour to schedule a command.

Values 0 to 23 (maximum 24 hour-numbers)

all

Specifies all hours.

Platforms

7705 SAR Gen 2

12.57 http-auth

http-auth

Syntax

http-auth password *password* [**hash | hash2**]

http-auth username *user-name*

http-auth username *user-name* **password** *password* [**hash | hash2**]

no http-auth

Context

[Tree] (config>system>security>pki>est-profile http-auth)

Full Context

configure system security pki est-profile http-auth

Description

This command configures HTTP authentication parameters. HTTP authentication is used by a client when requested by the server. When disabled, there is no HTTP-level client authentication.

The **no** form of the command reverts to the default value.

Default

no http-auth

Parameters

password

Specifies a text string containing the password. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

user-name

Specifies the name of the user to authenticate, up to 32 characters.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

Platforms

7705 SAR Gen 2

12.58 http-connections

http-connections

Syntax

http-connections [*ip-address/prefix-length*]

http-connections any

http-connections [*ipv6-address/prefix-length*]

no http-connections

Context

[Tree] (debug>system http-connections)

Full Context

debug system http-connections

Description

This command displays HTTP connections debug information.

Parameters

ip-address/prefix-length

Displays information for the specified host IP address and prefix length.

Values ip-address: a.b.c.d

prefix-length: 0 to 32

any

Specifies that any address can be used.

ipv6-address/prefix-length

Displays information for the specified host IPv6 address and prefix length.

Values ipv6-address:

- x:x:x:x:x:x:x:x: (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x [0 to FFFFF] H
- d [0 to 255] D

prefix-length: 0 to 128

Platforms

7705 SAR Gen 2

12.59 http-response-timeout

http-response-timeout

Syntax

http-response-timeout *timeout*

no http-response-timeout

Context

[Tree] (config>system>security>pki>ca-profile>cmpv2 http-response-timeout)

Full Context

configure system security pki ca-profile cmpv2 http-response-timeout

Description

This command specifies the timeout value for HTTP response that is used by CMPv2.

The **no** form of this command reverts to the default.

Default

http-response-timeout 30

Parameters

timeout

Specifies the HTTP response timeout, in seconds.

Values 1 to 3600

Platforms

7705 SAR Gen 2

http-response-timeout

Syntax

http-response-timeout *timeout*

no http-response-timeout

Context

[Tree] (config>system>security>pki>ca-profile>cmpv2 http-response-timeout)

Full Context

configure system security pki ca-profile cmpv2 http-response-timeout

Description

This command specifies the timeout value for HTTP response that is used by CMPv2.

The **no** form of this command reverts to the default.

Default

http-response-timeout 30

Parameters

timeout

Specifies the HTTP response timeout in seconds.

Values 1 to 3600

Platforms

7705 SAR Gen 2

12.60 http-version

http-version

Syntax

http-version [**1.0** | **1.1**]

Context

[Tree] (config>system>security>pki>ca-profile>cmpv2 http-version)

Full Context

configure system security pki ca-profile cmpv2 http-version

Description

This command configures the HTTP version for CMPv2 messages.

Default

http-version 1.1

Platforms

7705 SAR Gen 2

13 i Commands

13.1 i-counters

i-counters

Syntax

i-counters [**all**]

no i-counters

Context

[Tree] (config>log>acct-policy>cr>policer i-counters)

[Tree] (config>log>acct-policy>cr>queue i-counters)

[Tree] (config>log>acct-policy>cr>ref-queue i-counters)

[Tree] (config>log>acct-policy>cr>ref-policer i-counters)

Full Context

configure log accounting-policy custom-record policer i-counters

configure log accounting-policy custom-record queue i-counters

configure log accounting-policy custom-record ref-queue i-counters

configure log accounting-policy custom-record ref-policer i-counters

Description

This command configures ingress counter parameters for this custom record.

The **no** form of this command reverts all ingress counters to their default value.

Default

i-counters

Parameters

all

Specifies that all ingress counters should be included.

Platforms

7705 SAR Gen 2

13.2 ibgp-multipath

ibgp-multipath

Syntax

[**no**] **ibgp-multipath**

Context

[Tree] (config>service>vprn>bgp ibgp-multipath)

Full Context

configure service vprn bgp ibgp-multipath

Description

This command defines the type of IBGP multipath to use when adding BGP routes to the route table if the route resolving the BGP nexthop offers multiple next-hops.

The **no** form of this command disables the IBGP multipath load balancing feature.

Platforms

7705 SAR Gen 2

ibgp-multipath

Syntax

[**no**] **ibgp-multipath**

Context

[Tree] (config>router>bgp ibgp-multipath)

Full Context

configure router bgp ibgp-multipath

Description

This command enables IBGP multipath load balancing when adding BGP routes to the route table if the route resolving the BGP nexthop offers multiple next-hops.

The **no** form of this command disables the IBGP multipath load balancing feature.

Default

no ibgp-multipath

Platforms

7705 SAR Gen 2

13.3 icmp

icmp

Syntax

icmp

Context

[Tree] (config>service>vprn>nw-if icmp)

[Tree] (config>service>vprn>if icmp)

[Tree] (config>service>ies>if icmp)

Full Context

configure service vprn network-interface icmp

configure service vprn interface icmp

configure service ies interface icmp

Description

Commands in this context configure Internet Control Message Protocol (ICMP) parameters on a service.

Platforms

7705 SAR Gen 2

icmp

Syntax

icmp

Context

[Tree] (config>router>if icmp)

Full Context

configure router interface icmp

Description

This command enables access to the context to configure Internet Control Message Protocol (ICMP) parameters on a network IP interface. ICMP is a message control and error reporting protocol that also provides information relevant to IP packet processing.

Platforms

7705 SAR Gen 2

icmp

Syntax

[**no**] **icmp**

Context

[Tree] (debug>router>ip icmp)

Full Context

debug router ip icmp

Description

This command enables ICMP debugging.

Platforms

7705 SAR Gen 2

icmp

Syntax

icmp

Context

[Tree] (config>test-oam icmp)

Full Context

configure test-oam icmp

Description

Commands in this context configure test ICMP OAM parameters.

Platforms

7705 SAR Gen 2

13.4 icmp-code

icmp-code

Syntax

icmp-code *icmp-code*

no icmp-code

Context

[Tree] (config>filter>ip-exception>entry>match icmp-code)

[Tree] (config>filter>ipv6-filter>entry>match icmp-code)

[Tree] (config>filter>ipv6-exception>entry>match icmp-code)

[Tree] (config>filter>ip-filter>entry>match icmp-code)

Full Context

configure filter ip-exception entry match icmp-code

configure filter ipv6-filter entry match icmp-code

configure filter ipv6-exception entry match icmp-code

configure filter ip-filter entry match icmp-code

Description

Configures matching on /ICMPv6 code field in the /ICMPv6 header of an IPv4 or IPv6 packet as a filter match criterion or configures matching on the ICMP code field in the ICMP header of an IPv4 packet as an exception filter match criterion. An entry containing Layer 4 non-zero match criteria will not match non-initial (for example, 2nd, 3rd) fragments of a fragmented packet because only the first fragment contains the Layer 4 information. Similarly an entry containing " **icmp-code** 0" match criterion, may match non-initial fragments when the Layer 4 header is not present in a packet fragment and other match criteria are also met.

The **no** form of the command removes the criterion from the match entry.

Default

no icmp-code

Parameters

icmp-code

Specifies the /ICMPv6 code value that must be present to match. Value can be expressed as a decimal integer, as well as in hexadecimal or binary format, or even using keywords. The following value shows decimal integer only.

Values 0 to 255

Platforms

7705 SAR Gen 2

13.5 icmp-echo-reply

icmp-echo-reply

Syntax

[**no**] **icmp-echo-reply**

Context

[Tree] (config>service>vprn>nat>outside>pool icmp-echo-reply)

[Tree] (config>router>nat>outside>pool icmp-echo-reply)

Full Context

configure service vprn nat outside pool icmp-echo-reply

configure router nat outside pool icmp-echo-reply

Description

IPv4 addresses in a NAT pool can be configured to respond to ICMP Echo Requests (PINGs). The configuration can be toggled online while the pool is in use.

In L2-aware NAT when **port-block-extensions** is disabled, the reply from an outside IP address is generated only when this IP address has at least one host (binding) behind it.

In L2-aware NAT when **port-block-extensions** is enabled, the reply from an outside IP address is generated regardless if a binding is present.

In LSN, the reply from an outside IP address is generated regardless if a binding is present.

The **no** form of the command disables ICMP echo replies.

Default

no icmp-echo-reply

Platforms

7705 SAR Gen 2

13.6 icmp-generation

icmp-generation

Syntax

icmp-generation

Context

[Tree] (config>service>vprn>if>sap>ipsec-tunnel icmp-generation)

[Tree] (config>router>if>ipsec>ipsec-tunnel icmp-generation)

[Tree] (config>service>vprn>if>sap>ip-tunnel icmp-generation)

[Tree] (config>service>ies>if>sap>ip-tunnel icmp-generation)

[Tree] (config>ipsec>tunnel-template icmp-generation)

Full Context

configure service vprn interface sap ipsec-tunnel icmp-generation

configure router interface ipsec ipsec-tunnel icmp-generation

configure service vprn interface sap ip-tunnel icmp-generation

configure service ies interface sap ip-tunnel icmp-generation

configure ipsec tunnel-template icmp-generation

Description

This command enables the context to configure ICMP generation information.

Platforms

7705 SAR Gen 2

13.7 icmp-ping

icmp-ping

Syntax

icmp-ping {*ip-address* | *dns-name*} [{**bypass-routing** | {**interface** *interface-name*} | {**next-hop** *ip-address*}}] [**count** *requests*] [**do-not-fragment**] [**fc** *fc-name*] [**interval** { *centisecs* | *secs*}] [**pattern** *pattern*] [**rapid**] [{ **router** *router-or-service* | **router-instance** *router-instance* | **service-name** *service-name*}] [**size** *bytes*] [**source** *ip-address*] [**timeout** *timeout*] [**tos** *type-of-service*] [**ttl** *time-to-live*]

Context

[Tree] (config>saa>test>type icmp-ping)

Full Context

configure saa test type icmp-ping

Description

This command configures an ICMP traceroute test.

Parameters

ip-address* | *dns-name

Specifies the far-end IP address or DNS name to which to send the **svc-ping** request message in dotted-decimal notation.

Values

ipv4-address:	a.b.c.d	
ipv6-address:	x:x:x:x:x:x:x:x	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 to FFFF]H
	d:	[0 to 255]D
	interface	up to 32 characters. This is mandatory for link local addresses.
dns-name	up to 128 characters	

bypass-routing

Specifies whether to send the ping request to a host on a directly attached network bypassing the routing table.

interface-name

Specifies the name used to refer to the interface, up to 32 characters. The name must already exist in the **config>router>interface** context.

next-hop *ip-address*

Displays only static routes with the specified next-hop IP address.

Values

ipv4-address:	a.b.c.d (host bits must be 0)
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

requests

Specifies the number of times to perform an OAM ping probe operation. Each OAM echo message request must either time out or receive a reply before the next message request is sent.

Values 1 to 100000

Default 5

do-not-fragment

Sets the DF (Do Not Fragment) bit in the ICMP ping packet (does not apply to ICMPv6).

fc-name

Specifies the forwarding class of the SAA.

Values be, l2, af, l1, h2, ef, h1, nc

Default nc

interval {*centisecs* | *secs*}

Specifies the minimum amount of time, in seconds, that must expire before the next message request is sent. If the **rapid** parameter is configured, this value is measured in centiseconds (hundredths of a second) instead of seconds.

Values 1 to 10000

Default 1

pattern

Specifies the date portion in a ping packet is filled with the pattern value specified. If not specified, a system-generated sequential pattern is used.

Values 0 to 65535

rapid

Configures the *interval* parameter to use centiseconds (hundredths of a second) instead of seconds.

router-or-service

Specifies the numerical reference to the router instance or service. Well known router names "Base", "management", "vpm-vr-name", and " vpls-management" are allowed for convenience, but are mapped numerically.

Values {*router-name* | *vprn-svc-id*}

router-name: Base, management, cmp-vr-name, vpls-management

vprn-svc-id: 1 to 2147483647

cpm-vr-name: Up to 32 characters

The parameter *router-instance* is preferred for specifying the router or service.

Default Base

router-instance

Specifies the preferred method for entering a service name. Stored as the service name. Only the service linking function is allowed for both mixed-mode and model-driven configuration modes.

Values *router-name*, *vprn-svc-name*

router-name: Base, management, vpls-management, *cpm-vr-name*

vprn-svc-name: up to 64 characters

cpm-vr-name: up to 32 characters

service-name

Specifies the alias function that allows the service name to be used, converted and stored as a service ID, up to 64 characters.

The *router-instance* parameter is preferred for specifying the router or service.

bytes

Specifies the request packet size in bytes, expressed as a decimal integer.

Values 0 to 16384

Default 56

source *ip-address*

Specifies the IP address to be used.

Values

ipv4-address:	a.b.c.d	
ipv6-address:	x:x:x:x:x:x:x:x	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 to FFFF]H
	d:	[0 to 255]D

timeout

Specifies the override time that the router waits for a message reply after sending the last probe for a specific test. Upon the expiration of the time out, the test is marked complete and no more packets are processed for any of those request probes.

Values 1 to 10

Default 5

type-of-service

Specifies the service type.

Values 0 to 255

Default 0

time-to-live

Specifies the TTL value for the MPLS label, expressed as a decimal integer.

Values 1 to 128

Default 64

Platforms

7705 SAR Gen 2

13.8 icmp-query

icmp-query

Syntax

icmp-query [**min** *minutes*] [**sec** *seconds*]

no icmp-query

Context

[Tree] (config>service>nat>nat-policy>timeouts icmp-query)

Full Context

configure service nat nat-policy timeouts icmp-query

Description

This command configures the timeout applied to an ICMP query session.

Default

icmp-query min 1

Parameters

min *minutes*

Specifies the timeout, in minutes, applied to an ICMP query session.

Values 1 to 4

Default 1

sec *seconds*

Specifies the timeout, in seconds, applied to an ICMP query session.

Values 1 to 59

Platforms

7705 SAR Gen 2

13.9 icmp-tunneling

icmp-tunneling

Syntax

[**no**] **icmp-tunneling**

Context

[Tree] (config>router icmp-tunneling)

Full Context

configure router icmp-tunneling

Description

This command enables the tunneling of ICMP reply packets over MPLS LSP at a LSR node as per RFC 3032.

The LSR part of this feature consists of crafting the reply ICMP packet of type=11- 'time exceeded', with a source address set to a local address of the LSR node, and appending the IP header and leading payload octets of the original datagram. The system skips the lookup of the source address of the sender of the label TTL expiry packet, which becomes the destination address of the ICMP reply packet. Instead, CPM injects the ICMP reply packet in the forward direction of the MPLS LSP the label TTL expiry packet was received from. The TTL of pushed labels should be set to 255.

The source address of the ICMP reply packet is determined as follows. The LSR uses the address of the outgoing interface for the MPLS LSP. With LDP LSP or BGP LSP multiple ECMP next-hops can exist and in such a case the first outgoing interface is selected. If that interface does not have an address of the same family (IPv4 or IPv6) as the ICMP packet, then the system address of the same family is selected. If one is not configured, the packet is dropped.

When the packet is received by the egress LER, it performs a regular user packet lookup in the data path in the GRT context for BGP shortcut, 6PE, and BGP label route prefixes, or in VPRN context for VPRN and 6VPE prefixes. It then forwards it to the destination, which is the sender of the original packet which TTL expired at the LSR.

If the egress LER does not have a route to the destination of the ICMP packet, it drops the packets.

The rate of the tunneled ICMP replies at the LSR can be directly or indirectly controlled by the existing IOM level and CPM levels mechanisms. Specifically, the rate of the incoming UDP traceroute packets received with a label stack can be controlled at ingress IOM using the distributed CPU protection feature. The rate of the ICMP replies by CPM can also be directly controlled by configuring a system wide rate limit for packets ICMP replies to MPLS expired packets which are successfully forwarded to CPM using

the command 'configure system security vprn-network-exceptions'. While this command's name refers to VPRN service, this feature rate limits ICMP replies for packets received with any label stack, including VPRN and shortcuts.

The 7705 SAR Gen 2 implementation supports appending to the ICMP reply of type Time Exceeded the MPLS label stack object defined in RFC 4950. It does not include it in the ICMP reply type of Destination unreachable.

The new MPLS Label Stack object permits an LSR to include label stack information including label value, EXP, and TTL field values, from the encapsulation header of the packet that expired at the LSR node. The ICMP message continues to include the IP header and leading payload octets of the original datagram.

In order to include the MPLS Label Stack object, SR OS implementation adds support of RFC 4884 which defines extensions for a multi-part ICMPv4/v6 message of type Time Exceeded.

The **no** form of command disables the tunneling of ICMP reply packets over MPLS LSP at a LSR node.

Default

no icmp-tunneling

Platforms

7705 SAR Gen 2

13.10 icmp-type

icmp-type

Syntax

icmp-type *icmp-type*

no icmp-type

Context

[Tree] (config>filter>ip-exception>entry>match icmp-type)

[Tree] (config>filter>ipv6-filter>entry>match icmp-type)

[Tree] (config>filter>ipv6-exception>entry>match icmp-type)

[Tree] (config>filter>ip-filter>entry>match icmp-type)

Full Context

configure filter ip-exception entry match icmp-type

configure filter ipv6-filter entry match icmp-type

configure filter ipv6-exception entry match icmp-type

configure filter ip-filter entry match icmp-type

Description

This command configures matching on the /ICMPv6 type field in the /ICMPv6 header of an IPv4 or IPv6 packet as a filter match criterion or configures matching on the ICMP type field in the ICMP header of an IPv4 packet as an exception filter match criterion. An entry containing Layer 4 non-zero match criteria will not match non-initial (2nd, 3rd, etc.) fragments of a fragmented packet because only the first fragment contains the Layer 4 information. Similarly an entry containing " **icmp-type** 0" match criterion, may match non-initial fragments when the Layer 4 header is not present in a packet fragment and other match criteria are also met.

The **no** form of the command removes the criterion from the match entry.

Default

no icmp-type

Parameters

icmp-type

Specifies the /ICMPv6 type value that must be present to match. Value can be expressed as a decimal integer, as well as in hexadecimal or binary format, or even using keywords. The following value shows decimal integer only.

Values 0 to 255

Platforms

7705 SAR Gen 2

icmp-type

Syntax

icmp-type *icmp-type*

no icmp-type

Context

[Tree] (config>qos>network>egress>ipv6-criteria>entry>match icmp-type)

[Tree] (config>qos>network>egress>ip-criteria>entry>match icmp-type)

Full Context

configure qos network egress ipv6-criteria entry match icmp-type

configure qos network egress ip-criteria entry match icmp-type

Description

This command configures matching on the ICMP or ICMPv6 type field in the ICMP or ICMPv6 header of an IPv4 or IPv6 packet as a network QoS match criterion.

An entry containing Layer 4 non-zero match criteria will not match non-initial (2nd, 3rd, etc.) fragments of a fragmented packet since only the first fragment contains the Layer 4 information. Similarly, an entry

containing " **icmp-type** 0" match criterion, may match non-initial fragments when the Layer 4 header is not present in a packet fragment and other match criteria are also met.

The **no** form of the command removes the criterion from the match entry.

Default

no icmp-type

Parameters

icmp-type

Specifies the ICMP or ICMPv6 type value that must be present to match. Value can be expressed as a decimal integer, or in hexadecimal or binary format, or even using keywords.

Values 0 to 255 (Decimal)

0 to FF (Hexadecimal)

0 to 11111111 (Binary)

Platforms

7705 SAR Gen 2

13.11 icmp6

icmp6

Syntax

icmp6

Context

[Tree] (config>service>ies>if>ipv6 icmp6)

[Tree] (config>service>vprn>if>ipv6 icmp6)

Full Context

configure service ies interface ipv6 icmp6

configure service vprn interface ipv6 icmp6

Description

This command configures ICMPv6 parameters for the interface.

Platforms

7705 SAR Gen 2

icmp6

Syntax

icmp6

Context

[Tree] (config>router>if>ipv6 icmp6)

Full Context

configure router interface ipv6 icmp6

Description

Commands in this context configure ICMPv6 parameters for the interface.

Platforms

7705 SAR Gen 2

icmp6

Syntax

icmp6 [*ip-int-name*]

no icmp6

Context

[Tree] (debug>router>ip icmp6)

Full Context

debug router ip icmp6

Description

This command enables ICMPv6 debugging.

Platforms

7705 SAR Gen 2

13.12 icmp6-generation

icmp6-generation

Syntax

icmp6-generation

Context

[Tree] (config>service>vprn>if>sap>ipsec-tun icmp6-generation)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel icmp6-generation)

[Tree] (config>service>vprn>if>sap>ip-tunnel icmp6-generation)

[Tree] (config>ipsec>tnl-temp icmp6-generation)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel icmp6-generation)

[Tree] (config>router>if>ipsec>ipsec-tunnel icmp6-generation)

[Tree] (config>service>ies>if>sap>ip-tunnel icmp6-generation)

Full Context

configure service vprn interface sap ipsec-tunnel icmp6-generation

configure service vprn interface ipsec ipsec-tunnel icmp6-generation

configure service vprn interface sap ip-tunnel icmp6-generation

configure ipsec tunnel-template icmp6-generation

configure service ies interface ipsec ipsec-tunnel icmp6-generation

configure router interface ipsec ipsec-tunnel icmp6-generation

configure service ies interface sap ip-tunnel icmp6-generation

Description

This command enables the ICMPv6 packet generation configuration context.

Platforms

7705 SAR Gen 2

13.13 id

id

Syntax

[**no**] **id** *service-id*

Context

[Tree] (debug>service id)

Full Context

debug service id

Description

This command enables debugging for the specified service ID.

The **no** form of this command disables the debugging.

Parameters

service-id

The ID that uniquely identifies a service.

Values service-id: 1 to 214748364

svc-name: A string up to 64 characters in length

Platforms

7705 SAR Gen 2

13.14 idi

idi

Syntax

idi any

idi ipv4-prefix {**any** | *ipv4-prefix/ipv4-prefix-length*}

idi ipv6-prefix {**any** | *ipv6-prefix/ipv6-prefix-length*}

idi string-type *string-type* **string-value** *string-value*

no idi

Context

[Tree] (config>ipsec>client-db>client>client-id idi)

Full Context

configure ipsec client-db client client-identification idi

Description

This command specifies a match criteria that uses the peer's identification initiator (IDi) as the input, only one IDi criteria can be configured for a given client entry. This command supports the following matching methods:

- **idi any**: Matches any type of IDi with any value.
- **idi ipv4-prefix**: Matches an IDi with the type ID_IPV4_ADDR. If the **any** parameter is specified, then it will match any IPv4 address. If an IPv4 prefix is specified, then it will match an IPv4 address that is within the specified prefix.
- **idi ipv6-prefix**: Matches an IDi with the type ID_IPV6_ADDR. If the **any** parameter is specified, then it will match any IPv6 address. If an IPv6 prefix is specified, then it will match an IPv6 address that is within the specified prefix.
- **idi string-type**: Supports following type of IDi:
  - FQDN: Either a full match or a suffix match
  - RFC822: Either a full match or a suffix match

The **no** form of this command reverts to the default.

Default

no idi

Parameters

any

Matches any type of IDi with any value.

ipv4-prefix/ipv4-prefix-length

Matches any IPv4 address and prefix.

ipv6-prefix/ipv6-prefix-length

Matches any IPv6 address and prefix.

string-type

Matches the type of IDi value for this IPsec client entry.

Values fqdn, fqdn-suffix, rfc822, rfc822-suffix

string-value

Matches the IDi value within the client ID for this IPsec client entry up to 256 characters.

Platforms

7705 SAR Gen 2

idi

Syntax

[**no**] **idi**

Context

[Tree] (config>ipsec>client-db>match-list idi)

Full Context

configure ipsec client-db match-list idi

Description

This command enables the Identification Initiator (IDi) type in the IPsec client matching process.

The **no** form of this command disables the IDi matching process.

Default

no idi

Platforms

7705 SAR Gen 2

13.15 idle-time

idle-time

Syntax

idle-time *idle*

no idle-time

Context

[Tree] (config>system>grpc-tunnel>destination-group>tcp-keepalive idle-time)

[Tree] (config>system>grpc>tcp-keepalive idle-time)

[Tree] (config>system>telemetry>destination-group>tcp-keepalive idle-time)

Full Context

configure system grpc-tunnel destination-group tcp-keepalive idle-time

configure system grpc tcp-keepalive idle-time

configure system telemetry destination-group tcp-keepalive idle-time

Description

This command configures the amount of time, in seconds, that the connection must remain idle before TCP keepalive probes are sent.

The **no** form of this command reverts to the default value.

Default

idle-time 600

Parameters

idle

Specifies the number of seconds before the first TCP keepalive probe is sent.

Values 1 to 100000

Default 600

Platforms

7705 SAR Gen 2

13.16 idle-timeout

idle-timeout

Syntax

idle-timeout {*minutes* | **disable**}

no idle-timeout

Context

[Tree] (config>system>login-control idle-timeout)

Full Context

configure system login-control idle-timeout

Description

This command configures the idle timeout for console, FTP, Telnet, and SSH sessions before the session is terminated by the system.

By default, each idle console, FTP, Telnet, and SSH session times out after 30 minutes of inactivity.

The **no** form of this command reverts to the default value.

Default

idle-timeout 30

Parameters

minutes

Specifies the idle timeout in minutes. Allowed values are 1 to 1440.

Values 1 to 1440

disable

When the **disable** option is specified, a session will never timeout. To re-enable idle timeout, enter the command without the disable option.

Platforms

7705 SAR Gen 2

13.17 ies

ies

Syntax

ies *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**name** *name*]

no ies *service-id*

Context

[Tree] (config>service ies)

Full Context

configure service ies

Description

This command creates or edits an IES service instance.

The **ies** command creates or maintains an Internet Ethernet Service (IES). If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

IES services allow the creation of customer facing IP interfaces in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber must be unique between it and other addressing schemes used by the provider and potentially the entire Internet.

IP interfaces defined within the context of an IES service ID must have a SAP created as the access point to the subscriber network. This allows a combination of bridging and IP routing for redundancy purposes.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. Once a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

Once a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified results in an error.

Multiple IES services are created to separate customer owned IP interfaces. More than one IES service may be created for a single customer ID. More than one IP interface may be created within a single IES service ID. All IP interfaces created within an IES service ID belongs to the same customer.

By default, no IES service instances exist until they are explicitly created.

The **no** form of this command deletes the IES service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces defined within the service ID have been shut down and deleted.

Parameters

service-id

Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values *service-id*: 1 to 214748364

svc-name: A string up to 64 characters

customer-id

Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn-id

Specifies the VPN ID number used to identify virtual private networks (VPNs) by a VPN identification number.

Values 1 to 2147483647

Default **null** (0)

create

Keyword used to create the service ID. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

name

This parameter configures an optional service name, up to 64 characters, which adds a name identifier to a given service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider or administrator to identify and manage services within the SR OS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a given service once it is initially created.

If a name is not specified at creation time, then SR OS assigns a string version of the service-id as the name.

Service names may not begin with an integer (0 to 9).

Values *name*: up to 64 characters

Platforms

7705 SAR Gen 2

13.18 if-attribute

if-attribute

Syntax

if-attribute

Context

[Tree] (config>service>ies>interface if-attribute)

[Tree] (config>router if-attribute)

[Tree] (config>service>vprn>interface if-attribute)

[Tree] (config>router>interface if-attribute)

Full Context

configure service ies interface if-attribute

configure router if-attribute

configure service vprn interface if-attribute

configure router interface if-attribute

Description

This command creates the context to configure or apply IP interface attributes such as administrative group (admin-group) or Shared Risk Loss Group (SRLG).

Platforms

7705 SAR Gen 2

13.19 igmp

igmp

Syntax

[**no**] **igmp**

Context

[Tree] (config>service>vprn igmp)

Full Context

configure service vprn igmp

Description

Commands in this context configure IGMP parameters.

The **no** form of this command disables IGMP.

Default

no igmp

Platforms

7705 SAR Gen 2

igmp

Syntax

[**no**] **igmp**

Context

[Tree] (config>router igmp)

Full Context

configure router igmp

Description

This command enables the Internet Group Management Protocol (IGMP) context. When the context is created, the IGMP protocol is enabled.

The Internet Group Management Protocol (IGMP) is used by IPv4 systems (hosts and routers) to report their IP multicast group memberships to neighboring multicast routers. An IP multicast router can be a member of one or more multicast groups, in which case it performs both the "multicast router part" of

the protocol which collects the membership information needed by its multicast routing protocol, and the "group member part" of the protocol which informs itself and other neighboring multicast routers of its memberships.

The **no** form of the command disables the IGMP instance. To start or suspend execution of IGMP without affecting the configuration, use the **no shutdown** command.

Platforms

7705 SAR Gen 2

13.20 igmp-snooping

igmp-snooping

Syntax

igmp-snooping

Context

[Tree] (config>service>vpls>mesh-sdp igmp-snooping)

[Tree] (config>service>vpls>spoke-sdp igmp-snooping)

[Tree] (config>service>vpls>sap igmp-snooping)

[Tree] (config>service>vpls igmp-snooping)

Full Context

configure service vpls mesh-sdp igmp-snooping

configure service vpls spoke-sdp igmp-snooping

configure service vpls sap igmp-snooping

configure service vpls igmp-snooping

Description

This command enables the Internet Group Management Protocol (IGMP) snooping context.

Platforms

7705 SAR Gen 2

igmp-snooping

Syntax

[**no**] **igmp-snooping**

Context

[Tree] (debug>service>id igmp-snooping)

Full Context

debug service id igmp-snooping

Description

This command enables and configures IGMP-snooping debugging.

Platforms

7705 SAR Gen 2

igmp-snooping

Syntax

igmp-snooping

Context

[Tree] (config>service>pw-template igmp-snooping)

Full Context

configure service pw-template igmp-snooping

Description

This command enables the Internet Group Management Protocol (IGMP) snooping context.

Platforms

7705 SAR Gen 2

13.21 ignore-attached-bit

ignore-attached-bit

Syntax

ignore-attached-bit

no ignore-attached-bit

Context

[Tree] (config>service>vprn>isis ignore-attached-bit)

Full Context

configure service vprn isis ignore-attached-bit

Description

This command configures IS-IS to ignore the attached bit on received Level 1 LSPs to disable installation of default routes.

Platforms

7705 SAR Gen 2

ignore-attached-bit

Syntax

ignore-attached-bit

[**no**] **ignore-attached-bit**

Context

[Tree] (config>router>isis ignore-attached-bit)

Full Context

configure router isis ignore-attached-bit

Description

This command configures IS-IS to ignore the attached bit on received Level 1 LSPs to disable installation of default routes.

Platforms

7705 SAR Gen 2

13.22 ignore-default

ignore-default

Syntax

[**no**] **ignore-default**

Context

[Tree] (config>router>if>ipv6>urpf-check ignore-default)

[Tree] (config>router>if>urpf-check ignore-default)

Full Context

configure router interface ipv6 urpf-check ignore-default

configure router interface urpf-check ignore-default

Description

This command configures the uRPF check (if enabled) to ignore default routes for purposes of determining the validity of incoming packets. By default, default routes are considered eligible.

Platforms

7705 SAR Gen 2

13.23 ignore-dn-bit

ignore-dn-bit

Syntax

[**no**] **ignore-dn-bit**

Context

[Tree] (config>service>vprn>ospf ignore-dn-bit)

[Tree] (config>service>vprn>ospf3 ignore-dn-bit)

Full Context

configure service vprn ospf ignore-dn-bit

configure service vprn ospf3 ignore-dn-bit

Description

This command specifies whether to ignore the DN bit for OSPF LSA packets for this instance of OSPF on the router. When enabled, the DN bit for OSPF LSA packets are ignored.

The **no** form of this command does not ignore the DN bit for OSPF LSA packets.

Default

no ignore-dn-bit

Platforms

7705 SAR Gen 2

13.24 ignore-l2vpn-mtu-mismatch

ignore-l2vpn-mtu-mismatch

Syntax

ignore-l2vpn-mtu-mismatch

no ignore-l2vpn-mtu-mismatch

Context

[Tree] (config>service>epipe ignore-l2vpn-mtu-mismatch)

Full Context

configure service epipe ignore-l2vpn-mtu-mismatch

Description

This command enables the router to bring up a BGP-VPWS service regardless of any MTU mismatch. The router does not check the value of the Layer 2 MTU in the Layer2 Info Extended Community received in a BGP update message against the local service MTU or locally signaled MTU.

The **no** form of this command disables the functionality. When this command is disabled, the router does not bring up a BGP-VPWS service if an MTU mismatch occurs.

Default

no ignore-l2vpn-mtu-mismatch

Platforms

7705 SAR Gen 2

ignore-l2vpn-mtu-mismatch

Syntax

ignore-l2vpn-mtu-mismatch

no ignore-l2vpn-mtu-mismatch

Context

[Tree] (config>service>vpls ignore-l2vpn-mtu-mismatch)

Full Context

configure service vpls ignore-l2vpn-mtu-mismatch

Description

This command enables the router to bring up a VPLS service, regardless of any MTU mismatch. The router does not check the value of the Layer 2 MTU in the Layer 2 Info Extended Community received in a BGP update message or the value of the MTU interface parameter received in a LDP label mapping message against the local service MTU or locally signaled MTU.

The **no** form of this command disables the functionality. When this functionality is disabled, the router does not bring up a VPLS service if an MTU mismatch occurs.

Default

no ignore-l2vpn-mtu-mismatch

Platforms

7705 SAR Gen 2

13.25 ignore-lsp-errors

ignore-lsp-errors

Syntax

[**no**] **ignore-lsp-errors**

Context

[Tree] (config>router>isis ignore-lsp-errors)

[Tree] (config>service>vprn>isis ignore-lsp-errors)

Full Context

configure router isis ignore-lsp-errors

configure service vprn isis ignore-lsp-errors

Description

This command specifies that for this VPRN instance, ISIS will ignore LSP packets with errors. When enabled, IS-IS LSP errors will be ignored and the associated record will not be purged.

This command enables ISIS to ignore the ATT bit and therefore suppress the installation of default routes.

The **no** form of this command specifies that ISIS will not ignore LSP errors.

Platforms

7705 SAR Gen 2

13.26 ignore-match

ignore-match

Syntax

ignore-match

Context

[Tree] (config>filter>ipv6-filter>entry>action ignore-match)

[Tree] (config>filter>ip-filter>entry>action ignore-match)

Full Context

configure filter ipv6-filter entry action ignore-match

configure filter ip-filter entry action ignore-match

Description

This command sets the filter entry action to **ignore-match**, as a result this filter entry is ignored and not programmed in hardware.

Platforms

7705 SAR Gen 2

13.27 ignore-mclt-on-takeover

ignore-mclt-on-takeover

Syntax

[**no**] **ignore-mclt-on-takeover**

Context

[Tree] (config>service>vprn>dhcp6>server>pool>failover ignore-mclt-on-takeover)

[Tree] (config>router>dhcp>server>failover ignore-mclt-on-takeover)

[Tree] (config>service>vprn>dhcp>server>pool>failover ignore-mclt-on-takeover)

[Tree] (config>router>dhcp6>server>pool>failover ignore-mclt-on-takeover)

[Tree] (config>router>dhcp>server>pool>failover ignore-mclt-on-takeover)

[Tree] (config>router>dhcp6>server>failover ignore-mclt-on-takeover)

[Tree] (config>service>vprn>dhcp>server>failover ignore-mclt-on-takeover)

[Tree] (config>service>vprn>dhcp6>server>failover ignore-mclt-on-takeover)

Full Context

configure service vprn dhcp6 local-dhcp-server pool failover ignore-mclt-on-takeover

configure router dhcp local-dhcp-server failover ignore-mclt-on-takeover

configure service vprn dhcp local-dhcp-server pool failover ignore-mclt-on-takeover

configure router dhcp6 local-dhcp-server pool failover ignore-mclt-on-takeover

configure router dhcp local-dhcp-server pool failover ignore-mclt-on-takeover

configure router dhcp6 local-dhcp-server failover ignore-mclt-on-takeover

configure service vprn dhcp local-dhcp-server failover ignore-mclt-on-takeover

configure service vprn dhcp6 local-dhcp-server failover ignore-mclt-on-takeover

Description

With this flag enabled, the remote IP address or prefix can be taken over immediately upon entering the PARTNER-DOWN state of the intercommunication link, without having to wait for the Maximum Client Lead Time (MCLT) to expire. By setting this flag, the lease times of the existing DHCP clients, while the intercommunication link is in the PARTNER-DOWN state, will still be reduced to the MCLT over time and all new lease times are set to MCLT. This behavior remains the same as originally intended for MCLT.

Some deployments require that the remote IP address/prefix range starts delegating new IP addresses and prefixes upon the failure of the intercommunication link, without waiting for the intercommunication link to transition from the COMM-INT state into the PARTNER-DOWN state and the MCLT to expire while in PARTNER-DOWN state.

This can be achieved by enabling the **ignore-mclt-on-takeover** flag and by configuring the **partner-down-delay** to 0.

Enabling this functionality must be exercised with caution. One needs to keep in mind that the partner-down-delay and MCLT timers were originally introduced to prevent IP address duplication in cases where DHCP redundant nodes transition out-of-sync due to the failure of intercommunication link. These timers (**partner-down-delay** and MCLT) would ensure that during their duration, the new IP addresses and prefixes are delegated only from one node, the one with local IP address-range/prefix. This causes the new IP address delegation to be delayed and the service is impacted.

If it can be assured that the intercommunication link is always available, then the DHCP nodes would stay in sync and the two timers would not be needed. Therefore, it is important that in this mode of operation, the intercommunication link is well protected by providing multiple paths between the two DHCP nodes. The only event that should cause intercommunication link to fail is the entire nodal failure. This failure is acceptable since in this case only one DHCP node is available to provide new IP addresses and prefixes.

The **no** form of this command reverts to the default.

Platforms

7705 SAR Gen 2

13.28 ignore-mtu-mismatch

ignore-mtu-mismatch

Syntax

[**no**] **ignore-mtu-mismatch**

Context

[Tree] (config>service>vpls>bgp-evpn ignore-mtu-mismatch)

Full Context

configure service vpls bgp-evpn ignore-mtu-mismatch

Description

This command enables the system to ignore the received Layer 2 MTU in the L2 attributes extended community of the IMET route for a peer.

The **no** form of this command configures the system to compare the local service MTU against the received Layer 2 MTU and if there is a mismatch, keep the EVPN destination to the peer with operational state down.

Default

no ignore-mtu-mismatch

Platforms

7705 SAR Gen 2

13.29 ignore-narrow-metric

ignore-narrow-metric

Syntax

[**no**] **ignore-narrow-metric**

Context

[Tree] (config>service>vprn>isis ignore-narrow-metric)

Full Context

configure service vprn isis ignore-narrow-metric

Description

This command specifies that IS-IS ignores links with narrow metrics when wide-metrics support has been enabled.

The **no** form of this command specifies that IS-IS does not ignore these links.

Platforms

7705 SAR Gen 2

ignore-narrow-metric

Syntax

[**no**] **ignore-narrow-metric**

Context

[Tree] (config>router>isis ignore-narrow-metric)

Full Context

configure router isis ignore-narrow-metric

Description

This command specifies that IS-IS will ignore links with narrow metrics when wide-metrics support has been enabled.

The **no** form of this command specifies that IS-IS will not ignore these links.

Platforms

7705 SAR Gen 2

13.30 ignore-nh-metric

ignore-nh-metric

Syntax

[**no**] **ignore-nh-metric**

Context

[Tree] (config>service>vprn ignore-nh-metric)

[Tree] (config>service>vprn>bgp>best-path-selection ignore-nh-metric)

[Tree] (config>router>bgp>best-path-selection ignore-nh-metric)

Full Context

configure service vprn ignore-nh-metric

configure service vprn bgp best-path-selection ignore-nh-metric

configure router bgp best-path-selection ignore-nh-metric

Description

This command instructs BGP to disregard the resolved distance to the BGP next-hop in its decision process for selecting the best route to a destination. When configured in the config>router>bgp>best-path-selection context, this command applies to the comparison of two BGP routes with the same NLRI learned from base router BGP peers. When configured in the config>service>vprn context, this command applies to the comparison of two BGP-VPN routes for the same IP prefix imported into the VPRN from the base router BGP instance. When configured in the config>service>vprn>bgp>best-path-selection context, this command applies to the comparison of two BGP routes for the same IP prefix learned from VPRN BGP peers.

The **no** form of this command (no ignore-nh-metric) restores the default behavior whereby BGP factors distance to the next-hop into its decision process.

Default

no ignore-nh-metric

Platforms

7705 SAR Gen 2

13.31 ignore-oper-down

ignore-oper-down

Syntax

[**no**] **ignore-oper-down**

Context

[Tree] (config>service>epipe>sap ignore-oper-down)

Full Context

configure service epipe sap ignore-oper-down

Description

This command enables the ability to ignore the operationally down status for service oper state calculation. An Epipe service does not transition to Oper State: Down when a SAP fails and when this optional command is configured under that specific SAP. Only a single SAP in an Epipe may have this optional command included. The command can be used in Epipes with or without EVPN enabled.

The **no** form of this command disables whether a service ignores the operationally down state of the SAP.

Default
no ignore-oper-down

Platforms
7705 SAR Gen 2

13.32 ignore-rapid-commit

ignore-rapid-commit

Syntax
[**no**] **ignore-rapid-commit**

Context
[Tree] (config>service>vprn>dhcp6>server ignore-rapid-commit)
[Tree] (config>router>dhcp6>server ignore-rapid-commit)

Full Context
configure service vprn dhcp6 local-dhcp-server ignore-rapid-commit
configure router dhcp6 local-dhcp-server ignore-rapid-commit

Description
This command enables the Rapid Commit Option for DHCP6.
The **no** form of this command disables the Rapid Commit Option.

Platforms
7705 SAR Gen 2

13.33 ignore-router-id

ignore-router-id

Syntax
ignore-router-id include-internal *family* [*family*]
[**no**] **ignore-router-id**

Context
[Tree] (config>router>bgp>best-path-selection ignore-router-id)

[Tree] (config>service>vprn>bgp>best-path-selection ignore-router-id)

Full Context

configure router bgp best-path-selection ignore-router-id

configure service vprn bgp best-path-selection ignore-router-id

Description

When the **ignore-router-id** command is present, and the current best path to a destination was learned from EBGP peer X with BGP identifier x and a new path is received from EBGP peer Y with BGP identifier y, the best path remains unchanged if the new path is equivalent to the current best path up to the BGP identifier comparison – even if y is less than x.

The **no** form of this command restores the default behavior of selecting the route with the lowest BGP identifier (y) as best.

Default

no ignore-router-id

Parameters

family

Specifies up to two internal families to be included in this configuration.

Values mvpn-ipv4, mvpn-ipv6

include-internal

Specifies to ignore the router ID value even when comparing two IGBP paths or an EBGP and an IBGP path.

Platforms

7705 SAR Gen 2

13.34 ignore-standby-signaling

ignore-standby-signaling

Syntax

[**no**] **ignore-standby-signaling**

Context

[Tree] (config>service>vpls>endpoint ignore-standby-signaling)

[Tree] (config>service>vpls>spoke-sdp ignore-standby-signaling)

Full Context

configure service vpls endpoint ignore-standby-signaling

configure service vpls spoke-sdp ignore-standby-signaling

Description

When this command is enabled, the node ignores the standby-bit received from the TLDP peers for the specific spoke-SDP and performs internal tasks without taking it into account.

This command is present at the endpoint level and the spoke-SDP level. If the spoke-SDP is part of the explicit-endpoint, this setting cannot be changed at the spoke-SDP level. The existing spoke-SDP will become part of the explicit-endpoint only if the setting is not conflicting. The newly created spoke-SDP, which is a part of the specified explicit-endpoint, will inherit this setting from the endpoint configuration.

Default

no ignore-standby-signaling

Platforms

7705 SAR Gen 2

13.35 igp-instance

igp-instance

Syntax

igp-instance *igp-instance*

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls>sr-ospf igp-instance)

[Tree] (config>oam-pm>session>ip>tunnel>mpls>sr-isis igp-instance)

Full Context

configure oam-pm session ip tunnel mpls sr-ospf igp-instance

configure oam-pm session ip tunnel mpls sr-isis igp-instance

Description

This command configures the IGP instance to tunnel IP packets for the session test.

Default

igp-instance 0

Parameters

igp-instance

Specifies the IGP instance used to tunnel packets for the session.

Values

isis-inst	0 to 127
ospf-inst	0 to 31
ospf3-inst	0 to 31,64 to 95

Platforms

7705 SAR Gen 2

13.36 igp-shortcut

igp-shortcut

Syntax

igp-shortcut [**lfa-protect** | **lfa-only**] [**allow-sr-over-srte**]

igp-shortcut relative-metric [*offset*] [**allow-sr-over-srte**]

no igp-shortcut

Context

[Tree] (config>router>mpls>lsp-template igp-shortcut)

[Tree] (config>router>mpls>lsp igp-shortcut)

Full Context

configure router mpls lsp-template igp-shortcut

configure router mpls lsp igp-shortcut

Description

This command enables the use of a specific RSVP LSP by IS-IS and OSPF routing protocols as a shortcut or as a forwarding adjacency for resolving IGP routes.

When the **igp-shortcut** or the **advertise-tunnel-link** option is enabled at the IGP instance level, all RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in **config>router>mpls>lsp>to**, corresponds to a router-id of a remote node.

The **lfa-protect** option allows an LSP to be included in both the main SPF and the Loop-Free Alternate (LFA) SPF. For a given prefix, the LSP can be used either as a primary next-hop or as an LFA next-hop, but not both. If the main SPF computation selected a tunneled primary next-hop for a prefix, the LFA SPF will not select an LFA next-hop for this prefix and the protection of this prefix will rely on the RSVP LSP

FRR protection. If the main SPF computation selected a direct primary next-hop, then the LFA SPF will select an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.

The **lfa-only** option allows an LSP to be included in the LFA SPF only such that the introduction of IGP shortcuts does not impact the main SPF decision. For a given prefix, the main SPF always selects a direct primary next-hop. The LFA SPF selects an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.

When the **relative-metric** option is enabled, IGP will apply the shortest IGP cost between the endpoints of the LSP plus the value of the offset (instead of the LSP operational metric) when computing the cost of a prefix which is resolved to the LSP. The offset value is optional and it defaults to zero. The minimum net cost for a prefix is one (1) after applying the offset. The TTM continues the show the LSP operational metric as provided by MPLS. In other words, applications such as LDP-over-RSVP (when IGP shortcut is disabled) and BGP and static route shortcuts will continue to use the LSP operational metric.

The **relative-metric** option is mutually exclusive with the **lfa-protect** or the **lfa-only** options. In other words, an LSP with the **relative-metric** option enabled cannot be included in the LFA SPF, and vice-versa, when the **igp-shortcut** option is enabled in the IGP.

Finally, the **relative-metric** option is ignored when forwarding adjacency is enabled in IS-IS or OSPF. In this case, IGP advertises the LSP as a point-to-point unnumbered link along with the LSP operational metric as returned by MPLS and capped to maximum link metric allowed in that IGP. Both the main SPF and the LFA SPFs will use the local IGP database to resolve the routes.

When the router performs local SPF, the SR-TE LSP is used as an eligible IGP shortcut for SRv4 or SRv6 only if the LSP is explicitly allowed using the **allow-sr-over-srte** option when the top SID in the SR-TE LSP is an adjacency SID.

The **no** form of this command disables the use of a specific RSVP LSP by IS-IS and OSPF routing protocols as a shortcut or a forwarding adjacency for resolving IGP routes.

Default

igp-shortcut. All RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP corresponds to a router-id of a remote node.

Parameters

lfa-protect

Specifies an LSP is included in both the main SPF and the LFA SPF.

lfa-only

Specifies an LSP is included in the LFA SPF only.

relative-metric [*offset*]

Specifies the shortest IGP cost between the endpoints of the LSP plus the configured offset, instead of the LSP operational metric returned by MPLS, is used when calculating the cost of prefix resolved to this LSP. The offset parameter is an integer and is optional. An offset value of zero is used when the relative-metric option is enabled without specifying the offset parameter value.

Values [-10, +10]

allow-sr-over-srte

Specifies that the LSP or LSP template is eligible as an IGP shortcut.

Platforms

7705 SAR Gen 2

igp-shortcut

Syntax

igp-shortcut

Context

[Tree] (config>router>isis igp-shortcut)

Full Context

configure router isis igp-shortcut

Description

This command enables the use of an RSVP-TE or SR-TE shortcut for resolving IGP routes by OSPF or IS-IS routing protocols.

This command instructs IGP to include RSVP LSPs and SR-TE LSPs originating on this node and terminating on the router ID of a remote node as direct links with a metric equal to the metric provided by MPLS.

During the IP reach calculation to determine the reachability of nodes and prefixes, LSPs are overlaid and the LSP metric is used to determine the subset of paths that are equal lowest cost to reach a node or a prefix. If the user enabled the **relative-metric** option for this LSP, IGP will apply the shortest IGP cost between the endpoints of the LSP plus the value of the offset, instead of the LSP operational metric, when computing the cost of a prefix that is resolved to the LSP.

When a prefix is resolved to a tunnel next-hop, the packet is sent labeled with the label stack corresponding to the NHLFE of the RSVP-TE or SR-TE LSP, as well as the explicit-null IPv6 label at the bottom of the stack in the case of an IPv6 prefix. Any network event causing one or more IGP shortcuts to go down will trigger a full SPF computation, which may result in installing a new route over an updated set of tunnel next-hops and IP next-hops.

When **igp-shortcut** is enabled at the IGP instance level, all RSVP-TE and SR-TE LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in **config>router>mpls>lsp>to**, corresponds to a router ID of a remote node. LSPs with a destination corresponding to an interface address or any other loopback interface address of a remote node are automatically not considered by IGP. The user can, however, exclude a specific RSVP-TE or SR-TE LSP from being used as a shortcut for resolving IGP routes by entering the **config>router>mpls>lsp>no igp-shortcut** command.

The SPF in IGP only uses RSVP LSPs as forwarding adjacencies, IGP shortcuts, or as endpoints for LDP-over-RSVP. These applications of RSVP LSPs are mutually exclusive at the IGP instance level. If two or more options are enabled in the same IGP instance, then forwarding adjacency takes precedence over the shortcut application, which takes precedence over the LDP-over-RSVP application.

The SPF in IGP uses SR-TE LSPs as IGP shortcuts only.

When ECMP is enabled on the system and multiple equal-cost paths exist for a prefix, the following selection criteria are used to pick up the set of tunnel and IP next-hops to program in the data path.

- Where a destination is a tunnel-endpoint (including external prefixes with tunnel-endpoint as the next-hop), the tunnel with lowest tunnel-index is selected (the IP next-hop is never used in this case).
- Where a destination is not a tunnel-endpoint:
  - LSPs with metric higher than underlying IGP cost between the endpoint of the LSP are excluded
  - Tunnel next-hops are preferred over IP next-hops
  - Within tunnel next-hops, the following priority applies to selection:
    1. The lowest endpoint-to-destination cost is selected
    2. If the endpoint-to-destination costs are the same, the lowest endpoint node router ID is selected
    3. If the router IDs are the same, the lowest tunnel index is selected
  - Within IP next-hops, the following priority applies to selection:
    1. The lowest downstream router ID is selected
    2. If the downstream router IDs are the same, the lowest interface-index is selected

Note:
Although ECMP is not performed across both the IP and tunnel next-hops, the tunnel endpoint may lie in one of the shortest IGP paths for that prefix. In that case, the tunnel next-hop is always selected as long as the prefix cost using the tunnel is equal to or lower than the IGP cost.

When both RSVP-TE and SR-TE IGP shortcuts are available, the IP reach calculation, in the unicast routing table, will first follow the above ECMP tunnel and IP next-hop selection rules when resolving a prefix over IGP shortcuts. After the set of ECMP tunnel and IP next-hops have been selected, the preference of tunnel type is then applied based on the user setting for prefix family resolution. If the user enabled resolution of the prefix family to both RSVP-TE and SR-TE tunnel types, the TTM tunnel preference value is used to select one type for the prefix. In other words, an RSVP-TE LSP type is preferred to an SR-TE LSP type on a per-prefix basis.

The ingress IOM sprays the packets for this prefix over the set of tunnel next-hops and IP next-hops based on the hashing routine currently supported for IPv4 packets.

This feature provides IGP with the capability to populate the multicast RTM with the prefix IP next-hop when both the **igp-shortcut** and the **multicast-import** options are enabled in IGP. The unicast RTM can still use the tunnel next-hop for the same prefix. The SPF keeps track of both the direct first hop and the tunneled first hop of a node, which is added to the Dijkstra tree.

Platforms

7705 SAR Gen 2

igp-shortcut

Syntax

igp-shortcut

Context

[Tree] (config>router>ospf3 igp-shortcut)

[Tree] (config>router>ospf igp-shortcut)

Full Context

configure router ospf3 igp-shortcut

configure router ospf igp-shortcut

Description

This command enables the use of an RSVP-TE or SR-TE shortcut for resolving IGP routes by OSPF or IS-IS routing protocols.

This command instructs IGP to include RSVP LSPs and SR-TE LSPs originating on this node and terminating on the router ID of a remote node as direct links with a metric equal to the metric provided by MPLS.

During the IP reach calculation to determine the reachability of nodes and prefixes, LSPs are overlaid and the LSP metric is used to determine the subset of paths that are equal lowest cost to reach a node or a prefix. If the user enabled the **relative-metric** option for this LSP, IGP will apply the shortest IGP cost between the endpoints of the LSP plus the value of the offset, instead of the LSP operational metric, when computing the cost of a prefix that is resolved to the LSP.

When a prefix is resolved to a tunnel next hop, the packet is sent labeled with the label stack corresponding to the NHLFE of the RSVP-TE or SR-TE LSP, as well as the explicit-null IPv6 label at the bottom of the stack in the case of an IPv6 prefix. Any network event causing one or more IGP shortcuts to go down will trigger a full SPF computation, which may result in installing a new route over an updated set of tunnel next-hops and IP next-hops.

When **igp-shortcut** is enabled at the IGP instance level, all RSVP-TE and SR-TE LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in **config>router>mpls>lsp>to**, corresponds to a router ID of a remote node. LSPs with a destination corresponding to an interface address or any other loopback interface address of a remote node are automatically not considered by IGP. The user can, however, exclude a specific RSVP-TE or SR-TE LSP from being used as a shortcut for resolving IGP routes by entering the **config>router>mpls>lsp>no igp-shortcut** command.

The SPF in IGP only uses RSVP LSPs as forwarding adjacencies, IGP shortcuts, or as endpoints for LDP-over-RSVP. These applications of RSVP LSPs are mutually exclusive at the IGP instance level. If two or more options are enabled in the same IGP instance, then forwarding adjacency takes precedence over the shortcut application, which takes precedence over the LDP-over-RSVP application.

The SPF in IGP uses SR-TE LSPs as IGP shortcuts only.

When ECMP is enabled on the system and multiple equal-cost paths exist for a prefix, the following selection criteria are used to pick up the set of tunnel and IP next-hops to program in the data path.

- Where a destination is a tunnel-endpoint (including external prefixes with tunnel-endpoint as the next hop), the tunnel with lowest tunnel-index is selected (the IP next hop is never used in this case).
- Where a destination is not a tunnel-endpoint:
  - LSPs with metric higher than underlying IGP cost between the endpoint of the LSP are excluded
  - Tunnel next-hops are preferred over IP next-hops
  - Within tunnel next-hops:
    1. The lowest endpoint-to-destination cost is selected
    2. If the endpoint-to-destination costs are the same, the lowest endpoint node router ID is selected
    3. If the router IDs are the same, the lowest tunnel index is selected

– Within IP next-hops:

1. The lowest downstream router ID is selected
2. If the downstream router IDs are the same, the lowest interface-index is selected

Note:
Although ECMP is not performed across both the IP and tunnel next-hops, the tunnel endpoint may lie in one of the shortest IGP paths for that prefix. In that case, the tunnel next hop is always selected as long as the prefix cost using the tunnel is equal or lower than the IGP cost.

When both RSVP-TE and SR-TE IGP shortcuts are available, the IP reach calculation, in the unicast routing table, will first follow the above ECMP tunnel and IP next hop selection rules when resolving a prefix over IGP shortcuts. After the set of ECMP tunnel and IP next-hops have been selected, the preference of tunnel type is then applied based on the user setting of the resolution of the family of the prefix. If the user enabled resolution of the prefix family to both RSVP-TE and SR-TE tunnel types, the TTM tunnel preference value is used to select one type for the prefix. In other words, the RSVP-TE LSP type is preferred to an SR-TE LSP type on a per-prefix basis.

The ingress IOM sprays the packets for this prefix over the set of tunnel next-hops and IP next-hops based on the hashing routine currently supported for IPv4 packets.

This feature provides IGP with the capability to populate the multicast RTM with the prefix IP next hop when both the **igp-shortcut** and the **multicast-import** options are enabled in IGP. The unicast RTM can still make use of the tunnel next hop for the same prefix. This change is made possible with the enhancement by which SPF keeps track of both the direct first hop and the tunneled first hop of a node which is added to the Dijkstra tree.

Platforms

7705 SAR Gen 2

13.37 iid-tlv-enable

iid-tlv-enable

Syntax

[**no**] **iid-tlv-enable**

Context

[Tree] (config>service>vprn>isis iid-tlv-enable)

Full Context

configure service vprn isis iid-tlv-enable

Description

This command enables IS-IS multi-instance (MI) as described in draft-ietf-isis-mi-02. Multiple instances allow instance-specific adjacencies to be formed that support multiple network topologies on the same

physical interfaces. Each instance has an LSDB, and each PDU contains a TLV identifying the instance and the topology to which the PDU belongs.

The **iid-tlv-enable** (based on draft-ietf-isis-mi-02) and **standard-multi-instance** (based on draft-ginsberg-isis-mi-bis-01) commands cannot be configured in the same instance, because the MAC addresses and PDUs in each standard are incompatible.

Default

no iid-tlv-enable

Platforms

7705 SAR Gen 2

iid-tlv-enable

Syntax

[**no**] **iid-tlv-enable**

Context

[Tree] (config>router>isis iid-tlv-enable)

Full Context

configure router isis iid-tlv-enable

Description

This command enables IS-IS multi-instance (MI) as described in *draft-ietf-isis-mi-02*. Multiple instances allows the formation of instance-specific adjacencies that support multiple network topologies on the same physical interfaces. Each instance has an LSDB, and each PDU contains a TLV that identifies the instance and the topology to which the PDU belongs.

The **iid-tlv-enable** (based on *draft-ietf-isis-mi-02*) and **standard-multi-instance** (based on *draft-ginsberg-isis-mi-bis-01*) commands cannot be configured in the same instance, because the MAC addresses and PDUs in each standard are incompatible.

The **no** form of this command disables IS-IS MI.

Platforms

7705 SAR Gen 2

13.38 ike-auth-algorithm

ike-auth-algorithm

Syntax

ike-auth-algorithm {**md5** | **sha1** | **sha256** | **sha384** | **sha512** | **aes-xcbc** | **auth-encryption**}

Context

[Tree] (config>ipsec>ike-transform ike-auth-algorithm)

Full Context

configure ipsec ike-transform ike-auth-algorithm

Description

This command specifies the IKE authentication algorithm for the IKE transform

Default

ike-auth-algorithm sha1

Parameters

auth-algorithm

Specifies the values used to identify the hashing algorithm

Values

md5 — Configures the use of the hmac-md5 algorithm for authentication

sha1 — Configures the use of the hmac-sha1 algorithm for authentication

sha256 — Configures the use of the hmac-sha256 algorithm for authentication.

sha384 — Configures the use of the hmac-sha384 algorithm for authentication

sha512 — Configures the use of the hmac-sha512 algorithm for authentication.

aes-xcbc — Configures the use of aes-xcbc (RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*) algorithm for authentication.

Platforms

7705 SAR Gen 2

13.39 ike-encryption-algorithm

ike-encryption-algorithm

Syntax

ike-encryption-algorithm {**des** | **3des** | **aes128** | **aes192** | **aes256** | **aes128-gcm8** | **aes128-gcm16** | **aes256-gcm8** | **aes256-gcm16**}

Context

[Tree] (config>ipsec>ike-transform ike-encryption-algorithm)

Full Context

configure ipsec ike-transform ike-encryption-algorithm

Description

This command specifies the IKE encryption algorithm to be used in the IKE transform instance.

Default

ike-encryption-algorithm aes128

Parameters

encryption-algorithm

Specifies the IKE encryption algorithm.

Values **des** — Configures the 56-bit des algorithm for encryption. This is an older algorithm with relatively weak security. While better than nothing, it should only be used where a strong algorithm is not available on both ends at an acceptable performance level.

3des — Configures the 3-des algorithm for encryption. This is a modified application of the des algorithm which uses multiple des operations to make information more secure.

aes128 — Configures the aes algorithm with a block size of 128 bits. This is a mandatory implementation size for aes. This is a very strong algorithm choice.

aes192 — Configures the aes algorithm with a block size of 192 bits. This is a stronger version of aes.

aes256 — Configures the aes algorithm with a block size of 256 bits. This is the strongest available version of aes.

aes128-gcm8 - Configures ESP to use aes-gcm with a 128-bit key size and an 8-byte ICV for encryption and authentication.

aes128-gcm16 - Configures ESP to use aes-gcm with a 128-bit key size and a 16-byte ICV for encryption and authentication.

aes256-gcm8 - Configures ESP to use aes-gcm with a 256-bit key size and an 8-byte ICV for encryption and authentication.

aes256-gcm16 - This parameter configures ESP to use aes-gcm with a 256-bit key size and a 16-byte ICV for encryption and authentication.

Platforms

7705 SAR Gen 2

13.40 ike-mode

ike-mode

Syntax

ike-mode {**main** | **aggressive**}

no ike-mode

Context

[Tree] (config>ipsec>ike-policy ike-mode)

Full Context

configure ipsec ike-policy ike-mode

Description

This command specifies one of either two modes of operation. IKE version 1 can support main mode and aggressive mode. The difference lies in the number of messages used to establish the session.

The **no** form of this command reverts to the default.

Default

no ike-mode

Parameters

main

Specifies identity protection for the hosts initiating the IPsec session. This mode takes slightly longer to complete.

aggressive

Specifies that the aggressive mode provides no identity protection but is faster.

Platforms

7705 SAR Gen 2

13.41 ike-policy

ike-policy

Syntax

ike-policy *ike-policy-id* [**create**]

no ike-policy *ike-policy-id*

Context

[Tree] (config>ipsec ike-policy)

Full Context

configure ipsec ike-policy

Description

Commands in this context configure an IKE policy.

The **no** form of this command

Parameters

ike-policy-id

Specifies a policy ID value to identify the IKE policy.

Values 1 to 2048

Platforms

7705 SAR Gen 2

ike-policy

Syntax

ike-policy *ike-policy-id*

no ike-policy

Context

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>dyn ike-policy)

[Tree] (config>service>vprn>if>sap>ipsec-gw ike-policy)

[Tree] (config>router>if>ipsec>ipsec-tunnel>dyn ike-policy)

[Tree] (config>service>ies>if>sap>ipsec-gw ike-policy)

[Tree] (config>ipsec>trans-mode-prof>dyn ike-policy)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>dyn ike-policy)

Full Context

configure service ies interface ipsec ipsec-tunnel dynamic-keying ike-policy

configure service vprn interface sap ipsec-gw ike-policy

configure router interface ipsec ipsec-tunnel dynamic-keying ike-policy

configure service ies interface sap ipsec-gw ike-policy

configure ipsec ipsec-transport-mode-profile dynamic-keying ike-policy

configure service vprn interface ipsec ipsec-tunnel dynamic-keying ike-policy

Description

This command specifies the ID of the IKE policy used for IKE negotiation.

The **no** form of this command removes the IKE policy ID from the configuration.

Parameters

ike-policy-id

Specifies the IKE policy ID.

Values 1 to 2048

Platforms

7705 SAR Gen 2

13.42 ike-prf-algorithm

ike-prf-algorithm

Syntax

ike-prf-algorithm {**md5** | **sha1** | **sha256** | **sha384** | **sha512** | **aes-xcbc** | **same-as-auth**}

Context

[Tree] (config>ipsec>ike-transform ike-prf-algorithm)

Full Context

configure ipsec ike-transform ike-prf-algorithm

Description

This command specifies the PRF algorithm to use for IKE security association.

Note:
If an authenticated encryption algorithm like AES-GCM is used for IKE encryption algorithm, **same-as-auth** cannot be used for **ike-prf-algorithm**.

Default

ike-prf-algorithm same-as-auth

Parameters

md5

This parameter configures IKE to use the **hmac-md5** algorithm for PRF.

sha1

This parameter configures IKE to use the **hmac-sha1** algorithm for PRF.

sha256

This parameter configures IKE to use the **hmac-sha256** algorithm for PRF.

sha384

This parameter configures IKE to use the **hmac-sha384** algorithm for PRF.

sha512

This parameter configures IKE to use the **hmac-sha512** algorithm for PRF.

aes-xcbc

This parameter configures IKE to use the **aes128-xcbc** algorithm for PRF.

same-as-auth

This parameter configures the same algorithm as IKE authentication algorithm.

Platforms

7705 SAR Gen 2

13.43 ike-transform

ike-transform

Syntax

ike-transform *ike-transform-id* [*ike-transform-id* ...(up to 4 max)]

no ike-transform

Context

[Tree] (config>ipsec>ike-policy ike-transform)

Full Context

configure ipsec ike-policy ike-transform

Description

This command specifies the IKE transform to be used in the IKE policy. Up to four IKE transforms can be specified. If multiple IDs are specified, the system selects an IKE transform based on the peer's proposal. If the system is a tunnel initiator, it uses the configured IKE transform to generate the SA payload.

Default

no ike-transform

Parameters

ike-transform-id

Specifies up to four existing IKE transform instances to be associated with this IKE policy.

Values 1 to 4096

Platforms

7705 SAR Gen 2

ike-transform

Syntax

ike-transform *ike-transform-id* [**create**]

no ike-transform *ike-transform-id*

Context

[Tree] (config>ipsec ike-transform)

Full Context

configure ipsec ike-transform

Description

This commands creates a new or enters an existing IKE transform instance. The IKE transform include following configuration for IKE SA:

- DH Group
- IKE authentication algorithm
- IKE encryption algorithm
- IKE SA lifetime

The *ike-transform-id* is referenced in the **ike-policy** configuration.

Parameters

ike-transform

Specifies a number used to uniquely identify an IKE transform instance.

Values 1 to 4096

create

Keyword used to create the ike-transform instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

13.44 ike-version

ike-version

Syntax

ike-version {**1** | **2**}

Context

[Tree] (config>ipsec>ike-policy ike-version)

Full Context

configure ipsec ike-policy ike-version

Description

This command sets the IKE version (1 or 2) that the ike-policy will use.

Default

ike-version 1

Parameters

1 | 2

Specifies the version of IKE protocol.

Platforms

7705 SAR Gen 2

13.45 ikev1-ph1-responder-delete-notify

ikev1-ph1-responder-delete-notify

Syntax

[**no**] **ikev1-ph1-responder-delete-notify**

Context

[Tree] (config>ipsec>ike-policy ikev1-ph1-responder-delete-notify)

Full Context

configure ipsec ike-policy ikev1-ph1-responder-delete-notify

Description

This command specifies the system, when deleting an IKEv1 phase 1 SA for which it was the responder, to send a delete notification to the peer. This command only applies when the configured ike-version 1. This command is ignored with IKE version 2.

The **no** form of this command reverts to the default.

Default

ikev1-ph1-responder-delete-notify

Platforms

7705 SAR Gen 2

13.46 ikev2-fragment

ikev2-fragment

Syntax

ikev2-fragment mtu *octets* **reassembly-timeout** *seconds*

no ikev2-fragment

Context

[Tree] (config>ipsec>ike-policy ikev2-fragment)

Full Context

configure ipsec ike-policy ikev2-fragment

Description

This command enables IKEv2 protocol level fragmentation (RFC 7383). The specified MTU is the maximum size of IKEv2 packet.

Default

no ikev2-fragment

Parameters

octets

Specifies the MTU for IKEv2 messages.

Values 512 to 9000

seconds

Specifies the timeout for reassembly.

Values 1 to 5

Platforms

7705 SAR Gen 2

13.47 implicit-null-label

implicit-null-label

Syntax

[**no**] **implicit-null-label**

Context

[Tree] (config>router>ldp implicit-null-label)

Full Context

configure router ldp implicit-null-label

Description

This command enables the use of the implicit null label. Use this command to signal the implicit null option for all LDP FECs for which this node is the egress LER.

The **no** form of this command disables the signaling of the implicit null label.

Default

no implicit-null-label

Platforms

7705 SAR Gen 2

implicit-null-label

Syntax

[**no**] **implicit-null-label**

Context

[Tree] (config>router>rsvp implicit-null-label)

Full Context

configure router rsvp implicit-null-label

Description

This command enables the use of the implicit null label.

Signaling the IMPLICIT NULL label value for all RSVP LSPs can be enabled for which this node is the egress LER. RSVP must be shut down before being able to change this configuration option.

The egress LER does not signal the implicit null label value on P2MP RSVP LSPs. However, the Penultimate Hop Popping (PHP) node can honor a Resv message with the label value set to the implicit null.

The **no** form of this command disables the signaling of the implicit null label.

Default

no implicit-null-label

Platforms

7705 SAR Gen 2

implicit-null-label

Syntax

implicit-null-label [**enable** | **disable**]

no implicit-null-label

Context

[Tree] (config>router>rsvp>interface implicit-null-label)

Full Context

configure router rsvp interface implicit-null-label

Description

This command enables the use of the implicit null label over a specific RSVP interface.

All LSPs for which this node is the egress LER and for which the path message is received from the previous hop node over this RSVP interface will signal the implicit null label. This means that if the egress LER is also the merge-point (MP) node, then the incoming interface for the path refresh message over the bypass dictates if the packet will use the implicit null label or not. The same for a 1-to-1 detour LSP.

The user must shut down the RSVP interface before being able to change the implicit null configuration option.

The **no** form of this command returns the RSVP interface to use the RSVP level configuration value.

Default

no implicit-null-label

Parameters

enable

Enables the implicit null label.

disable

Disables the implicit null label.

Platforms

7705 SAR Gen 2

13.48 import

import

Syntax

import *policy-name*

no import

Context

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping import)

[Tree] (config>service>vpls>sap>igmp-snooping import)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping import)

[Tree] (config>service>vpls>sap>mld-snooping import)

[Tree] (config>service>vpls>mesh-sdp>mld-snooping import)

[Tree] (config>service>vpls>spoke-sdp>mld-snooping import)

Full Context

configure service vpls mesh-sdp igmp-snooping import

configure service vpls sap igmp-snooping import

configure service vpls spoke-sdp igmp-snooping import

configure service vpls sap mld-snooping import

configure service vpls mesh-sdp mld-snooping import

configure service vpls spoke-sdp mld-snooping import

Description

This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a SAP at any time.

The **no** form of this command removes the policy association from the SAP or SDP.

Default

no import

Parameters

policy-name

Specifies the routing policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Routing policies are configured in the **config>router>policy-options** context The router policy must be defined before it can be imported.

Platforms

7705 SAR Gen 2

import

Syntax

import *plcy-or-long-expr* [*plcy-or-expr*]

no import

Context

[Tree] (config>service>vprn>bgp import)

[Tree] (config>service>vprn>bgp>group>neighbor import)

[Tree] (config>service>vprn>bgp>group import)

Full Context

configure service vprn bgp import

configure service vprn bgp group neighbor import

configure service vprn bgp group import

Description

This command is used to specify route policies that control the handling of inbound routes received from certain peers. Route policies are configured in the **config>router>policy-options** context.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in a peer-group) or neighbor level (only applies to the specified peer). The most specific level is used

The **import** command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine the modifications of each route and the final action to accept or reject the route.

Only one of the 15 objects referenced by the **import** command can be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters while the remaining 14 objects have a maximum length of 64 characters each.

When multiple **import** commands are issued, the last command entered overrides the previous command.

When an import policy is not specified, BGP routes are accepted by default.

The **no** form of this command removes the policy association.

Default

no import

Parameters

plcy-or-long-expr

Specifies the route policy name (up to 64 characters) or a policy logical expression (up to 255 characters).

plcy-or-expr

Specifies the route policy name (up to 64 characters) or a policy logical expression (up to 255 characters).

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name*

no import

Context

[Tree] (config>service>vprn>igmp>if import)

Full Context

configure service vprn igmp interface import

Description

This command imports a policy to filter IGMP packets.

The **no** form of this command removes the policy association from the IGMP instance.

Default

no import — No import policy specified.

Parameters

policy-name

Specifies the import route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

The specified name(s) must already be defined.

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name* [*policy-name* ... (up to 5 max)]

no import

Context

[Tree] (config>service>vprn>isis import)

Full Context

configure service vprn isis import

Description

This command applies one or more (up to five) route polices as IS-IS import policies.

When a prefix received in an IS-IS LSP is accepted by an entry in an IS-IS import policy, it is installed in the routing table, if it is the most preferred route to the destination.

When a prefix received in an IS-IS LSP is rejected by an entry in an IS-IS import policy, it is not installed in the routing table, even if it has the lowest preference value among all the routes to that destination.

The flooding of LSPs is unaffected by IS-IS import policy actions.

The **no** form of this command removes all policies from the configuration.

Default

no import

Parameters

policy-name

Identifies the export route policy name. Allowed values are any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes. The specified name(s) must already be defined.

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name*

no import

Context

[Tree] (config>service>vprn>mld>if import)

Full Context

configure service vprn mld interface import

Description

This command specifies the import route policy to be used for determining which membership reports are accepted by the router. Route policies are configured in the **config>router>policy-options** context.

When an import policy is not specified, all the MLD reports are accepted.

The **no** form of this command removes the policy association from the MLD instance.

Default

no import

Parameters

policy-name

Specifies the route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes. Route policies are configured in the **config>router>policy-options** context.

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name* [*policy-name*]

no import

Context

[Tree] (config>service>vprn>ospf>area import)

[Tree] (config>service>vprn>ospf3>area import)

Full Context

configure service vprn ospf area import

configure service vprn ospf3 area import

Description

This command configures ABR import policies to filter OSPFv2 Type 3 Summary-LSAs or OSPFv3 Inter-Area-Prefix-LSA between areas, to only permit the specified routes from being imported into an area.

This command cannot be used in OSPF area 0.

The **no** form of this command reverts to the default value.

Default

no import

Parameters

policy-name

Specifies the export route policy name. A maximum of five policy names can be specified. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

The specified policy names must be predefined and already exist in the system.

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name* [*policy-name*]

no import

Context

[Tree] (config>service>vprn>ospf import)

[Tree] (config>service>vprn>ospf3 import)

Full Context

configure service vprn ospf import

configure service vprn ospf3 import

Description

This command applies one or more (up to five) route polices as OSPF import policies. When a prefix received in an OSPF LSA is accepted by an entry in an OSPF import policy it is installed in the routing table if it is the most preferred route to the destination. When a prefix received in an OSPF LSA is rejected by an entry in an OSPF import policy it is not installed in the routing table, even if it has the lowest preference value among all the routes to that destination. The flooding of LSAs is unaffected by OSPF import policy actions.

Default

If an OSPF route has the lowest preference value among all routes to a destination it is installed in the routing table.

Parameters

policy-name

Specifies the import route policy name. A maximum of five policy names can be specified. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

The specified policy name(s) must be predefined and already exist in the system.

Platforms

7705 SAR Gen 2

import

Syntax

import {**join-policy** | **register-policy**} *policy-name* [*policy-name* **...**(**up to 5 max**)]

no import {**join-policy** | **register-policy**}

Context

[Tree] (config>service>vprn>pim import)

Full Context

configure service vprn pim import

Description

This command specifies the import route policy to be used for determining which routes are accepted from peers. Route policies are configured in the **config>router>policy-options** context. When an import policy is not specified, BGP routes are accepted by default.

The **no** form of this command removes the policy association from the IGMP instance.

Default

no import join-policy

no import register-policy

Parameters

join-policy

Use this command to filter PIM join messages which prevents unwanted multicast streams from traversing the network.

register-policy

This keyword filters register messages. PIM register filters prevent register messages from being processed by the RP. This filter can only be defined on an RP. When a match is found, the RP immediately sends back a register-stop message.

policy-name

Specifies the route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes. Route policies are configured in the **config>router>policy-options** context.

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name* [*policy-name* **...**(**up to 5 max**)]

no import

Context

[Tree] (config>service>vprn>rip>group import)

[Tree] (config>service>vprn>rip import)

[Tree] (config>service>vprn>ripng import)

[Tree] (config>service>vprn>rip>group>neighbor import)

[Tree] (config>service>vprn>ripng>group import)

[Tree] (config>service>vprn>ripng>group>neighbor import)

Full Context

configure service vprn rip group import

configure service vprn rip import

configure service vprn ripng import

configure service vprn rip group neighbor import

configure service vprn ripng group import

configure service vprn ripng group neighbor import

Description

This command configures import route policies to determine routes that will be accepted from RIP neighbors. If no import policy is specified, RIP accepts all routes from configured RIP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

If multiple policy names are specified, the policies are evaluated in the order that they are specified. The first policy that matches is applied. If multiple import commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of this command removes all policies from the configuration.

Default

no import

Parameters

policy-name

The import route policy name. Allowed values are any string up to 32 characters in length and composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes. The specified names must already be defined.

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name* [*policy-name*]

no import

Context

[Tree] (config>router>ldp import)

Full Context

configure router ldp import

Description

This command configures import route policies to determine which label bindings (FECs) are accepted from LDP neighbors. Policies are configured in the **config>router>policy-options** context.

If no import policy is specified, LDP accepts all label bindings from configured LDP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple import commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of this command removes all policies from the configuration.

Default

no import

Parameters

policy-name

Specifies up to five import route policy names, up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name*

no import

Context

[Tree] (config>router>igmp>if import)

Full Context

configure router igmp interface import

Description

This command applies the referenced IGMP policy (filter) to an interface subscriber or a group-interface. An IGMP filter is also known as a black/white list and it is defined under the **config>router>policy-options**.

When redirection is applied, only the import policy from the subscriber will be in effect. The import policy under the group interface is applicable only for IGMP states received directly on the SAP (AN in IGMP proxy mode).

The **no** form of the command removes the policy association from the IGMP instance.

Default

no import

Parameters

policy-name

The route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name*

no import

Context

[Tree] (config>router>mld>if import)

Full Context

configure router mld interface import

Description

This command specifies the import route policy to determine which membership reports are accepted by the router. Route policies are configured in the **config>router>policy-options** context.

When an import policy is not specified, all the MLD reports are accepted.

The **no** form of this command removes the policy association from the MLD instance.

Default

no import

Parameters

policy-name

Specifies the route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

Platforms

7705 SAR Gen 2

import

Syntax

import {**join-policy** | **register-policy**} [*policy-name* [*policy-name*]]

no import {**join-policy** | **register-policy**}

Context

[Tree] (config>router>pim import)

Full Context

configure router pim import

Description

This command specifies the import route policy to be used. Route policies are configured in the **config>router>policy-options** context.

When an import policy is not specified, BGP routes are accepted by default. Up to five import policy names can be specified.

The **no** form of this command removes the policy association from the instance.

Default

no import

Parameters

join-policy

Filters PIM join messages which prevents unwanted multicast streams from traversing the network.

register-policy

Filters register messages. PIM register filters prevent register messages from being processed by the RP. This filter can only be defined on an RP. When a match is found, the RP immediately sends back a register-stop message.

policy-name

Specifies the route policy name, up to 32 characters. Route policies are configured in the **config>router>policy-options** context.

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name*

no import

Context

[Tree] (config>service>pw-template>igmp-snooping import)

Full Context

configure service pw-template igmp-snooping import

Description

This command specifies the import routing policy to be used for IGMP packets. Only a single policy can be imported at a time.

The **no** form of the command removes the policy association.

Default

no import

Parameters

policy-name

Specifies the import policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Routing policies are configured in the config>router>policy-options context The router policy must be defined before it can be imported.

Platforms

7705 SAR Gen 2

import

Syntax

import type {**cert** | **key** | **crl**} **input** *url-string* **output** *filename* **format** *input-format* [**password** [*32 chars max*]]

Context

[Tree] (admin>certificate import)

Full Context

admin certificate import

Description

This command converts an input file (key/certificate/CRL) to a system format file. The following list summarizes the formats supported by this command:

- Certificate
  - PKCS #12
  - PKCS #7 PEM encoded
  - PKCS #7 DER encoded
  - PEM
  - DER
- Key
  - PKCS #12
  - PEM
  - DER
- CRL
  - PKCS #7 PEM encoded
  - PKCS #7 DER encoded
  - PEM
  - DER

Note:
If there are multiple objects with the same type in the input file, only the first object is extracted and converted.

Parameters

input *url-string*

Specifies the URL for the input file. This URL could be either a local CF card URL file or a FP URL to download the input file.

Values

url-string	<local-url> up to 99 characters
local-url	<cflash-id>/<file-path>
cflash-id	cf1:| cf2:| cf3:

output *filename*

Specifies the name of output file up to 95 characters. The output directory depends on the file type like following:

- Key: cf3:\system-pki\key
- Cert: cf3:\system-pki\cert
- CRL: cf3:\system-pki\CRL

type

The type of input file.

Values cert, key, crl

format

Specifies the format of input file.

Values pkcs12, pkcs7-der, pkcs7-pem, pem, der

password

Specifies the password to decrypt the input file in case that it is an encrypted PKCS#12 file.

Platforms

7705 SAR Gen 2

import

Syntax

import *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]

no import

Context

[Tree] (config>router>bgp>group>neighbor import)

[Tree] (config>router>bgp>group import)

[Tree] (config>router>bgp import)

Full Context

configure router bgp group neighbor import

configure router bgp group import

configure router bgp import

Description

This command specifies route policies that control the handling of inbound routes received from certain peers. Route policies are configured in the **config>router>policy-options** context.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific level is used.

The **import** command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine the modifications of each route and the final action to accept or reject the route.

Only one of the 15 objects referenced by the **import** command is allowed to be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters; the remaining 14 objects have a maximum length of 64 characters each.

When multiple **import** commands are issued, the last command entered overrides the previous command.

When an import policy is not specified, BGP routes are accepted by default.

The **no** form of this command removes the policy association.

Default

no import

Parameters

plcy-or-long-expr

Specifies the route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). Allowed values are any string up to 255 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

plcy-or-expr

Specifies the route policy name (up to 64 characters long) or a policy logical expression (up to 64 characters long). Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name* [*policy-name*]

no import

Context

[Tree] (config>router>isis import)

Full Context

configure router isis import

Description

This command specifies up to five route polices as IS-IS import policies.

When a prefix received in an IS-IS LSP is accepted by an entry in an IS-IS import policy, it is installed in the routing table, if it is the most preferred route to the destination.

When a prefix received in an IS-IS LSP is rejected by an entry in an IS-IS import policy, it is not installed in the routing table, even if it has the lowest preference value among all the routes to that destination.

The flooding of LSPs is unaffected by IS-IS import policy actions.

The **no** form of this command removes all policies from the configuration.

Default

no import

Parameters

policy-name

Specifies the import route policy name. Allowed values are any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. The specified names must already be defined.

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name* [*policy-name*]

no import

Context

[Tree] (config>router>ospf3 import)

[Tree] (config>router>ospf import)

Full Context

configure router ospf3 import

configure router ospf import

Description

This command applies one or more (up to 5) route polices as OSPF import policies. When a prefix received in an OSPF LSA is accepted by an entry in an OSPF import policy, it is installed in the routing table if it is the most preferred route to the destination. When a prefix received in an OSPF LSA is rejected by an entry in an OSPF import policy, it is not installed in the routing table, even if it has the lowest preference value among all the routes to that destination. The flooding of LSAs is unaffected by OSPF import policy actions. The **no** form of this command removes all policies from the configuration.

Default

no import

Parameters

policy-name

Specifies up to 5 export route policy names. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. The specified names must already be defined.

Platforms

7705 SAR Gen 2

import

Syntax

[**no**] **import** *policy-name* [*policy-name*]

Context

[Tree] (config>router>ospf3>area import)

[Tree] (config>router>ospf>area import)

Full Context

configure router ospf3 area import

configure router ospf area import

Description

This command configures ABR import policies to filter OSPFv2 Type 3 Summary-LSAs or OSPFv3 Inter-Area-Prefix-LSA between areas, in order to only permit the specified routes from being imported into an area.

This command cannot be used in OSPF area 0.

The **no** form of this command reverts to the default value.

Default

no import

Parameters

policy-name

Specifies up to five import route policy names. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. The specified names must already be defined.

Platforms

7705 SAR Gen 2

import

Syntax

import *policy-name* [*policy-name*]

no import

Context

[Tree] (config>router>ripng>group>neighbor import)

[Tree] (config>router>rip import)

[Tree] (config>router>rip>group import)

[Tree] (config>router>ripng>group import)

[Tree] (config>router>rip>group>neighbor import)

[Tree] (config>router>ripng import)

Full Context

configure router ripng group neighbor import

configure router rip import

configure router rip group import

configure router ripng group import

configure router rip group neighbor import

configure router ripng import

Description

This command configures import route policies to determine which routes are accepted from RIP neighbors. If no import policy is specified, RIP accepts all routes from configured RIP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple import commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of the command removes all policies from the configuration.

Default

no import

Parameters

policy-name

Specifies up to five import route policy names. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

The specified names must already be defined.

Platforms

7705 SAR Gen 2

import

Syntax

import {**ignore** | **accept** | **drop**}

Context

[Tree] (config>service>vprn>bgp-ipvpn>attribute-set import)

Full Context

configure service vprn bgp-ipvpn attribute-set import

Description

This command configures the reception behavior for ATTR_SETs in received VPN-IP routes.

Default

import ignore

Parameters

accept

Keyword to configure BGP to accept and process ATTR_SETs in received unicast VPN-IP routes (MPLS or SRv6) when they are imported into the VPRN. The path attributes contained inside the ATTR_SET are used for best-path selection within the VPRN, instead of the outer-path attributes attached to the imported VPN-IP route. The path attributes inside the ATTR_SET determine the path attributes of BGP routes advertised to PE-CE peers of the VPRN. However, the ATTR_SET is removed at the time of advertisement. VPRN BGP routes with attributes derived from **accept** processing are only advertised to EBGP peers and IBGP route reflector client peers. VPRN BGP routes are not advertised to BGP confederation peers. If the origin AS in the ATTR_SET attribute does not match the configured ASN, VPRN BGP routes with attributes derived from **accept** processing are advertised to IBGP peers that are not covered by a cluster configuration.

drop

Keyword to configure BGP to ignore and silently discard ATTR_SETs in received VPN-IP routes when they are imported into the VPRN. The path attributes contained inside the ATTR_SET are not used for best path selection within the VPRN. If a VPRN is not involved in an independent domain Layer 3 VPN service, Nokia recommends configuring the **import** command to **drop**.

ignore

Keyword to configure BGP to ignore ATTR_SETs in received VPN-IP routes when they are imported into the VPRN. The path attributes contained inside the ATTR_SET are not used for best-path selection within the VPRN. With the **ignore** parameter, the ATTR_SET attribute is transmitted unchanged to the CE. Nokia does not recommend configuring the **import** command to **ignore** in most deployments.

Platforms

7705 SAR Gen 2

13.49 import-grt

import-grt

Syntax

import-grt *plcy-or-long-expr* [*plcy-or-expr*]

no import-grt

Context

[Tree] (config>service>vprn>grt import-grt)

Full Context

configure service vprn grt-lookup import-grt

Description

This command associates policies to control the leaking of GRT routes into the associated VPRN.

The GRT route must have first been leaked by a **leak-export** policy defined under the **config>router** context. Then the route must match a route entry in the specified **import-grt** policy with an accept action.

The **no** form of this command removes route leaking policy associations and disables the leaking of GRT routes into the local VPRN.

Parameters

plcy-or-long-expr

Specifies route policy names, up to 64 characters, or a policy logical expression, up to 255 characters.

Values *plcy-or-long-expr*: *policy-name* | *long-expr*

policy-name: up to 64 characters

long-expr: up to 255 characters

plcy-or-expr

Specifies up to four route policy names, up to 64 characters, or a policy logical expression, up to 64 characters.

Values *plcy-or-expr*: *policy-name* | *expr*

policy-name: up to 64 characters

expr: up to 64 characters

Platforms

7705 SAR Gen 2

13.50 import-pmsi-routes

import-pmsi-routes

Syntax

import-pmsi-routes

Context

[Tree] (config>router>ldp import-pmsi-routes)

Full Context

configure router ldp import-pmsi-routes

Description

Commands in this context configure import-pmsi-routes.

For option B, the leafs or ABR/ASBR that are not directly connected to the root have no visibility of the root. As such, for LDP to build the recursive FEC it needs to cache the MVPN PMSI AD routes, this command gives the user the ability to manually enable caching of MVPN PMSI AD routes internally in LDP for EVPN or MVPN inter-as or **mvpn_no_export_community** intra-as.

Platforms

7705 SAR Gen 2

13.51 import-prefixes

import-prefixes

Syntax

[**no**] **import-prefixes** *policy-name*

Context

[Tree] (config>router>ldp>session-params>peer import-prefixes)

Full Context

configure router ldp session-parameters peer import-prefixes

Description

This command configures the import FEC prefix policy to determine which prefixes received from this LDP peer are imported and installed by LDP on this node. If resolved these FEC prefixes are then re-distributed to other LDP and T-LDP peers. A FEC prefix that is filtered out (deny) will not be imported. A FEC prefix that is filtered in (accept) will be imported.

If no import policy is specified, the node will import all prefixes received from this LDP/T-LDP peer. This policy is applied in addition to the global LDP policy and targeted session policy.

Policies are configured in the **config>router>policy-options** context. A maximum of five policy names can be specified. Peer address has to be the peer LSR-ID address.

The **no** form of the command removes the policy from the configuration.

Default

no import-prefixes - no import route policy is specified

Parameters

policy-name

Specifies up to five import-prefix route policy names. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string. The specified name(s) must already be defined.

Platforms

7705 SAR Gen 2

import-prefixes

Syntax

import-prefixes *policy-name* [*policy-name*]

no import-prefixes

Context

[Tree] (config>router>ldp>targeted-session import-prefixes)

Full Context

configure router ldp targeted-session import-prefixes

Description

This command configures the import route policy to determine which FEC prefix label bindings are accepted from targeted LDP neighbors into this node. A label binding that is filtered out (deny) will not be imported. A route that is filtered in (accept) will be imported.

If no import policy is specified, this node session will accept all bindings from configured targeted LDP neighbors. This policy is applied in addition to the global LDP policy.

Policies are configured in the **config>router>policy-option**s context. A maximum of five policy names can be specified.

The **no** form of this command removes the policy from the configuration.

Parameters

policy-name

Specifies up to five import policy names. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

13.52 import-tunnel-table

import-tunnel-table

Syntax

import-tunnel-table *policy-name* [*policy-name*]

no import-tunnel-table

Context

[Tree] (config>router>ldp import-tunnel-table)

Full Context

configure router ldp import-tunnel-table

Description

This command controls the import, in the tunnel table, of LDP tunnels to non-host prefixes. This command is only intended for importing tunnels; it cannot be used for preventing the import of any specific prefix and only non-host prefixes will be considered when evaluating this policy in this context. The LDP tunnels to these non-host prefixes must be created before they can be imported.

This command does not affect the automatic import of LDP tunnels to host prefixes.

The **no** version of this command removes all of the import policies and, by consequence, any tunnels to non-host prefixes from the tunnel table. If a non-host prefix tunnel is currently being used for forwarding, disabling this command may be service-impacting.

Default

no import-tunnel-table

Parameters

policy-name

Specifies up to five import route policy names. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

The specified policy names must already be defined.

Platforms

7705 SAR Gen 2

13.53 imported-format

imported-format

Syntax

imported-format {**any** | **secure**}

Context

[Tree] (config>system>security>pki imported-format)

Full Context

configure system security pki imported-format

Description

This command specifies the allowed format of imported certificates or keys in the cf3:/system-pki directory.

Default

imported-format any

Parameters

any

Allows any imported format.

secure

Only allows enhanced secure imported formats.

Platforms

7705 SAR Gen 2

13.54 improved-assert

improved-assert

Syntax

[**no**] **improved-assert**

Context

[Tree] (config>service>vprn>pim>if improved-assert)

Full Context

configure service vprn pim interface improved-assert

Description

This command enables improved assert processing on this interface. The PIM assert process establishes a forwarder for a LAN and requires interaction between the control and forwarding planes.

The assert process is started when data is received on an outgoing interface. This could impact performance if data is continuously received on an outgoing interface.

When enabled, the PIM assert process is done entirely on the control-plane with no interaction between the control and forwarding plane.

Default

improved-assert

Platforms

7705 SAR Gen 2

improved-assert

Syntax

[**no**] **improved-assert**

Context

[Tree] (config>router>pim>interface improved-assert)

Full Context

configure router pim interface improved-assert

Description

This command enables improved assert processing. The PIM assert process establishes a forwarder for a LAN and requires interaction between the control and forwarding planes. The assert process is started when data is received on an outgoing interface meaning that duplicate traffic is forwarded to the LAN until the forwarder is negotiated among the routers.

When the **improved-assert** command is enabled, the PIM assert process is done entirely in the control plane. The advantages are that it eliminates duplicate traffic forwarding to the LAN. It also improves performance since it removes the required interaction between the control and data planes.

Note:
improved-assert is still fully interoperable with the RFC 4601, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)* and RFC 2362, *Protocol Independent Multicast-Sparse Mode (PIM-SM)*, implementations. However, there may be conformance tests that may fail if the tests expect control-data plane interaction in determining the assert winner. Disabling the **improved-assert** command when performing conformance tests is recommended.

Default

improved-assert

Platforms

7705 SAR Gen 2

13.55 in-plus-profile-octets-discarded-count

in-plus-profile-octets-discarded-count

Syntax

[**no**] **in-plus-profile-octets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters in-plus-profile-octets-discarded-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters in-plus-profile-octets-discarded-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters in-plus-profile-octets-discarded-count

configure log accounting-policy custom-record policer e-counters in-plus-profile-octets-discarded-count

Description

This command includes the in-plus profile octets discarded count.

The **no** form of this command excludes the in-plus profile octets discarded count.

Default

no in-plus-profile-octets-discarded-count

Platforms

7705 SAR Gen 2

13.56 in-plus-profile-octets-forwarded-count

in-plus-profile-octets-forwarded-count

Syntax

[**no**] **in-plus-profile-octets-forwarded-count**

Context

[Tree] (config>log>acct-policy>cr>policer>e-counters in-plus-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters in-plus-profile-octets-forwarded-count)

Full Context

configure log accounting-policy custom-record policer e-counters in-plus-profile-octets-forwarded-count

configure log accounting-policy custom-record ref-policer e-counters in-plus-profile-octets-forwarded-count

Description

This command includes the in-plus profile octets forwarded count.

The **no** form of this command excludes the in-plus profile octets forwarded count.

Default

no in-plus-profile-octets-forwarded-count

Platforms

7705 SAR Gen 2

13.57 in-plus-profile-octets-offered-count

in-plus-profile-octets-offered-count

Syntax

[**no**] **in-plus-profile-octets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters in-plus-profile-octets-offered-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters in-plus-profile-octets-offered-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters in-plus-profile-octets-offered-count

configure log accounting-policy custom-record policer e-counters in-plus-profile-octets-offered-count

Description

This command includes the in-plus profile octets offered count.

The **no** form of this command excludes the in-plus profile octets offered count.

Default

no in-plus-profile-octets-offered-count

Platforms

7705 SAR Gen 2

13.58 in-plus-profile-packets-discarded-count

in-plus-profile-packets-discarded-count

Syntax

[**no**] **in-plus-profile-packets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters in-plus-profile-packets-discarded-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters in-plus-profile-packets-discarded-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters in-plus-profile-packets-discarded-count

configure log accounting-policy custom-record policer e-counters in-plus-profile-packets-discarded-count

Description

This command includes the in-plus profile packets discarded count.

The **no** form of this command excludes the in-plus profile packets discarded count.

Default

no in-plus-profile-packets-discarded-count

Platforms

7705 SAR Gen 2

13.59 in-plus-profile-packets-forwarded-count

in-plus-profile-packets-forwarded-count

Syntax

[**no**] **in-plus-profile-packets-forwarded-count**

Context

[Tree] (config>log>acct-policy>cr>policer>e-counters in-plus-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters in-plus-profile-packets-forwarded-count)

Full Context

configure log accounting-policy custom-record policer e-counters in-plus-profile-packets-forwarded-count

configure log accounting-policy custom-record ref-policer e-counters in-plus-profile-packets-forwarded-count

Description

This command includes the in-plus profile packets forwarded count.

The **no** form of this command excludes the in-plus profile packets forwarded count.

Default

no in-plus-profile-packets-forwarded-count

Platforms

7705 SAR Gen 2

13.60 in-plus-profile-packets-offered-count

in-plus-profile-packets-offered-count

Syntax

[**no**] **in-plus-profile-packets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters in-plus-profile-packets-offered-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters in-plus-profile-packets-offered-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters in-plus-profile-packets-offered-count

configure log accounting-policy custom-record policer e-counters in-plus-profile-packets-offered-count

Description

This command includes the in-plus profile packets offered count.

The **no** form of this command excludes the in-plus profile packets offered count.

Default

no in-plus-profile-packets-offered-count

Platforms

7705 SAR Gen 2

13.61 in-profile-octets-discarded-count

in-profile-octets-discarded-count

Syntax

[**no**] **in-profile-octets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters in-profile-octets-discarded-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters in-profile-octets-discarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>e-counters in-profile-octets-discarded-count)

[Tree] (config>log>acct-policy>cr>queue>e-counters in-profile-octets-discarded-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters in-profile-octets-discarded-count

configure log accounting-policy custom-record policer e-counters in-profile-octets-discarded-count

configure log accounting-policy custom-record ref-queue e-counters in-profile-octets-discarded-count

configure log accounting-policy custom-record queue e-counters in-profile-octets-discarded-count

Description

This command includes the in-profile octets discarded count.

The **no** form of this command excludes the in-profile octets discarded count.

Default

no in-profile-octets-discarded-count

Platforms

7705 SAR Gen 2

in-profile-octets-discarded-count

Syntax

[**no**] **in-profile-octets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters in-profile-octets-discarded-count)

[Tree] (config>log>acct-policy>cr>policer>i-counters in-profile-octets-discarded-count)

Full Context

configure log accounting-policy custom-record ref-policer i-counters in-profile-octets-discarded-count

configure log accounting-policy custom-record policer i-counters in-profile-octets-discarded-count

Description

This command includes the in-profile octets discarded count.

The **no** form of this command excludes the in-profile octets discarded count.

Default

no in-profile-octets-discarded-count

Platforms

7705 SAR Gen 2

13.62 in-profile-octets-forwarded-count

in-profile-octets-forwarded-count

Syntax

[**no**] **in-profile-octets-forwarded-count**

Context

[Tree] (config>log>acct-policy>cr>queue>e-counters in-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters in-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters in-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>e-counters in-profile-octets-forwarded-count)

Full Context

configure log accounting-policy custom-record queue e-counters in-profile-octets-forwarded-count

configure log accounting-policy custom-record policer e-counters in-profile-octets-forwarded-count

configure log accounting-policy custom-record ref-policer e-counters in-profile-octets-forwarded-count

configure log accounting-policy custom-record ref-queue e-counters in-profile-octets-forwarded-count

Description

This command includes the in-profile octets forwarded count.

The **no** form of this command excludes the in-profile octets forwarded count.

Default

no in-profile-octets-forwarded-count

Platforms

7705 SAR Gen 2

in-profile-octets-forwarded-count

Syntax

[**no**] **in-profile-octets-forwarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters in-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>policer>i-counters in-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>queue>i-counters in-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters in-profile-octets-forwarded-count)

Full Context

configure log accounting-policy custom-record ref-queue i-counters in-profile-octets-forwarded-count

configure log accounting-policy custom-record policer i-counters in-profile-octets-forwarded-count

configure log accounting-policy custom-record queue i-counters in-profile-octets-forwarded-count

configure log accounting-policy custom-record ref-policer i-counters in-profile-octets-forwarded-count

Description

This command includes the in profile octets forwarded count.

The **no** form of this command excludes the in profile octets forwarded count.

Default

no in-profile-octets-forwarded-count

Platforms

7705 SAR Gen 2

13.63 in-profile-octets-offered-count

in-profile-octets-offered-count

Syntax

[**no**] **in-profile-octets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters in-profile-octets-offered-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters in-profile-octets-offered-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters in-profile-octets-offered-count

configure log accounting-policy custom-record policer e-counters in-profile-octets-offered-count

Description

This command includes the in profile octets offered count.

The **no** form of this command excludes the in-profile octets offered count.

Default

no in-profile-octets-offered-count

Platforms

7705 SAR Gen 2

in-profile-octets-offered-count

Syntax

[**no**] **in-profile-octets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>policer>i-counters in-profile-octets-offered-count)

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters in-profile-octets-offered-count)

Full Context

configure log accounting-policy custom-record policer i-counters in-profile-octets-offered-count

configure log accounting-policy custom-record ref-policer i-counters in-profile-octets-offered-count

Description

This command includes the in-profile octets offered count.

The **no** form of this command excludes the in-profile octets offered count.

Default

no in-profile-octets-offered-count

Platforms

7705 SAR Gen 2

13.64 in-profile-packets-discarded-count

in-profile-packets-discarded-count

Syntax

[**no**] **in-profile-packets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters in-profile-packets-discarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>e-counters in-profile-packets-discarded-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters in-profile-packets-discarded-count)

[Tree] (config>log>acct-policy>cr>queue>e-counters in-profile-packets-discarded-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters in-profile-packets-discarded-count

configure log accounting-policy custom-record ref-queue e-counters in-profile-packets-discarded-count

configure log accounting-policy custom-record policer e-counters in-profile-packets-discarded-count

configure log accounting-policy custom-record queue e-counters in-profile-packets-discarded-count

Description

This command includes the in-profile packets discarded count.

The **no** form of this command excludes the in-profile packets discarded count.

Default

no in-profile-packets-discarded-count

Platforms

7705 SAR Gen 2

in-profile-packets-discarded-count

Syntax

[**no**] **in-profile-packets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters in-profile-packets-discarded-count)

[Tree] (config>log>acct-policy>cr>policer>i-counters in-profile-packets-discarded-count)

Full Context

configure log accounting-policy custom-record ref-policer i-counters in-profile-packets-discarded-count

configure log accounting-policy custom-record policer i-counters in-profile-packets-discarded-count

Description

This command includes the in-profile packets discarded count.

The **no** form of this command excludes the in-profile packets discarded count.

Default

no in-profile-packets-discarded-count

Platforms

7705 SAR Gen 2

13.65 in-profile-packets-forwarded-count

in-profile-packets-forwarded-count

Syntax

[**no**] **in-profile-packets-forwarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters in-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>queue>e-counters in-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters in-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>e-counters in-profile-packets-forwarded-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters in-profile-packets-forwarded-count

configure log accounting-policy custom-record queue e-counters in-profile-packets-forwarded-count

configure log accounting-policy custom-record policer e-counters in-profile-packets-forwarded-count

configure log accounting-policy custom-record ref-queue e-counters in-profile-packets-forwarded-count

Description

This command includes the in-profile packets forwarded count.

The **no** form of this command excludes the in-profile packets forwarded count.

Default

no in-profile-packets-forwarded-count

Platforms

7705 SAR Gen 2

in-profile-packets-forwarded-count

Syntax

[**no**] **in-profile-packets-forwarded-count**

Context

[Tree] (config>log>acct-policy>cr>policer>i-counters in-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>queue>i-counters in-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters in-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters in-profile-packets-forwarded-count)

Full Context

configure log accounting-policy custom-record policer i-counters in-profile-packets-forwarded-count

configure log accounting-policy custom-record queue i-counters in-profile-packets-forwarded-count

configure log accounting-policy custom-record ref-queue i-counters in-profile-packets-forwarded-count

configure log accounting-policy custom-record ref-policer i-counters in-profile-packets-forwarded-count

Description

This command includes the in profile packets forwarded count.

The **no** form of this command excludes the in profile packets forwarded count.

Default

no in-profile-packets-forwarded-count

Platforms

7705 SAR Gen 2

13.66 in-profile-packets-offered-count

in-profile-packets-offered-count

Syntax

[**no**] **in-profile-packets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters in-profile-packets-offered-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters in-profile-packets-offered-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters in-profile-packets-offered-count

configure log accounting-policy custom-record policer e-counters in-profile-packets-offered-count

Description

This command includes the in profile packets offered count.

The **no** form of this command excludes the in profile packets offered count.

Default

no in-profile-packets-offered-count

Platforms

7705 SAR Gen 2

in-profile-packets-offered-count

Syntax

[**no**] **in-profile-packets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>policer>i-counters in-profile-packets-offered-count)

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters in-profile-packets-offered-count)

Full Context

configure log accounting-policy custom-record policer i-counters in-profile-packets-offered-count

configure log accounting-policy custom-record ref-policer i-counters in-profile-packets-offered-count

Description

This command includes the in-profile packets offered count.

The **no** form of this command excludes the in-profile packets offered count.

Default

no in-profile-packets-offered-count

Platforms

7705 SAR Gen 2

13.67 in-remark

in-remark

Syntax

in-remark {**dscp** *dscp-name* | **prec** *ip-prec-value*}

no in-remark

Context

[Tree] (config>qos>sap-ingress>fc in-remark)

Full Context

configure qos sap-ingress fc in-remark

Description

This command is used in a SAP ingress QoS policy to define an explicit in-profile remark action for a forwarding class or subclass. While the SAP ingress QoS policy may be applied to any SAP, the remarking functions are only enforced when the SAP is associated with an IP or subscriber interface (in an IES or VPRN). When the policy is applied to a Layer 2 SAP (i.e., Epipe or VPLS), the remarking definitions are silently ignored.

In the case where the policy is applied to a Layer 3 SAP, the in-profile remarking definition will be applied to packets that have been classified to the forwarding class or subclass. It is possible for a packet to match a classification command that maps the packet to a particular forwarding class or subclass, only to have a more explicit (higher priority match) override the association. Only the highest priority match forwarding class or subclass association will drive the in-profile marking.

The in-remark command is only applicable to ingress IP routed packets that are considered in-profile. The profile of a SAP ingress packet is affected by either the explicit in-profile/out-of-profile definitions or the ingress policing function applied to the packet. Table 43: Effect of In-Remark Command on Received SAP Ingress Packets shows the effect of the in-remark command on received SAP ingress packets. Within the in-profile IP packet's ToS field, either the six DSCP bits or the three precedence bits are remarked.

Table 43: Effect of In-Remark Command on Received SAP Ingress Packets

SAP Ingress Packet State	in-remark Command Effect
Non-Routed, Policed In-Profile	No Effect (non-routed packet)
Non-Routed, Policed Out-of-Profile	No Effect (non-routed packet)
Non-Routed, Explicit In-Profile	No Effect (non-routed packet)
Non-Routed, Explicit Out-of-Profile	No Effect (non-routed packet)
IP Routed, Policed In-Profile	in-remark value applied to IP header ToS field
IP Routed, Policed Out-of-Profile	No Effect (out-of-profile packet)
IP Routed, Explicit In-Profile	in-remark value applied to IP header ToS field
IP Routed, Explicit Out-of-Profile	No Effect (out-of-profile packet)

The **no** form of this command disables ingress remarking of in-profile packets classified to the forwarding class or subclass.

Parameters

dscp *dscp-name*

Specifies that the matching packet's DSCP bits should be overridden with the value represented by dscp-name.

The *dscp-name* parameter is a 6-bit value. It must be one of the predefined DSCP names defined on the system.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

prec *ip-prec-value*

Specifies that the matching packet's precedence bits should be overridden with the value represented by *ip-prec-value*.

Values 0 to 7

Platforms

7705 SAR Gen 2

13.68 inactivity-timeout

inactivity-timeout

Syntax

inactivity-timeout *seconds*

no inactivity-timeout

Context

[Tree] (config>test-oam>twamp>server inactivity-timeout)

Full Context

configure test-oam twamp server inactivity-timeout

Description

This command configures the inactivity time out for all TWAMP-control connections. If no TWAMP control message is exchanged over the TCP connection for this duration of time the connection is closed and all in-progress tests are terminated.

The **no** form of this command returns the value to the default.

Default

inactivity-timeout 900

Parameters

seconds

Specifies the duration of the inactivity time out.

Values 60 to 3600

Default 900

Platforms

7705 SAR Gen 2

inactivity-timeout

Syntax

inactivity-timeout *seconds*

no inactivity-timeout

Context

[Tree] (config>test-oam>twamp>twamp-light inactivity-timeout)

Full Context

configure test-oam twamp twamp-light inactivity-timeout

Description

This command configures the length of time to maintain stale state on the session reflector. Stale state is test data that has not been refreshed or updated by newly arriving probes for that specific test in a predetermined length of time. Any single reflector can maintain up state for a maximum of 12000 tests. If the maximum value is exceeded, the session reflector lacks memory to allocate to new tests.

The **no** form of this command returns the value to the default.

Default

inactivity-timeout 100

Parameters

seconds

Specifies the value in seconds for maintaining stale state.

Values 10 to 100

Default 100

Platforms

7705 SAR Gen 2

13.69 inband

inband

Syntax

inband *service-id*

no inband

Context

[Tree] (config>system>security>vprn-aaa-server inband)

Full Context

configure system security vprn-aaa-server inband

Description

This command configures TACACS+ or RADIUS servers in a VPRN to be used for AAA by that VPRN and by sessions in the Base routing instance.

The **no** form of this command disables the use of servers for in-band management.

Default

no inband

Parameters

service-id

Specifies the VPRN server for AAA to use for in-band sessions.

Values *service-id*: 1 to 2147483648

svc-name: 64 characters maximum

Platforms

7705 SAR Gen 2

13.70 inbound-max-sessions

inbound-max-sessions

Syntax

inbound-max-sessions *number-of-sessions*

no inbound-max-sessions

Context

[Tree] (config>system>login-control>ftp inbound-max-sessions)

Full Context

configure system login-control ftp inbound-max-sessions

Description

This command configures the maximum number of concurrent inbound FTP sessions.

This value is the combined total of inbound and outbound sessions.

The **no** form of this command reverts to the default value.

Default

inbound-max-sessions 3

Parameters

value

Specifies the maximum number of concurrent FTP sessions on the node.

Values 0 to 5

Platforms

7705 SAR Gen 2

inbound-max-sessions

Syntax

inbound-max-sessions *number-of-sessions*

no inbound-max-sessions

Context

[Tree] (config>system>login-control>telnet inbound-max-sessions)

[Tree] (config>system>login-control>ssh inbound-max-sessions)

Full Context

configure system login-control telnet inbound-max-sessions

configure system login-control ssh inbound-max-sessions

Description

This parameter limits the number of inbound Telnet and SSH sessions. A maximum of 30 telnet and ssh connections can be established to the router. The local serial port cannot be disabled.

Telnet and SSH maximum sessions can also use the combined total of both inbound sessions (SSH +Telnet). While it is acceptable to continue to internally limit the combined total of SSH and Telnet sessions to N, either SSH or Telnet sessions can use the inbound maximum sessions, if so required by the Operator.

The **no** form of this command reverts to the default value.

Default

inbound-max-sessions 5

Parameters

number-of-sessions

The maximum number of concurrent inbound Telnet sessions, expressed as an integer.

Values 0 **to** 50 (default = 5) or 0 **to** N where N is the new total number of SSH +Telnet sessions if they are scaled

Platforms

7705 SAR Gen 2

13.71 incl-mcast-l2-attributes-advertisement

incl-mcast-l2-attributes-advertisement

Syntax

[**no**] **incl-mcast-l2-attributes-advertisement**

Context

[Tree] (config>service>vpls>bgp-evpn incl-mcast-l2-attributes-advertisement)

Full Context

configure service vpls bgp-evpn incl-mcast-l2-attributes-advertisement

Description

This command triggers the advertisement of the Layer 2 attributes extended community including:

- the service MTU in the Layer 2 MTU field
- the F bit, which is set to 1 if the **hash-label** command is set to true (in the **configure service vpls bgp-evpn mpls** context); otherwise, the F bit is set to 0
- the C bit, which is set to 1 if the **control-word** command is set to true (in the **configure service vpls bgp-evpn mpls** context); otherwise, the C bit is set to 0

The router compares the received Layer 2 MTU from a peer with the local service MTU. If there is a mismatch, the operation state of the EVPN destination is set to down, except if the **configure service vpls bgp-evpn ignore-mtu-mismatch** command is enabled.

A mismatch between the received C bit and the local **control-word** setting (in the **configure service vpls bgp-evpn mpls** context) results in the operational state of the EVPN destination being set to down.

A mismatch between the received F bit and the local F bit (via the hash label configuration) results in the operational state of the EVPN destination being set to down.

The **no** form of this command prevents the router from advertising the Layer 2 attributes extended community along with the IMET route for the service.

Default

no incl-mcast-l2-attributes-advertisement

Platforms

7705 SAR Gen 2

13.72 incl-mcast-orig-ip

incl-mcast-orig-ip

Syntax

incl-mcast-orig-ip *ip-address*

no incl-mcast-orig-ip

Context

[Tree] (config>service>vpls>bgp-evpn incl-mcast-orig-ip)

Full Context

configure service vpls bgp-evpn incl-mcast-orig-ip

Description

The IP address configured by the user in the **incl-mcast-orig-ip** command is encoded in the **originating-ip** field of EVPN Inclusive Multicast Routes with tunnel type Ingress Replication (value 6), mLDP (2), and Composite IR and mLDP (130).

The configured address does not need to be reachable in the base router or have an interface in the base router. The originating-ip address is used solely for BGP route-key selection.

The originating-ip is never changed for Inclusive Multicast Routes with tunnel type AR (Assisted Replication, value 10).

The **no** version of the command withdraws the affected Inclusive Multicast Routes and re-advertises it with the default system-ip address in the originating-ip field.

Default

incl-mcast-orig-ip 1

Parameters

ip-address

Specifies the IPv4 address value.

Values a.b.c.d

Platforms

7705 SAR Gen 2

13.73 include

include

Syntax

include *group-name* [*group-name*]

no include [*group-name* [*group-name*]]

Context

[Tree] (config>router>mpls>lsp>secondary include)

[Tree] (config>router>mpls>lsp include)

[Tree] (config>router>mpls>lsp>primary include)

[Tree] (config>router>mpls>lsp-template include)

Full Context

configure router mpls lsp secondary include

configure router mpls lsp include

configure router mpls lsp primary include

configure router mpls lsp-template include

Description

This command specifies the admin groups to be included when an LSP is set up. Up to five groups per operation can be specified, up to 32 maximum. The **include** statement instructs the CSPF algorithm to pick TE links among the links which belong to one or more of the specified admin groups. A link that does not belong to at least one of the specified admin groups is excluded and thus pruned from the TE database before the CSPF computation. However, a link can still be selected if it belongs to one of the groups in a **include** statement but also belongs to other groups which are not part of any **include** statement in the LSP or primary/secondary path configuration. In other words, the **include** statements implements the "include-any" behavior.

The **no** form of this command deletes the specified groups in the specified context.

Default

no include

Parameters

group-name

Specifies admin groups to be included when an LSP is set up.

Platforms

7705 SAR Gen 2

include

Syntax

[**no**] **include** *tag*

Context

[Tree] (config>router>admin-tags>route-admin-tag-policy include)

Full Context

configure router admin-tags route-admin-tag-policy include

Description

This configures an admin tag to be included when matching a route against an LSP.

Up to eight inclusion statements are supported per policy.

The **no** form of this command removes the admin tag from the include statement.

Parameters

tag

Specifies the value of the admin tag, up to 32 characters.

Platforms

7705 SAR Gen 2

13.74 include-all

include-all

Syntax

include-all

Context

[Tree] (config>router>fad>flex-algo include-all)

Full Context

configure router flexible-algorithm-definitions flex-algo include-all

Description

Commands in this context configure administrative groups to include in the flexible algorithm topology graph. Administrative groups are attributes associated with a link and are generally referred to as link colors.

Flexible algorithms provide the possibility to restrict inclusion into the topology graph to links that have a pre-defined combination of associated administrative groups. The **include-all** command requires that all configured administrative groups must be present in a link before the link can be included in the topology graph.

Platforms

7705 SAR Gen 2

13.75 include-any

include-any

Syntax

include-any

Context

[Tree] (config>router>fad>flex-algo include-any)

Full Context

configure router flexible-algorithm-definitions flex-algo include-any

Description

Commands in this context configure administrative groups to include in the flexible algorithm topology graph. Administrative groups are attributes associated with a link and are generally referred to as link colors.

Flexible algorithms provide the possibility to restrict inclusion into the topology graph to links that have a pre-defined combination of associated administrative groups. The **include-any** command requires that one of the configured administrative groups must be present on a link before the link can be included in the topology graph.

Platforms

7705 SAR Gen 2

13.76 include-dns

include-dns

Syntax

[**no**] **include-dns**

Context

[Tree] (config>service>vprn>router-advert>if>dns-options include-dns)

Full Context

configure service vprn router-advertisement interface dns-options include-dns

Description

This command enables the Recursive DNS Server (RDNSS) Option in router advertisements. This must be enabled for each interface on which the RDNSS option is required in router advertisement messages.

The **no** form of this command disables the RDNSS option in router advertisements.

Default

include-dns

Platforms

7705 SAR Gen 2

include-dns

Syntax

[**no**] **include-dns**

Context

[Tree] (config>router>router-advert>if>dns-opt include-dns)

Full Context

configure router router-advertisement interface dns-options include-dns

Description

This command enables the Recursive DNS Server (RDNSS) Option in router advertisements. This must be enabled for each interface on which the RDNSS option is required in router advertisement messages.

The **no** form of this command disables the RDNSS option in router advertisements.

Default

include-dns

Platforms

7705 SAR Gen 2

13.77 include-group

include-group

Syntax

include-group *ip-admin-group-name* [**pref** *preference*]

no include-group *ip-admin-group-name*

Context

[Tree] (config>router>route-next-hop-policy>template include-group)

Full Context

configure router route-next-hop-policy template include-group

Description

This command configures the admin group constraint into the route next-hop policy template.

Each group is entered individually. The **include-group** statement instructs the LFA SPF selection algorithm to pick up a subset of LFA next-hops among the links which belong to one or more of the specified admin groups. A link which does not belong to at least one of the admin-groups is excluded. However, a link can still be selected if it belongs to one of the groups in a include-group statement but also belongs to other groups which are not part of any include-group statement in the route next-hop policy.

The **pref** option is used to provide a relative preference for the admin group to select. A lower preference value means that LFA SPF will first attempt to select a LFA backup next-hop which is a member of the corresponding admin group. If none is found, then the admin group with the next higher preference value is evaluated. If no preference is configured for a given admin group name, then it is supposed to be the least preferred, that is, numerically the highest preference value.

When evaluating multiple **include-group** statements within the same preference, any link which belongs to one or more of the included admin groups can be selected as an LFA next-hop. There is no relative preference based on how many of those included admin groups the link is a member of.

The **exclude-group** statement simply prunes all links belonging to the specified admin group before making the LFA backup next-hop selection for a prefix.

If the same group name is part of both include and exclude statements, the exclude statement will win. It other words, the exclude statement can be viewed as having an implicit preference value of 0.

The admin-group criteria are applied before running the LFA next-hop selection algorithm.

The **no** form deletes the admin group constraint from the route next-hop policy template.

Parameters

ip-admin-group-name

Specifies the name of the group, up to 32 characters.

preference

An integer specifying the relative preference of a group.

Values	1 to 255
Default	255

Platforms

7705 SAR Gen 2

13.78 include-radius-attribute

include-radius-attribute

Syntax

[**no**] **include-radius-attribute**

Context

[Tree] (config>ipsec>rad-auth-plcy include-radius-attribute)

[Tree] (config>ipsec>rad-acct-plcy include-radius-attribute)

Full Context

configure ipsec radius-authentication-policy include-radius-attribute

configure ipsec radius-accounting-policy include-radius-attribute

Description

Commands in this context specify the RADIUS attributes that the system should include into RADIUS Access-Request (for authentication) and Accounting-Request (for accounting) messages.

Platforms

7705 SAR Gen 2

13.79 include-system-info

include-system-info

Syntax

[**no**] **include-system-info**

Context

[Tree] (config>log>accounting-policy include-system-info)

Full Context

configure log accounting-policy include-system-info

Description

This command allows the operator to optionally include router information at the top of each accounting file generated for a given accounting policy.

The **no** form of this command configures the router to not include optional router information at the top of the file.

Default

no include-system-info

Platforms

7705 SAR Gen 2

13.80 incremental-spf-wait

incremental-spf-wait

Syntax

incremental-spf-wait *incremental-spf-wait*

no incremental-spf-wait

Context

[Tree] (config>router>ospf>timers incremental-spf-wait)

[Tree] (config>router>ospf3>timers incremental-spf-wait)

Full Context

configure router ospf timers incremental-spf-wait

configure router ospf3 timers incremental-spf-wait

Description

This command sets the delay before an incremental SPF calculation is performed when LSA types 3, 4, 5, or 7 are received. This allows multiple updates to be processed in the same SPF calculation. Type 1 or type 2 LSAs are considered a topology change and will always trigger a full SPF calculation.

The **no** form of this command resets the timer value back to the default value.

Note:
The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is ≥ 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Default

incremental-spf-wait 1000

Parameters

incremental-spf-wait

Specifies the OSPF incremental SPF calculation delay, in milliseconds.

Values 0 to 1000

Platforms

7705 SAR Gen 2

13.81 indirect

indirect

Syntax

[**no**] **indirect** *ip-address*

Context

[Tree] (config>service>vprn>static-route-entry indirect)

Full Context

configure service vprn static-route-entry indirect

Description

This command specifies that the route is indirect and specifies the next hop IP address used to reach the destination.

The configured *ip-addr*ess is not directly connected to a network configured on this node. The destination can be reached via multiple paths. The indirect address can only be resolved from dynamic routing protocol. Another static route cannot be used to resolve the indirect address.

The *ip-address* configured here can be either on the network side or the access side and is typically at least one hop away from this node.

Default

no indirect

Parameters

ip-address

The IP address of the IP interface.

Values

ipv4-address	a.b.c.d
ipv6-address	x:x:x:x:x:x:x:x-[interface]

Platforms

7705 SAR Gen 2

indirect

Syntax

[**no**] **indirect** *ip-address*

Context

[Tree] (config>router>static-route-entry indirect)

Full Context

configure router static-route-entry indirect

Description

This command specifies that the route is indirect and specifies the next hop IP address used to reach the destination.

The configured *ip-addr*ess is not directly connected to a network configured on this node. The destination can be reached via multiple paths. The indirect address can only be resolved from dynamic routing protocol. Another static route cannot be used to resolve the indirect address.

The *ip-address* configured here can be either on the network side or the access side and is typically at least one hop away from this node.

Default

no indirect

Parameters

ip-address

Specifies the IP address of the IP interface.

Values

ipv4-address	a.b.c.d
ipv6-address	x:x:x:x:x:x:x:x-[interface]

Platforms

7705 SAR Gen 2

13.82 info

info

Syntax

info [**detail**] [**objective**]

info [**detail**] [**objective**] **operational**

Context

[Tree] (info)

Full Context

info

Description

This command displays the running configuration for the configuration context where it is entered and all branches below that context level. It can be used in any branch under configure, but not with **configure** itself.

By default, the command only enters the configuration parameters that vary from the default values.

The **detail** keyword causes all configuration parameters to be displayed. The **include-dynamic** objective keyword includes configuration parameters from dynamic sources such as dynamic data services Python scripts. These dynamic configuration parameters are not saved in the configuration file.

The **operational** keyword is available in edit-cfg mode only, in which case the keyword is mandatory when using the **info** command.

Example:

```
A:ALA-48>config>router>if-attr# info
----------------------------------
            admin-group "green" value 15
            admin-group "red" value 25
            admin-group "yellow" value 20
A:ALA-48>config>router>mpls# info
----------------------------------
            interface "system"
            exit
            interface "to-104"
                admin-group "green"
                admin-group "red"
                admin-group "yellow"
                label-map 35
                    swap 36 nexthop 10.10.10.91
                    no shutdown
                exit
            exit
            path "secondary-path"
                hop 1 10.10.0.111  strict
                hop 2 10.10.0.222  strict
                hop 3 10.10.0.123  strict
```

```
                no shutdown
            exit
            path "to-NYC"
                hop 1 10.10.10.104  strict
                hop 2 10.10.0.210  strict
                no shutdown
            exit
            path "to-104"
                no shutdown
            exit
            lsp "to-104"
                to 10.10.10.104
                from 10.10.10.103
                rsvp-resv-style ff
                cspf
...
-----------------------------------------------
A:ALA-48>config>router>mpls#
A:ALA-48>config>router>mpls# info detail
-----------------------------------------------
            frr-object
            no resignal-timer
            interface "system"
                no admin-group
                no shutdown
            exit
            interface "to-104"
                admin-group "green"
                admin-group "red"
                admin-group "yellow"
                label-map 35
                    swap 36 nexthop 10.10.10.91
                    no shutdown
                exit
                no shutdown
            exit
            path "secondary-path"
                hop 1 10.10.0.111  strict
                hop 2 10.10.0.222  strict
                hop 3 10.10.0.123  strict
                no shutdown
            exit
            path "to-NYC"
                hop 1 10.10.10.104  strict
                hop 2 10.10.0.210  strict
                no shutdown
            exit
            path "to-104"
                no shutdown
            exit
            lsp "to-104"
                to 10.10.10.104
                from 10.10.10.103
                rsvp-resv-style ff
                adaptive
                cspf
                include "red"
                exclude "green"
                adspec
                fast-reroute one-to-one
                    no bandwidth
                    no hop-limit
                    node-protect
                exit
```

```
                hop-limit 10
                retry-limit 0
                retry-timer 30
                secondary "secondary-path"
                    no standby
                    no hop-limit
                    adaptive
                    no include
                    no exclude
                    record
                    record-label
                    bandwidth 50000
                    no shutdown
                exit
                primary "to-NYC"
                    hop-limit 50
                    adaptive
                    no include
                    no exclude
                    record
                    record-label
                    no bandwidth
                    no shutdown
                exit
                no shutdown
            exit
...
-------------------------------------------
A:ALA-48>config>router>mpls#
```

Parameters

detail

Displays all configuration parameters including parameters at their default values.

objective

Provides an output objective that controls the configuration parameters to be displayed.

Values **include-dynamic**: includes configuration parameters from dynamic sources such as dynamic data services Python scripts.

Platforms

7705 SAR Gen 2

13.83 info-output

info-output

Syntax

info-output

Context

[Tree] (config>system>management-interface>cli>md-cli>environment info-output)

Full Context

configure system management-interface cli md-cli environment info-output

Description

Commands in this context configure the elements that are displayed in the MD-CLI session.

Platforms

7705 SAR Gen 2

13.84 ingress

ingress

Syntax

ingress

Context

[Tree] (config>card>fp ingress)

Full Context

configure card fp ingress

Description

This command enables access to the ingress fp CLI context.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>service>vpls>sap ingress)

[Tree] (config>service>ies>if>spoke-sdp ingress)

[Tree] (config>service>vpls>mesh-sdp ingress)

[Tree] (config>service>ies>if>sap ingress)

[Tree] (config>service>vpls>spoke-sdp ingress)

Full Context

configure service vpls sap ingress

configure service ies interface spoke-sdp ingress

configure service vpls mesh-sdp ingress

configure service ies interface sap ingress

configure service vpls spoke-sdp ingress

Description

Commands in this context configure ingress Quality of Service (QoS) policies and filter policies.

If no QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>port>access ingress)

[Tree] (config>card>mda>access ingress)

Full Context

configure port access ingress

configure card mda access ingress

Description

Commands in this context configure ingress buffer pool parameters which define the percentage of the pool buffers that are used for CBS calculations and specify the slope policy that is configured in the **config>qos>slope-policy** context.

On the MDA level, access ingress pools are only allocated on channelized MDAs.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>port>ethernet>access ingress)

Full Context

configure port ethernet access ingress

Description

This command configures Ethernet access ingress port parameters.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>service>epipe>sap ingress)

Full Context

configure service epipe sap ingress

Description

Commands in this context configure ingress SAP Quality of Service (QoS) policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>service>epipe>spoke-sdp ingress)

Full Context

configure service epipe spoke-sdp ingress

Description

This command configures the ingress SDP context.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>service>ies>if>vpls ingress)

Full Context

configure service ies interface vpls ingress

Description

The ingress node in this context under the vpls binding is used to define the routed IPv4 and IPv6 optional filter overrides.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>service>ies>if ingress)

Full Context

configure service ies interface ingress

Description
This command enters context to configure ingress parameters for network interfaces.

Platforms
7705 SAR Gen 2

ingress

Syntax
ingress

Context
[Tree] (config>service>vprn>network ingress)

Full Context
configure service vprn network ingress

Description
Commands in this context configure network ingress parameters for the VPRN service.

Platforms
7705 SAR Gen 2

ingress

Syntax
ingress

Context
[Tree] (config>service>vprn>if ingress)

Full Context
configure service vprn interface ingress

Description
This command enters context to configure ingress parameters for network interfaces.

Platforms
7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>service>vprn>if>sap ingress)

Full Context

configure service vprn interface sap ingress

Description

Commands in this context configure ingress SAP Quality of Service (QoS) policies and filter policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>service>vprn>if>vpls ingress)

Full Context

configure service vprn interface vpls ingress

Description

The ingress node in this context under the vpls binding is used to define the routed IPv4 and IPv6 optional filter overrides.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>card>mda>network ingress)

Full Context

configure card mda network ingress

Description

Commands in this context configure MDA-level IOM Quality of Service (QoS).

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>service>vprn>ipmirrorif>spoke-sdp ingress)

[Tree] (config>mirror>mirror-dest>remote-src>spoke-sdp ingress)

Full Context

configure service vprn ip-mirror-interface spoke-sdp ingress

configure mirror mirror-dest remote-source spoke-sdp ingress

Description

Commands in this context configure spoke SDP ingress parameters.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>qos>network ingress)

Full Context

configure qos network ingress

Description

This command is used to enter the CLI node that creates or edits policy entries that specify the DiffServ code points-to-forwarding class mapping for all IP packets and define the MPLS EXP bits-to-forwarding class mapping for all labeled packets.

When premarked IP or MPLS packets ingress on a network port, they get a Per Hop Behavior (that is, the QoS treatment through the router, based on the mapping defined under the current node).

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>qos>queue-group-templates ingress)

Full Context

configure qos queue-group-templates ingress

Description

Commands in this context create ingress queue group templates. Ingress queue group templates can be applied to ingress ports to create an ingress queue group of the same name.

An ingress template must be created for a group-name prior to creating a queue group with the same name on an ingress port.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>router>if ingress)

Full Context

configure router interface ingress

Description

This command enables access to the context to configure ingress network filter policies for the IP interface. If an ingress filter is not defined, no filtering is performed.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>service>cust>multi-service-site ingress)

Full Context

configure service customer multi-service-site ingress

Description

Commands in this context configure the ingress node associate an existing scheduler policy name with the customer site. The ingress node is an entity to associate commands that complement the association.

Platforms

7705 SAR Gen 2

ingress

Syntax

ingress

Context

[Tree] (config>service>pw-template ingress)

Full Context

configure service pw-template ingress

Description

Commands in this context configure spoke SDP binding ingress filter parameters.

Platforms

7705 SAR Gen 2

13.85 ingress-rate

ingress-rate

Syntax

ingress-rate *sub-rate*

no ingress-rate

Context

[Tree] (config>port>ethernet ingress-rate)

Full Context

configure port ethernet ingress-rate

Description

This command configures the maximum amount of ingress bandwidth that this port can receive with the configured sub-rate using packet-based accounting.

The **no** form of this command returns the value to the default.

Default

no ingress-rate

Parameters

sub-rate

Specifies the ingress rate, in Mb/s.

Values 1 to 400000

Platforms

7705 SAR Gen 2

13.86 ingress-repl-inc-mcast-advertisement

ingress-repl-inc-mcast-advertisement

Syntax

[**no**] **ingress-repl-inc-mcast-advertisement**

Context

[Tree] (config>service>vpls>bgp-evpn ingress-repl-inc-mcast-advertisement)

Full Context

configure service vpls bgp-evpn ingress-repl-inc-mcast-advertisement

Description

This command enables and disables the advertisement of the Inclusive Multicast Ethernet Tag route (IMET route) with tunnel-type Ingress-Replication in the PMSI Tunnel Attribute, or with the tunnel-type Composite Point-to-Multipoint and Ingress-Replication (P2MP+IR) in the root-and-leaf nodes. The following must be considered:

- When **no ingress-repl-inc-mcast-advertisement** is configured, no IMET routes will be sent for the service unless the **provider-tunnel** is configured with **owner bgp-evpn-mpls** and **root-and-leaf**, in which case, an IMET-P2MP route is sent.
- When **ingress-repl-inc-mcast-advertisement** and **provider-tunnel** are configured for **bgp-evpn-mpls** with **root-and-leaf**, the system will send an IMET-P2MP-IR route, that is, an IMET route with a composite P2MP+IR tunnel type.
- When **no ingress-repl-inc-mcast-advertisement** and **assisted-replication replicator** are configured, the system will send IMET-AR routes, but IMET-IR routes will not be sent.

Default

ingress-repl-inc-mcast-advertisement

Platforms

7705 SAR Gen 2

13.87 ingress-replication-bum-label

ingress-replication-bum-label

Syntax

[**no**] **no-ingress-replication-bum-label**

Context

[Tree] (config>service>vpls>bgp-evpn>mpls ingress-replication-bum-label)

Full Context

configure service vpls bgp-evpn mpls ingress-replication-bum-label

Description

This command allows the user to configure the system so that a separate label is sent for BUM (Broadcast, Unknown unicast and Multicast) traffic in a specified service. By default (**no ingress-replication-bum-label**), the same label is used for unicast and flooded BUM packets when for-warding traffic to remote PEs.

When saving labels, this might cause transient traffic duplication for all-active multi-homing. By enabling **ingress-replication-bum-label**, the system will advertise two labels per EVPN VPLS instance, one for unicast and one for BUM traffic. The ingress PE will use the BUM label for flooded traffic to the advertising egress PE, so that the egress PE can determine if the unicast traffic has been flooded by the ingress PE. Depending on the scale required in the network, the user may choose between saving label space or avoiding transient packet duplication sent to an all-active multi-homed CE for certain macs.

Default

no ingress-replication-bum-label

Platforms

7705 SAR Gen 2

13.88 ingress-statistics

ingress-statistics

Syntax

[**no**] **ingress-statistics**

Context

[Tree] (config>router>mpls>fwd-policies>fwd-policy ingress-statistics)

Full Context

configure router mpls forwarding-policies forwarding-policy ingress-statistics

Description

This command configures ingress statistics in an MPLS forwarding policy.

The ingress statistics are associated with a binding label, that is the ILM of the forwarding policy, and provides aggregate packet and byte counters for packets matching the binding label.

The **no** form of this command removes the statistics from the MPLS forwarding policy.

Platforms

7705 SAR Gen 2

13.89 init

init

Syntax

init [**detail**]

no init

Context

[Tree] (debug>router>ldp>peer>packet init)

Full Context

debug router ldp peer packet init

Description

This command enables debugging for LDP Init packets.

The **no** form of the command disables the debugging output.

Parameters

detail

Displays detailed information.

Platforms

7705 SAR Gen 2

13.90 init-delay

init-delay

Syntax

init-delay *seconds*

no init-delay

Context

[Tree] (config>service>ies>if>ipv6>vrrp init-delay)

Full Context

configure service ies interface ipv6 vrrp init-delay

Description

This command configures a VRRP initialization delay timer.

Default

no init-delay

Parameters

seconds

Specifies the initialization delay timer for VRRP, in seconds.

Values 1 to 65535

Platforms

7705 SAR Gen 2

init-delay

Syntax

init-delay *seconds*

no init-delay

Context

[Tree] (config>service>ies>if>vrrp init-delay)

Full Context

configure service ies interface vrrp init-delay

Description

This command configures a VRRP initialization delay timer.

Default

no init-delay

Parameters

seconds

Specifies the initialization delay timer for VRRP, in seconds.

Values 1 to 65535

Platforms

7705 SAR Gen 2

init-delay

Syntax

init-delay *seconds*

no init-delay

Context

[Tree] (config>service>vprn>if>vrrp init-delay)

[Tree] (config>service>vprn>if>ipv6>vrrp init-delay)

Full Context

configure service vprn interface vrrp init-delay

configure service vprn interface ipv6 vrrp init-delay

Description

This command configures a VRRP initialization delay timer.

Default

no init-delay

Parameters

seconds

Specifies the initialization delay timer for VRRP, in seconds.

Values 1 to 65535

Platforms

7705 SAR Gen 2

init-delay

Syntax

init-delay *seconds*

no init-delay

Context

[Tree] (config>router>if>vrrp init-delay)

[Tree] (config>router>if>ipv6>vrrp init-delay)

Full Context

configure router interface vrrp init-delay

configure router interface ipv6 vrrp init-delay

Description

This command configures a VRRP initialization delay timer.

Default

no init-delay

Parameters

seconds

Specifies the initialization delay timer for VRRP, in seconds.

Values 1 to 65535

Platforms

7705 SAR Gen 2

13.91 init-extract-prio-mode

init-extract-prio-mode

Syntax

init-extract-prio-mode {**uniform** | **l3-classify**}

Context

[Tree] (config>card>fp init-extract-prio-mode)

Full Context

configure card fp init-extract-prio-mode

Description

This command determines the scheme used to select the initial drop priority of extracted control plane traffic. The initial drop priority of extracted packets can be either low or high priority. The drop priority of the extracted packets can be subsequently altered by mechanisms such as CPU protection. High-priority traffic receives preferential treatment in control plane congestion situations over low-priority traffic.

Default

init-extract-prio-mode uniform

Parameters

uniform

Initializes the drop priority of all extracted control traffic as high priority. Drop priority can then be altered (marked low priority) by distributed CPU protection (DCP) or centralized CPU protection rate-limiting functions in order to achieve protocol and interface isolation.

l3-classify

Initializes the drop priority of Layer 3 extracted control traffic (BGP and OSPF) based on the QoS classification of the packets. This is useful in networks where the DSCP and EXP markings can be trusted as the primary method to distinguish, protect, and isolate good terminating protocol traffic from unknown or potentially harmful protocol traffic instead of using the rate-based DCP and centralized CPU protection traffic marking/coloring mechanisms (for example, **out-profile-rate** and **exceed-action low-priority**).

For network interfaces, the QoS classification profile result selects the drop priority (in = high priority, out = low priority) for extracted control traffic, and the default QoS classification maps different DSCP and EXP values to different in/out profile states.

For access interfaces, the QoS classification priority result typically selects the drop priority for extracted control traffic. The default access QoS classification (**default-priority**) maps all traffic to **low**. If the queues in the access QoS policy are configured as **profile-mode** queues (rather than the default **priority-mode**) extracted traffic will use the QoS classification profile value configured against the associated FC (rather than the priority result) to select the drop priority.

Layer 2 extracted control traffic (ARP or ETH-CFM) and protocols that cannot always be QoS-classified, such as IS-IS, are initialized as low drop priority in order to protect Layer 2 protocol traffic on uniform interfaces (which would typically be subject to centralized CPU protection). Alternately, DCP can be used (by configuring a non-zero rate with **exceed-action** of **low-priority** for the **all-unspecified** protocol) to mark some of this traffic as high priority.

Platforms

7705 SAR Gen 2

13.92 initial-registration

initial-registration

Syntax

initial-registration ca *ca-profile-name* **key-to-certify** *key-filename* **protection-alg** {**password** *password* **reference** *ref-number* | **signature** [**cert** *cert-file-name* [**send-chain** [**with-ca** *ca-profile-name*]]] [**protection-key** *key-file-name*] [**hash-alg** {**md5** | **sha1** | **sha224** | **sha256** | **sha384** | **sha512**}]} **subject-dn** *dn* [**domain-name** *domain-names*] [**ip-addr** *ip-address* | *ipv6-address*] **save-as** *save-path-of-result-cert*

Context

[Tree] (admin>certificate>cmpv2 initial-registration)

Full Context

admin certificate cmpv2 initial-registration

Description

This command request initial certificate from CA by using CMPv2 initial registration procedure.

The **ca** parameter specifies a CA-profile which includes CMP server information.

The **key-to-certify** is an imported key file to be certified by the CA.

The protection-key is an imported key file used to for message protection if protection-alg is signature.

The request is authenticated either of following methods:

- A password and a reference number that pre-distributed by CA via out-of-band means.
- The specified password and reference number are not necessarily in the cmp-keylist configured in the corresponding CA-Profile
- A signature signed by the protection-key or key-to-certify, optionally along with the corresponding certificate. If the protection-key is not specified, system will use the key-to-certify for message protection. The hash algorithm used for signature is depends on key type:
- DSA key: SHA1
- RSA key: MD5/SHA1/SHA224 | SHA256 | SHA384 | SHA512, by default is SHA1

Optionally, the system could also send a certificate or a chain of certificates in extraCerts field. Certificate is specified by the "cert” parameter, it must include the public key of the key used for message protection.

Sending a chain is enabled by specify the **send-chain** parameter.

subject-dn specifies the subject of the requesting certificate.

save-as specifies full path name of saving the result certificate.

In some cases, CA may not return certificate immediately, due to reason like request processing need manual intervention. In such cases, the **admin certificate cmpv2** poll command could be used to poll the status of the request. If key-list is not configured in the corresponding **ca-profile**, then the system will use the existing password to authenticate the CMPv2 packets from server if it is in password protection.

If key-list is configured in the corresponding **ca-profile** and server does not send SenderKID, then the system will use lexicographical first key in the key-list to authenticate the CMPv2 packets from server in case it is in password protection.

Parameters

ca-profile-name

Specifies a ca-profile name which includes CMP server information up to 32 characters.

key-filename

Specifies the file name of the key to certify up to 95 characters.

password

Specifies an ASCII string up to 64 characters.

ref-number

Specifies the reference number for this CA initial authentication key up to 64 characters.

cert-file-name

specifies the certificate file up to 95 characters.

ca-profile-name

Specifies to send the chain.

key-file-name

Specifies the protection key associated with the action on the CA profile.

hash-algorithm

Specifies the hash algorithm for RSA key.

Values md5,sha1,sha224,sha256,sha384,sha512

dn

Specifies the subject of the requesting certificate up to 256 characters.

Values attr1 equals val1

attr2 equals val2 where: attrN equals {C | ST | O | OU | CN}

save-path-of-result-cert

Specifies the save full path name of saving the result certificate up to 200 characters.

domain-name *domain-names*

Specifies FQDNs for SubjectAltName of the requesting certificate, separated by commas, up to 512 characters.

ip-address* | *ipv6-address

Specifies an IPv4 or IPv6 address for SubjectAtName of the requesting certificate.

Platforms

7705 SAR Gen 2

13.93 initial-send-delay-zero

initial-send-delay-zero

Syntax

[**no**] **initial-send-delay-zero**

Context

[Tree] (config>service>vprn>bgp>group initial-send-delay-zero)

[Tree] (config>service>vprn>bgp>group>neighbor initial-send-delay-zero)

[Tree] (config>service>vprn>bgp initial-send-delay-zero)

Full Context

configure service vprn bgp group initial-send-delay-zero

configure service vprn bgp group neighbor initial-send-delay-zero

configure service vprn bgp initial-send-delay-zero

Description

This command configures BGP to send UPDATE messages announcing reachability information to a peer or set of peers immediately after the sessions come up (become established) with these peers.

The default behavior, provided by the **no** form of this command, is to wait for min-route-advertisement time after each session is established before sending the first set of UPDATE messages.

Platforms

7705 SAR Gen 2

initial-send-delay-zero

Syntax

[**no**] **initial-send-delay-zero**

Context

[Tree] (config>router>bgp initial-send-delay-zero)

[Tree] (config>router>bgp>group initial-send-delay-zero)

[Tree] (config>router>bgp>group>neighbor initial-send-delay-zero)

Full Context

configure router bgp initial-send-delay-zero

configure router bgp group initial-send-delay-zero

configure router bgp group neighbor initial-send-delay-zero

Description

This command configures BGP to send UPDATE messages announcing reachability information to a peer or set of peers immediately after the sessions become established with these peers.

The **no** form of this command waits for **min-route-advertisement** time after each session is established before sending the first set of UPDATE messages.

Platforms

7705 SAR Gen 2

13.94 inner-tag

inner-tag

Syntax

inner-tag *value* [*vid-mask*]

no inner-tag

Context

[Tree] (config>qos>sap-ingress>mac-criteria>entry>match inner-tag)

Full Context

configure qos sap-ingress mac-criteria entry match inner-tag

Description

This command configures the matching of the second tag that is carried transparently through the service. The inner tag on ingress is the second tag on the frame if there are no service delimiting tags. The inner tag is the second tag before any service delimiting tags on egress but is dependent in the ingress configuration and may be set to 0 even in cases where additional tags are on the frame. This allows matching VLAN tags for explicit filtering or QoS setting when using default or null encapsulations.

The inner tag is not applicable in ingress on dot1Q SAPs. The inner tag may be populated on egress depending on the ingress SAP type.

On QinQ SAPs of null and default that do not strip tags, the inner-tag will contain the second tag (which is still the second tag carried transparently through the service.) On ingress SAPs that strip any tags, the inner tag will contain 0 even if there are more than two tags on the frame.

The optional vid_mask is defaulted to 4095 (exact match) but may be specified to allow pattern matching. The masking operation is ((value and vid-mask) = = (tag and vid-mask)). A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6.

For QoS, the VID type cannot be specified on the default QoS policy.

The default vid-mask is set to 4095 for exact match.

Platforms

7705 SAR Gen 2

13.95 insert

insert

Syntax

insert [*line*]

Context

[Tree] (candidate insert)

Full Context

candidate insert

Description

This command inserts the contents of the temporary buffer (populated by a previous copy or delete command) into the candidate configuration. The contents are inserted by default after the current edit point. Optional parameters allow the insertion after some other point of the candidate. The contents of the temporary buffer are deleted when the operator exits candidate edit mode.

Insertions are context-aware. The temporary buffer always stores the CLI context (such as the current CLI branch) for each line deleted or copied. If the lines to be inserted are supported at the context of the insertion point then the lines are simply inserted into the configuration. If the lines to be inserted are not supported at the context of the insertion point, then the context at the insertion point is first closed using multiple exit statements, the context of the lines to be inserted is built (added) into the candidate at the insertion point, then the lines themselves are added, the context of the inserted lines is closed using exit statements and finally the context from the original insertion point is built again leaving the context at the same point as it was before the insertion.

Parameters

line

Indicates where to insert the line starting at the point indicated by the following options.

Values line, offset, **first**, **edit-point**, **last**

line	absolute line number
offset	relative line number to current edit point. Prefixed with '+' or '-'
first	keyword - first line
edit-point	keyword - current edit point
last	keyword - last line that is not 'exit'

Platforms
7705 SAR Gen 2

13.96 inside

inside

Syntax
inside

Context
[Tree] (config>service>vprn>nat inside)
[Tree] (config>router>nat inside)

Full Context
configure service vprn nat inside
configure router nat inside

Description
Commands in this context the inside NAT instance.

Platforms
7705 SAR Gen 2

13.97 install-backup-path

install-backup-path

Syntax
install-backup-path
no install-backup-path

Context
[Tree] (config>router>policy-options>policy-statement>default-action install-backup-path)
[Tree] (config>router>policy-options>policy-statement>entry>action install-backup-path)

Full Context
configure router policy-options policy-statement default-action install-backup-path

configure router policy-options policy-statement entry action install-backup-path

Description

When the best BGP route for an IPv4 or IPv6 prefix is matched by a policy entry or policy default action with this command, BGP attempts to find and install a preprogrammed backup path for the prefix in order to provide BGP fast reroute protection.

The **install-backup-path** command overrides and has no dependency on commands such as the BGP instance **backup-path** command or the VPRN-level **enable-bgp-vpn-backup** command, which enable BGP fast reroute for an entire address family. The **install-backup-path** command provides more precise control over which IP prefixes are supported with preprogrammed backup paths.

In VPRN, if the best path for an IP prefix is provided by a VPRN BGP route, the backup path can be provided by another VPRN BGP route or an imported VPN-IP route. If the best path for an IP prefix is provided by an imported VPN-IP route, the backup path can be provided by another VPN-IP route.

The **install-backup-path** command is supported only in BGP and VRF import policies and has no effect on other types. The **install-backup-path** command applies only to the following types of matched routes: IPv4, IPv6, label-IPv4, label-IPv6, VPN-IPv4, and VPN-IPv6.

The **no** form of this command disables the install-backup-path functionality.

Default

no install-backup-path

Platforms

7705 SAR Gen 2

13.98 instant-prune-echo

instant-prune-echo

Syntax

[**no**] **instant-prune-echo**

Context

[Tree] (config>service>vprn>pim>if instant-prune-echo)

Full Context

configure service vprn pim interface instant-prune-echo

Description

This command enables PIM to send an instant prune echo when the router starts the prune pending timer for a group on the interface. All downstream routers will see the prune message immediately, and can send a join override if they are interested in receiving the group. Configuring instant-prune-

echo is recommended on broadcast interfaces with more than one PIM neighbor to optimize multicast convergence.

The **no** form of this command disables instant Prune Echo on the PIM interface.

Default

no instant-prune-echo

Platforms

7705 SAR Gen 2

instant-prune-echo

Syntax

[**no**] **instant-prune-echo**

Context

[Tree] (config>router>pim>interface instant-prune-echo)

Full Context

configure router pim interface instant-prune-echo

Description

This command enables PIM to send an instant prune echo when the router starts the prune pending timer for a group on the interface. All downstream routers will see the prune message immediately, and can send a join override if they are interested in receiving the group. Configuring instant-prune-echo is recommended on broadcast interfaces with more than one PIM neighbor to optimize multicast convergence.

The **no** form of this command disables instant Prune Echo on the PIM interface.

Default

no instant-prune-echo

Platforms

7705 SAR Gen 2

13.99 interactive-authentication

interactive-authentication

Syntax

[**no**] **interactive-authentication**

Context

[Tree] (config>system>security>radius interactive-authentication)

[Tree] (config>service>vprn>aaa>rmt-srv>radius interactive-authentication)

Full Context

configure system security radius interactive-authentication

configure service vprn aaa remote-servers radius interactive-authentication

Description

This command enables RADIUS interactive authentication for the system. Enabling interactive-authentication forces RADIUS to fall into challenge/response mode.

Default

no interactive-authentication

Platforms

7705 SAR Gen 2

interactive-authentication

Syntax

[**no**] **interactive-authentication**

Context

[Tree] (config>system>security>tacplus interactive-authentication)

[Tree] (config>service>vprn>aaa>rmt-srv>tacplus interactive-authentication)

Full Context

configure system security tacplus interactive-authentication

configure service vprn aaa remote-servers tacplus interactive-authentication

Description

This configuration instructs the SR OS to send no username nor password in the TACACS+ start message, and to display the *server_msg* in the GETUSER and GETPASS response from the TACACS+ server. Interactive authentication can be used to support a One Time Password scheme (such as an S/Key). An example flow (such as with a telnet connection) is as follows:

- The SR OS sends an authentication start request to the TACACS+ server with no username nor password.
- TACACS+ server replies with TAC_PLUS_AUTHEN_STATUS_GETUSER and a *server_msg*.
- The SR OS displays the *server_msg*, and collects the username.
- The SR OS sends a continue message with the username.
- TACACS+ server replies with TAC_PLUS_AUTHEN_STATUS_GETPASS and a *server_msg*.

- The SR OS displays the *server_msg* (which may contain, for example, an S/Key for One Time Password operation), and collects the password.
- The SR OS sends a continue message with the password.
- TACACS+ server replies with PASS or FAIL.

When interactive-authentication is disabled, the SR OS will send the username and password in the *tacplus* start message. An example flow (such as with a telnet connection) is as follows:

- TAC_PLUS_AUTHEN_TYPE_ASCII.
  - the login username in the "user" field.
  - the password in the *user_msg* field (while this is non-standard, it does not cause interoperability problems).
- TACACS+ server ignores the password and replies with TAC_PLUS_AUTHEN_STATUS_GETPASS.
- The SR OS sends a continue packet with the password in the *user_msg* field.
- TACACS+ server replies with PASS or FAIL.

When interactive-authentication is enabled, tacplus must be the first method specified in the authentication-order configuration.

Default

no interactive-authentication

Platforms

7705 SAR Gen 2

13.100 interface

interface

Syntax

interface *ip-int-name* [**create**]

interface *ip-int-name* [**create**] **tunnel**

no interface *ip-int-name*

Context

[Tree] (config>service>vprn interface)

[Tree] (config>service>ies interface)

Full Context

configure service vprn interface

configure service ies interface

Description

This command creates a logical IP routing interface. Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.

The **interface** command, under the context of services, is used to create and maintain IP routing interfaces within service IDs. The **interface** command can be executed in the context of a service ID. The IP interface created is associated with the service core network routing instance and default routing table. The typical use for IP interfaces created in this manner is for subscriber Internet access. An IP address cannot be assigned to an IES interface. Multiple SAPs can be assigned to a single group interface.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config>router>interface**, **config>service>ies>interface** and **config>service>vprn>interface** (that is, the network core router instance). Interface names must not be in the dotted decimal notation of an IP address. For example, the name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The **no** form of this command removes IP the interface and all the associated configuration. The interface must be administratively shut down before issuing the **no interface** command.

The IP interface must be shut down before the SAP on that interface may be removed. IES and VPRN services do not have the **shutdown** command in the SAP CLI context. The service SAPs rely on the interface status to enable and disable them.

Parameters

ip-int-name

Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and config service vprn interface commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

If *ip-int-name* already exists within the service ID, the context will be changed to maintain that IP interface. If *ip-int-name* already exists within another service ID or is an IP interface defined within the **config router** commands, an error will occur and context will not be changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

tunnel

Specifies that the interface is configured as tunnel interface, which could be used to terminate IPsec or GRE runnels in the private service.

create

Creates the IPsec interface instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>router>igmp interface)

Full Context

configure router igmp interface

Description

Commands in this context configure an IGMP interface. The interface is a local identifier of the network interface on which reception of the specified multicast address is to be enabled or disabled.

The **no** form of the command deletes the IGMP interface. The **shutdown** command in the **config>router>igmp>interface** context can be used to disable an interface without removing the configuration for the interface.

Default

no interface

Parameters

ip-int-name

The IP interface name. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

If the IP interface name does not exist or does not have an IP address configured an error message will be returned.

If the IP interface exists in a different area it will be moved to this area.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>router>mld interface)

Full Context

configure router mld interface

Description

Commands in this context configure an Multicast Listener Discovery (MLD) interface. The interface is a local identifier of the network interface on which reception of the specified multicast address is to be enabled or disabled.

The **no** form of this command deletes the MLD interface. The **shutdown** command in the **config>router>mld>interface** context can be used to disable an interface without removing the configuration for the interface.

Default

no interface — No interfaces are defined.

Parameters

ip-int-name

Specifies the IP interface name. Interface names must be unique within the group of defined IP interfaces for **config>router>interface** and **config>service>ies>interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

If the IP interface name does not exist or does not have an IP address configured an error message will be returned.

If the IP interface exists in a different area it will be moved to this area.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>service>vpls interface)

Full Context

configure service vpls interface

Description

This command creates a logical IP routing interface for a VPLS service. Once created, attributes such as IP address and service access points (SAP) can be associated with the IP interface.

The interface command, under the context of services, is used to create and maintain IP routing interfaces within the VPLS service IDs. The IP interface created is associated with the VPLS management routing instance. This instance does not support routing.

Interface names are case-sensitive and must be unique within the group of defined IP interfaces defined for the network core router instance. Interface names in the dotted decimal notation of an IP address are not allowed. For example, the name "1.1.1.1” is not allowed, but "int-1.1.1.1” is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. Duplicate interface names can exist in different router instances.

Enter a new name to create a logical router interface. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, no default IP interface names are defined within the system. All VPLS IP interfaces must be explicitly defined in an enabled state.

The no form of this command removes the IP interface and the entire associated configuration. The interface must be administratively shut down before issuing the no interface command.

For VPLS services, the IP interface must be shut down before the SAP on that interface is removed.

For VPLS service, ping and traceroute are the only applications supported.

Parameters

ip-int-name

Specifies the name of the IP interface. Interface names must be unique within the group of defined IP.

An interface name:

- Should not be in the form of an IP address.
- Can be from 1 to 32 alphanumeric characters.
- If the string contains special characters (such as #,$,spaces), the entire string must be enclosed within double quotes.

If ip-int-name already exists within the service ID, the context changes to maintain that IP interface. If ip-int-name already exists within another service ID, an error occurs and the context does not change to that IP interface. If ip-int-name does not exist, the interface is created and the context is changed to that interface for further command processing.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** [*ip-int-name* | *ip-address*]

Context

[Tree] (debug>router>igmp interface)

Full Context

debug router igmp interface

Description

This command enables debugging for IGMP interfaces.

The **no** form of this command disables the IGMP interface debugging for the specifies interface name or IP address.

Parameters

ip-int-name

Debugs the information associated with the specified IP interface name.

ip-address

Debugs the information associated with the specified IP address.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>service>vprn>igmp interface)

Full Context

configure service vprn igmp interface

Description

Commands in this context configure interface parameters.

Parameters

ip-int-name

Specifies the name of the IP interface, up to 32 characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>service>vprn>isis interface)

Full Context

configure service vprn isis interface

Description

This command creates the context to configure an IS-IS interface.

When an area is defined, the interfaces belong to that area. Interfaces cannot belong to separate areas.

When the interface is a POS channel, the OSI Network Layer Control Protocol (OSINLCP) is enabled when the interface is created and removed when the interface is deleted.

The **no** form of this command removes IS-IS from the interface.

The **shutdown** command in the **config>router>isis>if** context administratively disables IS-IS on the interface without affecting the IS-IS configuration.

Default

no interface — No IS-IS interfaces are defined.

Parameters

ip-int-name

Identify the IP interface name created in the **config>router>if** context. The IP interface name must already exist.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>service>vprn>mld interface)

Full Context

configure service vprn mld interface

Description

Commands in this context configure an Multicast Listener Discovery (MLD) interface. The interface is a local identifier of the network interface on which reception of the specified multicast address is to be enabled or disabled.

The **no** form of this command deletes the MLD interface. The **shutdown** command in the **config>router>mld>if** context can be used to disable an interface without removing the configuration for the interface.

Default

no interface

Parameters

ip-int-name

Specifies the IP interface name. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

If the IP interface name does not exist or does not have an IP address configured an error message will be returned.

If the IP interface exists in a different area it will be moved to this area.

Platforms

7705 SAR Gen 2

interface

Syntax

interface *ip-int-name* [**secondary**]

no interface *ip-int-name*

Context

[Tree] (config>service>vprn>ospf>area interface)

[Tree] (config>service>vprn>ospf3>area interface)

Full Context

configure service vprn ospf area interface

configure service vprn ospf3 area interface

Description

This command creates a context to configure an OSPF interface.

By default interfaces are not activated in any interior gateway protocol, such as OSPF, unless explicitly configured.

The **no** form of this command deletes the OSPF interface configuration for this interface. The **shutdown** command in the **config>router>ospf>if** context can be used to disable an interface without removing the configuration for the interface.

Default

no interface

Parameters

ip-int-name

Specifies the IP interface name. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service vprn interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

If the IP interface name does not exist or does not have an IP address configured an error message will be returned.

If the IP interface exists in a different area it will be moved to this area.

secondary

Keyword used to allow multiple secondary adjacencies, in addition to the primary adjacency, to be established over a single IP interface. This keyword can also be applied to the system interface and to loopback interfaces to allow them to participate in multiple areas, although no adjacencies are formed over these types of interfaces.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>service>vprn>pim interface)

Full Context

configure service vprn pim interface

Description

This command enables PIM on an interface and enables the context to configure interface-specific parameters. By default interfaces are activated in PIM based on the **apply-to** command, and do not have to be configured on an individual basis unless the default values must be changed.

The **no** form of this command deletes the PIM interface configuration for this interface. If the **apply-to** command parameter is configured, then the **no interface** form must be saved in the configuration to avoid automatic (re)creation after the next **apply-to** is executed as part of a reboot.

The **shutdown** command can be used to disable an interface without removing the configuration for the interface.

Default

Interfaces are activated in PIM based on the apply-to command.

Parameters

ip-int-name

Specifies the interface name. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>service>vprn>router-advertisement interface)

Full Context

configure service vprn router-advertisement interface

Description

This command configures router advertisement properties on a specific interface. The interface must already exist in the **config>router>if** context.

Default

No interfaces are configured by default.

Parameters

ip-int-name

Specifies the interface name. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name* [**dual-stack**]

Context

[Tree] (config>router>ldp>interface-parameters interface)

Full Context

configure router ldp interface-parameters interface

Description

This command enables LDP on the specified IP interface.

The **no** form of the command deletes the LDP interface and all configuration information associated with the LDP interface.

The LDP interface must be disabled using the **shutdown** command before it can be deleted.

The user can configure different parameters for IPv4 and IPv6 LDP interfaces by entering **ipv4** or **ipv6** as the next command.

Parameters

ip-int-name

Specifies the name of an existing interface. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

dual-stack

This optional keyword allows the user to explicitly indicate if this interface should create the IPv4 context automatically or not. With the introduction of LDP IPv6, the creation of the interface does not automatically mean it is to be used for IPv4 like with legacy IPv4 only LDP interface. Thus the dual-stack keyword is an indication to the system that user will manually enable the IPv4, IPv6, or the dual-stack IPv4/IPv6 contexts manually.

The following are some of the key points for this keyword:

- If the keyword is provided, then IPv4 interface context will not be created automatically. If it is not provided, the IPv4 interface context will be created like in the legacy single stack LDP IPv4 interface behavior.
- This new keyword will always show in a configuration.
- When entering an already configured interface, there is no need to provide the keyword, but it will be ignored if provided.
- When deleting a configured interface, the keyword will not be accepted in the **no** version of the **interface** command.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *interface-name family*

Context

[Tree] (debug>router>ldp interface)

Full Context

debug router ldp interface

Description

Use this command for debugging an LDP interface.

Parameters

interface-name

The name of an existing interface.

family

Specifies the family type.

Values ipv4, ipv6

Platforms

7705 SAR Gen 2

interface

Syntax

interface *ip-address* **srlg-group** *group-name* [*group-name*]

no interface *ip-address* [**srlg-group** *group-name*]

Context

[Tree] (config>router>mpls>srlg-database>router-id interface)

Full Context

configure router mpls srlg-database router-id interface

Description

This command allows the operator to manually enter the SRLG membership information for any link in the network, including links on this node, into the user SRLG database.

An interface can be associated with up to five SRLG groups for each execution of this command. The operator can associate an interface with up to 64 SRLG groups by executing the command multiple times.

CSPF will not use entered SRLG membership if an interface is not validated as part of a router ID in the routing table.

The **no** form of this command deletes a specific interface entry in this user SRLG database. The *group-name* must already exist in the config>router>if-attribute>srlg-group context.

Parameters

ip-address

Specifies the IPv4 address in a.b.c.d

srlg-group group-name

Specifies the SRLG group name. Up to 1024 group names can be defined in the config>router>if-attribute context. The SRLG group names must be identical across all routers in a single domain.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>router>mpls interface)

Full Context

configure router mpls interface

Description

This command specifies MPLS protocol support on an IP interface. No MPLS commands are executed on an IP interface where MPLS is not enabled. An MPLS interface must be explicitly enabled (**no shutdown**).

The **no** form of this command deletes all MPLS commands such as **label-map** which are defined under the interface. The MPLS interface must be shutdown first in order to delete the interface definition. If the

interface is not shutdown, the **no interface** *ip-int-name* command does nothing except issue a warning message on the console indicating that the interface is administratively up.

Default

shutdown

Parameters

ip-int-name

Specifies the name of the network IP interface. An interface name cannot be in the form of an IP address. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Values 1 to 32 alphanumeric characters.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>router>rsvp interface)

Full Context

configure router rsvp interface

Description

This command enables RSVP protocol support on an IP interface. No RSVP commands are executed on an IP interface where RSVP is not enabled.

The **no** form of this command deletes all RSVP commands such as **hello-interval** and **subscription**, which are defined for the interface. The RSVP interface must be **shutdown** it can be deleted. If the interface is not shut down, the **no interface** *ip-int-name* command does nothing except issue a warning message on the console indicating that the interface is administratively up.

Default

shutdown

Parameters

ip-int-name

Specifies the name of the network IP interface. An interface name cannot be in the form of an IP address. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Values 1 to 32

Platforms

7705 SAR Gen 2

interface

Syntax

interface [*ip-int-name* | *mt-int-name* | *ip-address*] [**detail**]

no interface

Context

[Tree] (debug>router>pim interface)

Full Context

debug router pim interface

Description

This command enables debugging for PIM interface information.

The **no** form of this command disables PIM interface debugging.

Parameters

ip-int-name

Debugs the information associated with the specified IP interface name.

Values IPv4 or IPv6 interface address

mt-int-name

Debugs the information associated with the specified VPRN ID and group address.

ip-address

Debugs the information associated with the specified IP address.

detail

Debugs detailed IP interface information.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>router>pim interface)

Full Context

configure router pim interface

Description

This command creates a PIM interface.

Interface names are case-sensitive and must be unique within the group of defined IP interfaces defined for **config>router>interface**, **config>service>ies>interface**, and **config>service>ies>subscriber-interface>group-interface**. Interface names must not be in the dotted decimal notation of an IP address. For example, the name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed. Show commands for router interfaces use either the interface names or the IP addresses. Ambiguity can exist if an IP address is used as an IP address and an interface name. Duplicate interface names can exist in different router instances, although this is not recommended because it may be confusing.

By default, no interfaces or names are defined within PIM.

The **no** form of this command removes the IP interface and all the associated configurations.

Parameters

ip-int-name

Specifies the name of the IP interface, up to 32 characters. Interface names must be unique within the group of defined IP interfaces for **config router interface**, **config service ies interface**, and **config service ies subscriber-interface group-interface** commands. An interface name cannot be in the form of an IP address. If the string contains special characters (#, $, spaces, and so on.), the entire string must be enclosed within double quotes.

If the *ip-int-name* already exists, the context is changed to maintain that IP interface. If *ip-int-name* does not exist, the interface is created and the context is changed to that interface for further command processing.

Platforms

7705 SAR Gen 2

interface

Syntax

interface *interface-name* [**unnumbered-mpls-tp**]

interface *interface-name* **pdn**

no interface *interface-name*

Context

[Tree] (config>router interface)

Full Context

configure router interface

Description

This command creates a logical IP routing or unnumbered MPLS-TP interface. Once created, attributes like IP address, port, or system can be associated with the IP interface.

Interface names are case-sensitive and must be unique within the group of IP interfaces defined for **config router interface** and **config service ies interface**. Interface names must not be in the dotted decimal notation of an IP address.; for example, the name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed. Show commands for router interfaces use either the interface names or the IP addresses. Ambiguity can exist if an IP address is used as an IP address and an interface name. Duplicate interface names can exist in different router instances, although this is not recommended because it is confusing. Nokia recommends that names are meaningful and unique to remove ambiguity when displaying the state associated with IP interfaces through show commands.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

Although not a keyword, the ip-int-name "**system**" is associated with the network entity (such as a specific router), not a specific interface. The system interface is also referred to as the loopback address.

An unnumbered MPLS-TP interface is a special type of interface that is only intended for MPLS-TP LSPs. IP routing protocols are blocked on interfaces of this type. If an interface is configured as unnumbered-mpls-tp, then it can only be associated with an Ethernet port or VLAN, using the port command, then either a unicast, multicast, or broadcast remote MAC address may be configured. Only static ARP is supported.

The control-tunnel parameter creates a loopback interface representing a GRE tunnel. One IP tunnel can be created in this interface.

Only the primary IPv4 interface address and only one IP tunnel per interface are allowed. Multiple tunnels can be configured using up to four controlTunnel loopback interfaces. A static route can take the new controlTunnel interface as a next hop.

The **no** form of this command removes the IP interface and all the associated configurations. The interface must be administratively shut down before issuing the **no interface** command.

Parameters

interface-name

Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Values 1 to 32 alphanumeric characters

If the *ip-int-name* already exists, the context is changed to maintain that IP interface. If *ip-int-name* already exists within another service ID or is an IP interface defined within the **config router** commands, an error will occur and the context will not be changed to that IP interface. If *ip-int-name* does not exist, the interface is created and the context is changed to that interface for further command processing.

unnumbered-mpls-tp

Specifies that an interface is an unnumbered MPLS-TP. An unnumbered MPLS-TP interface is a special type of interface that is only intended for MPLS-TP LSPs. IP routing protocols are blocked on interfaces of this type. If an interface is configured as **unnumbered-mpls-tp**, then it can only be associated with an Ethernet port or VLAN, using the **port** command. A unicast, multicast, or broadcast remote MAC address can be configured using the **static-arp** command. Only static ARP is supported.

pdn

Specifies that the interface is a PDN.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>router>router-advert interface)

Full Context

configure router router-advertisement interface

Description

This command configures router advertisement properties on a specific interface. The interface must already exist in the **config**>**router**>**if** context.

Parameters

ip-int-name

Specifies the interface name. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** [{*ip-int-name* | *ip-address*}]

Context

[Tree] (debug>router>ip interface)

Full Context

debug router ip interface

Description

This command displays the router IP interface table sorted by interface index.

Parameters

ip-int-name

Only displays the interface information associated with the specified IP interface name.

Values 32 characters maximum

ip-address

Only displays the interface information associated with the specified IP address.

Values

ipv4-address	a.b.c.d (host bits must be 0)
ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>router>isis interface)

Full Context

configure router isis interface

Description

This command creates the context to configure an IS-IS interface.

When an area is defined, the interfaces belong to that area. Interfaces cannot belong to separate areas.

When the interface is a POS channel, the OSINLCP is enabled when the interface is created and removed when the interface is deleted.

The **no** form of this command removes IS-IS from the interface.

The **shutdown** command in the **config>router>isis>interface** context administratively disables IS-IS on the interface without affecting the IS-IS configuration.

Parameters

ip-int-name

Identify the IP interface name created in the **config>router>interface** context. The IP interface name must already exist.

Platforms

7705 SAR Gen 2

interface

Syntax

interface [*ip-int-name* | *ip-address*]

no interface

Context

[Tree] (debug>router>isis interface)

Full Context

debug router isis interface

Description

This command enables debugging for IS-IS interface.

The **no** form of the command disables debugging.

Parameters

ip-address

When specified, only the interface with the specified interface address is debugged.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

ip-int-name

When specified, only the interface with the specified interface name is debugged.

Platforms

7705 SAR Gen 2

interface

Syntax

interface *ip-int-name* [**secondary**]

no interface *ip-int-name*

Context

[Tree] (config>router>ospf>area interface)

[Tree] (config>router>ospf3>area interface)

Full Context

configure router ospf area interface

configure router ospf3 area interface

Description

This command configures an OSPF interface.

Unless they are explicitly configured, interfaces are not activated, by default, in any interior gateway protocol, such as OSPF.

The **no** form of this command deletes the OSPF interface configuration for this interface. Use the **shutdown** command in the **config**>**router**>**ospf**>**interface** context to disable an interface without removing the configuration for the interface.

Default

no interface

Parameters

ip-int-name

Specifies the IP interface name. Interface names must be unique within the group of defined IP interfaces for the **configure router interface** and **configure service ies interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string, up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

If the IP interface name does not exist or does not have an IP address configured, an error message is returned.

If the IP interface exists in a different area it is moved to this area.

secondary

Keyword used to allow multiple secondary adjacencies, in addition to the primary adjacency, to be established over a single IP interface. This keyword can also be applied

to the system interface and to loopback interfaces to allow them to participate in multiple areas, although no adjacencies are formed over these types of interfaces.

Platforms

7705 SAR Gen 2

interface

Syntax

interface [*ip-int-name* | *ip-address*]

interface [**interface-name**]

no interface

Context

[Tree] (debug>router>ospf3 interface)

[Tree] (debug>router>ospf interface)

Full Context

debug router ospf3 interface

debug router ospf interface

Description

This command enables debugging for an OSPF and OSPF3 interface.

Parameters

ip-int-name

Specifies the IP interface name, in the **debug**>**router**>**ospf** context. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

ip-address

Specifies the interface's IP address, in the **debug**>**router**>**ospf** context.

interface-name

Specifies the interface name, in the **debug**>**router**>**ospf3** context.

Platforms

7705 SAR Gen 2

interface

Syntax

interface *interface-name*

no interface

Context

[Tree] (config>router>policy-options>policy-statement>entry>from interface)

Full Context

configure router policy-options policy-statement entry from interface

Description

This command specifies the router interface, specified either by name or address, as a filter criterion.

The **no** form of this command removes the criterion from the configuration.

Default

no interface

Parameters

ip-int-name

Specifies the name of the interface as a match criterion for this entry. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>router>bgp>group>dynamic-neighbor interface)

Full Context

configure router bgp group dynamic-neighbor interface

Description

Commands in this context configure an unnumbered base router network interface for dynamic neighbors.

If this interface connects to a network with other BGP routers, sessions with the other routers can be set up automatically without explicitly configuring them as BGP neighbors. The interface must be IPv6 enabled, but because the interface is considered unnumbered, it does not require an IPv4 address or a global-unicast IPv6 address. The sessions are set up using IPv6 link-local addresses.

The BGP unnumbered feature supports all address families that allow IPv6 link-local BGP next-hop addresses. This includes IPv4 with the use of RFC 8950 extensions.

When an interface is added to the list of dynamic-neighbor interfaces, an outgoing connection attempt is initiated toward any directly connected router on the interface that announces itself using an ICMPv6 router advertisement message. The session attempt is unsuccessful if the peer type is not EBGP, the reported AS number of the peer does not match one of the allowed values, or the maximum session limit of the interface would be exceeded.

The **no** form of this command removes the interface from the list of dynamic-neighbor interfaces.

Parameters

ip-int-name

Specifies the name of a base router IP interface, up to 32 characters.

Platforms

7705 SAR Gen 2

interface

Syntax

[**no**] **interface** *ip-int-name*

Context

[Tree] (config>service>vprn>bgp>group>dynamic-neighbor interface)

Full Context

configure service vprn bgp group dynamic-neighbor interface

Description

Commands in this context configure an unnumbered VPRN access IP interface for dynamic neighbors.

If this interface connects to a network with other BGP routers, sessions with the other routers can be set up automatically without explicitly configuring them as BGP neighbors. The interface must be IPv6 enabled, but because the interface is considered unnumbered, it does not require an IPv4 address or a global-unicast IPv6 address. The sessions are set up using IPv6 link-local addresses.

The BGP unnumbered feature supports all address families that allow IPv6 link-local BGP next-hop addresses. This includes IPv4 with the use of RFC 8950 extensions.

When an interface is added to the list of dynamic-neighbor interfaces, an outgoing connection attempt is initiated toward any directly connected router on the interface that announces itself using an ICMPv6 router advertisement message. The session attempt is unsuccessful if the peer type is not EBGP, the reported AS number of the peer does not match one of the allowed values, or the maximum session limit of the interface would be exceeded.

The **no** form of this command removes the interface from the list of dynamic-neighbor interfaces.

Parameters

ip-int-name

Specifies the name of a VPRN access IP interface, up to 32 characters.

Platforms

7705 SAR Gen 2

13.101 interface-ful

interface-ful

Syntax

interface-ful

Context

[Tree] (config>service>system>bgp-evpn>ip-prefix-routes interface-ful)

Full Context

configure service system bgp-evpn ip-prefix-routes interface-ful

Description

Commands in this context configure IP prefix routes for Interface-ful (IFF) configurations.

Platforms

7705 SAR Gen 2

13.102 interface-id

interface-id

Syntax

interface-id [**ascii-tuple**]

interface-id ifindex

interface-id sap-id

interface-id string *string*

no interface-id

Context

[Tree] (config>service>vprn>if>ipv6>dhcp6>option interface-id)

[Tree] (config>service>ies>if>ipv6>dhcp6>option interface-id)

Full Context

configure service vprn interface ipv6 dhcp6-relay option interface-id

configure service ies interface ipv6 dhcp6-relay option interface-id

Description

This command enables the sending of interface ID options in the DHCPv6 relay packet.

The **no** form of this command disables the sending of interface ID options in the DHCPv6 relay packet.

Parameters

ascii-tuple

Specifies that the ASCII-encoded concatenated tuple is used which consists of the access-node-identifier, service-id, and interface-name, separated by "| ".

ifindex

Specifies that the interface index is used. The If Index of a router interface can be displayed using the **show>router>interface>detail** command.

sap-id

Specifies that the SAP identifier is used.

string

Specifies that a string is used.

string

Specifies a string of up to 80 characters long, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

interface-id

Syntax

interface-id [**ascii-tuple**]

interface-id [**vlan-ascii-tuple**]

no interface-id

Context

[Tree] (config>service>vpls>sap>dhcp6>ldra>options interface-id)

Full Context

configure service vpls sap dhcp6 ldra options interface-id

Description

This command enables the sending of interface ID options in the DHCPv6 LDRA.

The **no** form of this command disables the sending of interface ID options in the DHCPv6 LDRA.

Parameters

ascii-tuple

Specifies the use of the ASCII-encoded concatenated tuple, which consists of the *system name*, *service-id*, and *sap-id* separated by "| ".

vlan-ascii-tuple

Specifies the use of the ASCII-encoded concatenated tuple enhanced with VLAN ID and dot1p bits, consisting of the *system name*, *service-id*, *sap-id*, *dot1p-inner-vlan*, and *inner-vplan-id*, separated by "| ".

Platforms

7705 SAR Gen 2

13.103 interface-id-mapping

interface-id-mapping

Syntax

[**no**] **interface-id-mapping**

Context

[Tree] (config>router>dhcp6>server interface-id-mapping)

[Tree] (config>service>vprn>dhcp6>server interface-id-mapping)

Full Context

configure router dhcp6 local-dhcp-server interface-id-mapping

configure service vprn dhcp6 local-dhcp-server interface-id-mapping

Description

This command enables the behavior where unique /64 prefix is allocated per interface-id, and all clients having the same interface-id get an address allocated out of this /64 prefix for DHCP6. This is relevant for bridged clients behind the same local-loop (and same SAP), where sharing the same prefix allows communication between bridged clients behind the same local-loop to stay local. For SLAAC based assignment, downstream neighbor-discovery is automatically enabled to resolve the assigned address.

The **no** form of this command reverts to the default.

Platforms

7705 SAR Gen 2

13.104 interface-parameters

interface-parameters

Syntax

interface-parameters

Context

[Tree] (config>router>ldp interface-parameters)

Full Context

configure router ldp interface-parameters

Description

Commands in this context configure LDP interfaces and parameters applied to LDP interfaces. The user can configure different default parameters for IPv4 and IPv6 LDP interfaces by entering **ipv4** or **ipv6** as the next command.

Platforms

7705 SAR Gen 2

13.105 interface-subnets

interface-subnets

Syntax

interface-subnets [**service** *service-id*] *interface-name*

no interface-subnets

Context

[Tree] (config>router>policy-options>policy-statement>entry>from interface-subnets)

Full Context

configure router policy-options policy-statement entry from interface-subnets

Description

This command configures the applied router instance and interfaces that are used as matching condition within each policy-statement entry. A maximum of 10 *interface-name* entries is supported, and all entries must belong to the same routing context (either **base** or **service**). The interface subnet policy-statement match criterion is applied to the following unicast use case contexts:

- **export**, when used with OSPFv2, OSPFv3, IS-IS, RIP, RIPng, and BGP
- **route-table-import**, when used with BGP
- **vrf-export**, when used with MP-BGP

The **no** form of this command removes all policies from the configuration.

Default

no interface-subnets

Parameters

service

Specifies the context in which the configured interface exists. By default, the base routing instance is assumed. However, the configured service context is used only when the service is configured.

service-id

Specifies the service ID of the service to match.

Values *service-id* — 1 to 2147483647

svc-name — 64 characters maximum

interface-name

Specifies the interface name, up to 32 characters, to match when exporting the IP address of the associated interface to a routing protocol.

Platforms

7705 SAR Gen 2

13.106 interface-type

interface-type

Syntax

interface-type {**broadcast** | **point-to-point**}

no interface-type

Context

[Tree] (config>service>vprn>isis>if interface-type)

Full Context

configure service vprn isis interface interface-type

Description

This command configures the IS-IS interface type as either broadcast or point-to-point.

Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the designated IS-IS overhead if the link is used as a point-to-point.

If the interface type is not known at the time the interface is added to IS-IS and subsequently the IP interface is bound (or moved) to a different interface type, then this command must be entered manually.

The **no** form of this command reverts to the default value.

Default

point-to-point — For IP interfaces on SONET channels.

broadcast — For IP interfaces on Ethernet or unknown type physical interfaces.

Parameters

broadcast

Configures the interface to maintain this link as a broadcast network.

point-to-point

Configures the interface to maintain this link as a point-to-point link.

Platforms

7705 SAR Gen 2

interface-type

Syntax

interface-type {**broadcast** | **point-to-point** | **non-broadcast** | **p2mp-nbma**}

no interface-type

Context

[Tree] (config>service>vprn>ospf3>area>if interface-type)

[Tree] (config>service>vprn>ospf>area>if interface-type)

Full Context

configure service vprn ospf3 area interface interface-type

configure service vprn ospf area interface interface-type

Description

This command configures the interface type to:

- broadcast

- non-broadcast
- point-to-point
- point-to-multipoint on a link without broadcast or multicast support

Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the broadcast adjacency maintenance overhead if the Ethernet link provided the link is used as a point-to-point.

For subscriber interfaces, configure the adjacent interface (CPE) with interface type point-to-point. For subscriber interfaces, when the interface is configured as P2MP-NBMA, the subscriber interface becomes an active OPSF interface, allowing it to both send and receive OSPF LSAs. For all other interface types, subscriber interfaces remain as passive OSPF interfaces by default.

The **no** form of this command reverts to the default value.

Default

point-to-point — If the physical interface is SONET.

broadcast — If the physical interface is Ethernet or unknown.

Parameters

broadcast

Specifies the interface as a broadcast network. To significantly improve adjacency forming and network convergence, configure the network as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.

point-to-point

Specifies the interface as a point-to-point link. Set the interface type of an Ethernet link to point-to-point to avoid having to carry the broadcast adjacency maintenance overhead if the Ethernet link provided is used as a point-to-point.

non-broadcast

Specifies the interface as a non-broadcast network.

p2mp-nbma

Specifies the interface as a point-to-multipoint on a link without broadcast or multicast support. No designated router or backup designated router is elected on this type of interface and all OSPF neighbors connect through individual point-to-point links. Only VPRN and IES services interfaces support this interface type.

Platforms

7705 SAR Gen 2

interface-type

Syntax

interface-type {**broadcast** | **point-to-point**}

no interface-type

Context

[Tree] (config>router>isis>interface interface-type)

Full Context

configure router isis interface interface-type

Description

This command configures the IS-IS interface type as either broadcast or point-to-point.

Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the designated IS-IS overhead if the link is used as a point-to-point.

If the interface type is not known at the time the interface is added to IS-IS and subsequently the IP interface is bound (or moved) to a different interface type, then this command must be entered manually.

The **no** form of this command reverts to the default value.

Default

interface-type point-to-point — For IP interfaces on SONET channels.

interface-type broadcast — For IP interfaces on Ethernet or unknown type physical interfaces.

Parameters

broadcast

Configures the interface to maintain this link as a broadcast network.

point-to-point

Configures the interface to maintain this link as a point-to-point link.

Platforms

7705 SAR Gen 2

interface-type

Syntax

interface-type {**broadcast** | **point-to-point** | **non-broadcast** | **p2mp-nbma**}

no interface-type

Context

[Tree] (config>router>ospf>area>interface interface-type)

[Tree] (config>router>ospf3>area>interface interface-type)

Full Context

configure router ospf area interface interface-type

configure router ospf3 area interface interface-type

Description

This command configures the interface type to:

- broadcast
- non-broadcast
- point-to-point
- point-to-multipoint on a link without broadcast or multicast support

Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the broadcast adjacency maintenance overhead of the Ethernet link provided the link is used as point-to-point.

For subscriber interfaces, configure the adjacent interface (CPE) with interface type point-to-point. For subscriber interfaces, when the interface is configured as P2MP-NBMA, the subscriber interface becomes an active OPSF interface, allowing it to both send and receive OSPF LSAs. For all other interface types, subscriber interfaces remain as passive OSPF interfaces by default.

The **no** form of this command returns the setting to the default value.

Default

interface-type point-to-point (if the physical interface is SONET)

interface-type broadcast (if the physical interface is Ethernet or unknown)

Parameters

broadcast

Specifies the interface as a broadcast network. To significantly improve adjacency forming and network convergence, configure a network as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.

point-to-point

Specifies the interface as a point-to-point link. Set the interface type of an Ethernet link to point-to-point to avoid having to carry the broadcast adjacency maintenance overhead if the Ethernet link provided is used as a point-to-point.

non-broadcast

Specifies the interface as a non-broadcast network.

p2mp-nbma

Specifies the interface as a point-to-multipoint on a link without broadcast or multicast support. No designated router or backup designated router is elected on this type of interface and all OSPF neighbors connect through individual point-to-point links. Only VPRN and IES services interfaces support this interface type.

Platforms

7705 SAR Gen 2

interface-type

Syntax

interface-type {**client-facing** | **network-facing**}

no interface-type

Context

[Tree] (config>service>vpls>sap>dhcp6>ldra interface-type)

Full Context

configure service vpls sap dhcp6 ldra interface-type

Description

This command configures LDRA interface type as either client or network facing.

The **no** form of this command reverts to the default value.

Default

no interface-type

Parameters

client-facing

Configures the SAP as an untrusted client-facing interface. Only DHCPv6 client messages are accepted and encapsulated in a Relay-Forward message. It is mandatory to configure an interface ID for client-facing SAPs. Relay-Forward, Relay-Reply, and DHCPv6 server messages are silently dropped when received on a client-facing SAP

network-facing

Configures the SAP as a network-facing interface. Only Relay-Reply messages are accepted: the server message is extracted from the Relay-Reply message and forwarded in the VPLS. All other DHCPv6 message types are silently dropped when received on a network-facing SAP.

Platforms

7705 SAR Gen 2

13.107 internal-ip4-address

internal-ip4-address

Syntax

[**no**] **internal-ip4-address**

Context

[Tree] (config>ipsec>ike-policy>relay-unsol-attr internal-ip4-address)

Full Context

configure ipsec ike-policy relay-unsolicited-cfg-attribute internal-ip4-address

Description

This command will return IPv4 address from source (such as a RADIUS server) to IKEv2 remote-access tunnel client regardless if the client has requested it in the CFG_REQUEST payload.

Default

no internal-ip4-address

Platforms

7705 SAR Gen 2

13.108 internal-ip4-dns

internal-ip4-dns

Syntax

[**no**] **internal-ip4-dns**

Context

[Tree] (config>ipsec>ike-policy>relay-unsol-attr internal-ip4-dns)

Full Context

configure ipsec ike-policy relay-unsolicited-cfg-attribute internal-ip4-dns

Description

This command will return IPv4 DNS server address from source (such as a RADIUS server) to IKEv2 remote-access tunnel client regardless if the client has requested it in the CFG_REQUEST payload.

Default

no internal-ip4-dns

Platforms

7705 SAR Gen 2

13.109 internal-ip4-netmask

internal-ip4-netmask

Syntax

[**no**] **internal-ip4-netmask**

Context

[Tree] (config>ipsec>ike-policy>relay-unsol-attr internal-ip4-netmask)

Full Context

configure ipsec ike-policy relay-unsolicited-cfg-attribute internal-ip4-netmask

Description

This command will return IPv4 netmask from source (such as a RADIUS server) to IKEv2 remote-access tunnel client regardless if the client has requested it in the CFG_REQUEST payload.

Default

no internal-ip4-netmask

Platforms

7705 SAR Gen 2

13.110 internal-ip6-address

internal-ip6-address

Syntax

[**no**] **internal-ip6-address**

Context

[Tree] (config>ipsec>ike-policy>relay-unsol-attr internal-ip6-address)

Full Context

configure ipsec ike-policy relay-unsolicited-cfg-attribute internal-ip6-address

Description

This command will return IPv6 address from source (such as a RADIUS server) to IKEv2 remote-access tunnel client regardless if the client has requested it in the CFG_REQUEST payload.

Default

no internal-ip6-address

Platforms

7705 SAR Gen 2

13.111 internal-ip6-dns

internal-ip6-dns

Syntax

[**no**] **internal-ip6-dns**

Context

[Tree] (config>ipsec>ike-policy>relay-unsol-attr internal-ip6-dns)

Full Context

configure ipsec ike-policy relay-unsolicited-cfg-attribute internal-ip6-dns

Description

This command will return IPv6 DNS server address from source (RADIUS server) to IKEv2 remote-access tunnel client regardless if the client has requested it in the CFG_REQUEST payload.

Default

no internal-ip6-dns

Platforms

7705 SAR Gen 2

13.112 internal-lease-ipsec

internal-lease-ipsec

Syntax

[**no**] **internal-lease-ipsec**

Context

[Tree] (config>router>dhcp6>server>lease-hold-time-for internal-lease-ipsec)

[Tree] (config>service>vprn>dhcp>server>lease-hold-time-for internal-lease-ipsec)

[Tree] (config>service>vprn>dhcp6>server>lease-hold-time-for internal-lease-ipsec)

[Tree] (config>router>dhcp>server>lease-hold-time-for internal-lease-ipsec)

Full Context

configure router dhcp6 local-dhcp-server lease-hold-time-for internal-lease-ipsec

configure service vprn dhcp local-dhcp-server lease-hold-time-for internal-lease-ipsec

configure service vprn dhcp6 local-dhcp-server lease-hold-time-for internal-lease-ipsec

configure router dhcp local-dhcp-server lease-hold-time-for internal-lease-ipsec

Description

This command enables the server to hold up the lease of local IPsec clients.

The **no** form of this command disables the ability of the server to hold up the lease of local IPsec clients.

Platforms

7705 SAR Gen 2

13.113 interval

interval

Syntax

interval *seconds*

no interval

Context

[Tree] (config>service>vprn>static-route-entry>next-hop>cpe-check interval)

[Tree] (config>service>vprn>static-route-entry>indirect>cpe-check interval)

Full Context

configure service vprn static-route-entry next-hop cpe-check interval

configure service vprn static-route-entry indirect cpe-check interval

Description

This optional parameter specifies the interval between ICMP pings to the target IP address.

Default

interval 1

Parameters

seconds

An integer interval value.

Values 1 to 255

Platforms

7705 SAR Gen 2

interval

Syntax

interval *seconds*

Context

[Tree] (config>ipsec>tnl-temp>icmp6-gen>pkt-too-big interval)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>icmp6-gen>pkt-too-big interval)

Full Context

configure ipsec tunnel-template icmp6-generation pkt-too-big interval

configure service ies interface ipsec ipsec-tunnel icmp6-generation pkt-too-big interval

Description

This command configures the maximum interval during which messages can be sent.

Parameters

seconds

Specifies the maximum interval during which messages can be sent, in seconds.

Values 1 to 60

Default 10

Platforms

7705 SAR Gen 2

interval

Syntax

interval *seconds*

Context

[Tree] (config>router>if>ipsec-tunnel>icmp6-gen>pkt-too-big interval)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>icmp6-gen>pkt-too-big interval)

[Tree] (config>ipsec>tnl-temp>icmp6-gen>pkt-too-big interval)

[Tree] (config>service>vprn>if>sap>ipsec-tun>icmp6-gen>pkt-too-big interval)

Full Context

configure router interface ipsec ipsec-tunnel icmp6-generation pkt-too-big interval

configure service vprn interface ipsec ipsec-tunnel icmp6-generation pkt-too-big interval

configure ipsec tunnel-template icmp6-generation pkt-too-big interval

configure service vprn interface sap ipsec-tunnel icmp6-generation pkt-too-big interval

Description

This command configures the interval for sending ICMPv6 Packet Too Big (code 2) messages. The maximum number of messages that can be sent during the interval is configured by the **message-count** command.

The **no** form of the command reverts to the default value.

Default

interval 10

Parameters

seconds

Specifies the time, in seconds, for sending 'message-count' ICMPv6 messages.

Values 1 to 60

Platforms

7705 SAR Gen 2

interval

Syntax

interval *interval*

no interval

Context

[Tree] (config>saa>test>type-multi-line>lsp-ping>sr-policy interval)

[Tree] (config>saa>test>type-multi-line>lsp-ping interval)

Full Context

configure saa test type-multi-line lsp-ping sr-policy interval

configure saa test type-multi-line lsp-ping interval

Description

This command configures the number of seconds to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

The **no** form of this command reverts to the default value.

Default

interval 1

Parameters

interval

Specifies the number of seconds to wait before the next message request is sent.

Values 1 to 10

Default 1

Platforms

7705 SAR Gen 2

interval

Syntax

interval *milliseconds*

no interval

Context

[Tree] (config>oam-pm>session>ip>twamp-light interval)

Full Context

configure oam-pm session ip twamp-light interval

Description

This command defines the message period, or probe spacing, for transmitting a TWAMP Light frame.

The **no** form of this command sets the interval to the default value.

Default

interval 1000

Parameters

milliseconds

Specifies the number of milliseconds between TWAMP Light frame transmission.

Values 50, 100, 200, 300, 400, 500, 600, 700, 800, 900, 1000, 10000

Default 1000

Platforms

7705 SAR Gen 2

interval

Syntax

interval [*seconds*]

no interval

Context

[Tree] (config>filter>redirect-policy>dest>ping-test interval)

Full Context

configure filter redirect-policy destination ping-test interval

Description

This command specifies the amount of time, in seconds, between consecutive requests sent to the far end host.

Default

interval 1

Parameters

seconds

Specifies the amount of time, in seconds, between consecutive requests sent to the far end host.

Values 1 to 60

Platforms

7705 SAR Gen 2

interval

Syntax

interval *seconds*

no interval

Context

[Tree] (config>router>static-route-entry>indirect>cpe-check interval)

[Tree] (config>router>static-route-entry>next-hop>cpe-check interval)

Full Context

configure router static-route-entry indirect cpe-check interval

configure router static-route-entry next-hop cpe-check interval

Description

This optional parameter specifies the interval between ICMP pings to the target IP address.

Default

interval 1

Parameters

seconds

Specifies the interval value, in seconds.

Values 1 to 255

Platforms

7705 SAR Gen 2

interval

Syntax

interval *seconds*

no interval

Context

[Tree] (config>vrrp>policy>priority-event>host-unreachable interval)

Full Context

configure vrrp policy priority-event host-unreachable interval

Description

This command configures the number of seconds between host unreachable priority event ICMP echo request messages directed to the host IP address.

The **no** form of the command reverts to the default value.

Default

interval 1

Parameters

seconds

Specifies the number of seconds between the ICMP echo request messages sent to the host IP address for the host unreachable priority event.

Values 1 to 60

Platforms

7705 SAR Gen 2

interval

Syntax

interval *seconds*

no interval

Context

[Tree] (config>system>cron>sched interval)

Full Context

configure system cron schedule interval

Description

This command specifies the interval between runs of an event.

Default

no interval

Parameters

seconds

Specifies the interval, in seconds, between runs of an event.

Values 30 to 42949672

Platforms

7705 SAR Gen 2

interval

Syntax

interval *interval*

no interval

Context

[Tree] (config>system>grpc>tcp-keepalive interval)

[Tree] (config>system>telemetry>destination-group>tcp-keepalive interval)

[Tree] (config>system>grpc-tunnel>destination-group>tcp-keepalive interval)

Full Context

configure system grpc tcp-keepalive interval

configure system telemetry destination-group tcp-keepalive interval

configure system grpc-tunnel destination-group tcp-keepalive interval

Description

This command configures the amount of time, in seconds, between successive TCP keepalive probes sent by the router.

The **no** form of this command reverts to the default value.

Default

interval 15

Parameters

interval

Specifies the number of seconds between TCP keepalive probes.

Values 1 to 100000

Default 15

Platforms

7705 SAR Gen 2

interval

Syntax

interval *seconds*

Context

[Tree] (config>router>if>ipsec>ipsec-tunnel>icmp-generation>frag-required interval)

[Tree] (config>ipsec>tnl-temp>icmp-gen>frag-required interval)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>icmp-generation>frag-required interval)

[Tree] (config>service>vprn>if>sap>ip-tunnel>icmp-generation>frag-required interval)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel>icmp-generation>frag-required interval)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>icmp-generation>frag-required interval)

Full Context

configure router interface ipsec ipsec-tunnel icmp-generation frag-required interval

configure ipsec tunnel-template icmp-generation frag-required interval

configure service ies interface ipsec ipsec-tunnel icmp-generation frag-required interval

configure service vprn interface sap ip-tunnel icmp-generation frag-required interval

configure service vprn interface sap ipsec-tunnel icmp-generation frag-required interval

configure service vprn interface ipsec ipsec-tunnel icmp-generation frag-required interval

Description

This command configures the interval for sending ICMP Destination Unreachable "fragmentation needed and DF set" messages (type 3, code 4). The maximum number of messages that can be sent during the interval is configured by the **message-count** command.

The **no** form of the command reverts to the default value.

Default

interval 10

Parameters

seconds

Specifies the time, in seconds, for sending ICMPv6 Destination Unreachable "fragmentation needed and DF set" messages (type 3, code 4).

Values 1 to 60

Platforms

7705 SAR Gen 2

13.114 intervals-stored

intervals-stored

Syntax

intervals-stored *intervals*

no intervals-stored

Context

[Tree] (config>oam-pm>session>meas-interval intervals-stored)

Full Context

configure oam-pm session meas-interval intervals-stored

Description

This command defines the number of completed measurement intervals per session to be stored in volatile system memory. The entire block of memory is allocated for the measurement interval when the test is active (**no shutdown**) to ensure memory is available. The numbers are increasing from 1 to the configured value + 1. The active pm data is stored in the interval number 1 and older runs are stored, in

order, to the upper most number with the oldest rolling off when the number of completed measurement intervals exceeds the configured value+1. As new test measurement intervals complete for the session, the stored intervals are renumbered to maintain the described order. Use caution when setting this value. There must be a balance between completed runs stored in volatile memory and the use of the write-to-flash function of the accounting policy.

The **5-mins** and **15-mins** measurement intervals share the same (1 to 96) retention pool. In the event that both intervals are required, the sum total of both intervals cannot exceed 96. The **1-hour** and **1-day** measurement intervals utilize their own ranges.

If this command is omitted when configuring the measurement interval, the default value is used.

The **no** form of the command reverts to the default.

Default

intervals-stored 1

Parameters

intervals

Specifies the number of measurement intervals.

Values **5-mins**: 1 to 96
15-mins: 1 to 96
1-hour: 1 to 24
1-day: 1

Default **5-mins**: 32
15-mins: 32
1-hour: 8
1-day: 1

Platforms

7705 SAR Gen 2

13.115 iom

iom

Syntax

iom [**detail**]

no iom

Context

[Tree] (debug>router>mpls>event iom)

Full Context

debug router mpls event iom

Description

This command reports MPLS debug events originating from the XMA.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about MPLS events originating from the XMA.

Platforms

7705 SAR Gen 2

13.116 ip

ip

Syntax

ip *address*

no ip

Context

[Tree] (config>service>vpls>mcr-default-gtw ip)

Full Context

configure service vpls mcr-default-gtw ip

Description

This command relates to a system configured for Dual Homing in L2-TPSDA. It defines the IP address used when the system sends out a gratuitous ARP on an active SAP after a ring heals or fails in order to attract traffic from subscribers on the ring with connectivity to that SAP.

The **no** form of this command reverts to the default.

Default

no ip

Parameters

address

Specifies the IP address in a.b.c.d. format.

Platforms

7705 SAR Gen 2

ip

Syntax

ip *name*

no ip

Context

[Tree] (config>service>template>vpls-sap-template>ingress>filter-name ip)

[Tree] (config>service>template>vpls-sap-template>egress>filter-name ip)

Full Context

configure service template vpls-sap-template ingress filter-name ip

configure service template vpls-sap-template egress filter-name ip

Description

This command associates an existing IP filter policy with the template.

Parameters

name

Specifies the IP filter policy name, up to 64 characters.

Platforms

7705 SAR Gen 2

ip

Syntax

ip

Context

[Tree] (config>oam-pm>session ip)

Full Context

configure oam-pm session ip

Description

Commands in this context configure the IP-specific source and destination information, the priority, and the IP test tools on the launch point.

Platforms

7705 SAR Gen 2

ip

Syntax

[**no**] **ip** *ip-filter-id*

Context

[Tree] (config>filter>system-filter ip)

Full Context

configure filter system-filter ip

Description

This command activates an IPv4 system filter policy. Once activated, all IPv4 ACL filter policies that chain to the system filter (**config>filter>ip-filter>chain-to-system-filter**) will automatically execute system filter policy rules first.

The **no** form of the command deactivates the system filter policy.

Parameters

ip-filter-id

Specifies the existing IPv4 filter policy with scope **system**. This parameter can either be expressed as a decimal integer, or as an ASCII string of up to 64 characters.

Values 1 to 65535 or the filter policy name (*filter-name*, 64 char max)

Platforms

7705 SAR Gen 2

ip

Syntax

[**no**] **ip**

Context

[Tree] (debug>router ip)

Full Context

debug router ip

Description

This command configures debugging for IP.

Platforms

7705 SAR Gen 2

ip

Syntax

ip

Context

[Tree] (config>system ip)

Full Context

configure system ip

Description

This command configures system-wide IP router parameters.

Platforms

7705 SAR Gen 2

ip

Syntax

ip *ip-address netmask*

ip *ip-address/mask*

ip ip-prefix-list *ip-prefix-list-name*

no ip

Context

[Tree] (config>filter>ip-filter>entry>match ip)

Full Context

configure filter ip-filter entry match ip

Description

This command configures a destination or source IP address to be used as an IP match criterion.

Parameters

ip-address/mask

Specifies the IPv4 address and mask.

Values ip-address a.b.c.d

netmask

Specifies the name of the IP prefix list, up to 256 characters.

ip-prefix-list-name

Specifies the name of an IP prefix list, up to 32 characters.

Platforms

7705 SAR Gen 2

ip

Syntax

ip *ipv6-address ipv6-address-mask*

ip *ipv6-address/mask*

ip ipv6-prefix-list *prefix-list-name*

no ip

Context

[Tree] (config>filter>ipv6-filter>entry>match ip)

Full Context

configure filter ipv6-filter entry match ip

Description

This command configures a destination or source IP address to be used as an IP match criterion.

Parameters

ipv6-address/mask

Specifies the IPv6 address and mask.

Values ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

ip-prefix-list-name

Specifies the name of an IPv6 prefix list, up to 32 characters.

Platforms

7705 SAR Gen 2

13.117 ip-criteria

ip-criteria

Syntax

[**no**] **ip-criteria**

Context

[Tree] (config>qos>sap-ingress ip-criteria)

[Tree] (config>qos>sap-egress ip-criteria)

Full Context

configure qos sap-ingress ip-criteria

configure qos sap-egress ip-criteria

Description

IP criteria-based SAP ingress or egress policies are used to select the appropriate ingress or egress queue or policer and corresponding forwarding class and packet profile for matched traffic.

This command is used to enter the context to create or edit policy entries that specify IP criteria such as IP quintuple lookup or DiffServ code point.

The software implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason, entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. When IP criteria entries are removed from a SAP ingress or egress policy, the IP criteria is removed from all services where that policy is applied.

Platforms

7705 SAR Gen 2

ip-criteria

Syntax

[**no**] **ip-criteria**

Context

[Tree] (config>qos>network>ingress ip-criteria)

[Tree] (config>qos>network>egress ip-criteria)

Full Context

configure qos network ingress ip-criteria

configure qos network egress ip-criteria

Description

IP criteria-based network ingress and egress policies are used to select the appropriate ingress or egress queue or policer, and the corresponding forwarding class and packet profile for matched traffic. This command is used to enter the context to create or edit policy entries that specify IP criteria such as IP quintuple lookup or DSCP.

The SR OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. Entries must be sequenced correctly from most to least explicit.

The ingress classification only applies to the outer IP header of non-tunneled traffic. The only exception is for traffic received on a Draft Rosen tunnel, for which only classification on the outer IP header is supported.

Attempting to apply a network QoS policy containing an **ip-criteria** statement to any object except a network IP interface will result in an error.

The **no** form of this command deletes all entries specified under this node. When IP criteria entries are removed from a network policy, the IP criteria are removed from all network interfaces to which that policy is applied.

Platforms

7705 SAR Gen 2

13.118 ip-exception

ip-exception

Syntax

ip-exception *filter-id*

no ip-exception

Context

[Tree] (config>service>ies>if>ipsec ip-exception)

[Tree] (config>service>vprn>if>ipsec ip-exception)

[Tree] (config>router>if>ipsec ip-exception)

Full Context

configure service ies interface ipsec ip-exception

configure service vprn interface ipsec ip-exception

configure router interface ipsec ip-exception

Description

This command configures the IP exception filter for the secured interface. All ingress traffic matching by the specified filter bypasses IPsec processing.

The **no** form of this command removes the policy from the configuration.

Default

no ip-exception

Parameters

filter-id

Specifies IP filter policy that will be used to bypass encryption.

Platforms

7705 SAR Gen 2

ip-exception

Syntax

ip-exception *filter-id* [**create**]

no ip-exception *filter-id*

Context

[Tree] (config>filter ip-exception)

Full Context

configure filter ip-exception

Description

Commands in this context configure the specified IPv4 exception filter.

The **no** form of the command deletes the IPv4 exception filter.

Parameters

filter-id

Specifies the IPv4 filter policy ID expressed as a decimal integer.

Values 1 to 65535

create

This keyword is required to create the configuration context. Once it is created, the context can be enabled with or without the **create** keyword.

Platforms

7705 SAR Gen 2

ip-exception

Syntax

ip-exception *filter-id* **direction** {**inbound** | **outbound**}

no ip-exception direction {**inbound** | **outbound**}

Context

[Tree] (config>router>if>group-encryption ip-exception)

Full Context

configure router interface group-encryption ip-exception

Description

This command associates an IP exception filter policy with an NGE-enabled router interface to allow packets matching the exception criteria to transit the NGE domain as clear text.

When an exception filter is added for inbound traffic, packets matching the criteria in the IP exception filter policy are allowed to be received in clear text even if an inbound key group is configured. If no inbound key group is configured, then associated inbound IP exception filter policies will be ignored.

When an exception filter is added for outbound traffic, packets matching the criteria in the IP exception filter policy are not encrypted when sent out of the router interface even if an outbound key group is configured. If no outbound key group is configured, then associated outbound IP exception filter policies will be ignored.

The **no** form of this command removes the IP exception filter policy from the specified direction.

Default

no ip-exception direction inbound

no ip-exception direction outbound

Parameters

filter-id

Specifies the IP exception filter policy. The IP exception ID or exception name must have already been created.

Values 1 to 6553, *filter-name* (64 characters maximum)

inbound

Binds the exception filter policy in the inbound direction.

outbound

Binds the exception filter policy in the outbound direction.

Platforms

7705 SAR Gen 2

13.119 ip-filter

ip-filter

Syntax

ip-filter *ip-filter-id* **entry** *entry-id* [*entry-id*]

no ip-filter *ip-filter-id* [**entry** *entry-id*]

Context

[Tree] (config>mirror>mirror-source ip-filter)

Full Context

configure mirror mirror-source ip-filter

Description

This command enables mirroring of packets that match specific entries in an existing IP filter.

The **ip-filter** command directs packets which match the defined list of entry IDs to be mirrored to the mirror destination referenced by the *mirror-dest-service-id* of the **mirror-source**.

The IP filter must already exist in order for the command to execute. Filters are configured in the **config>filter** context. If the IP filter does not exist, an error will occur. If the filter exists but has not been associated with a SAP or IP interface, an error is not generated but mirroring will not be enabled (there are no packets to mirror). Once the IP filter is defined to a SAP or IP interface, mirroring is enabled.

If the IP filter is defined as ingress, only ingress packets are mirrored. Ingress mirrored packets are mirrored to the mirror destination prior to any ingress packet modifications.

If the IP filter is defined as egress, only egress packets are mirrored. Egress mirrored packets are mirrored to the mirror destination after all egress packet modifications.

An *entry-id* within an IP filter can only be mirrored to a single mirror destination. If the same *entry-id* is defined multiple times, an error occurs and only the first **mirror-source** definition is in effect.

By default, no packets matching any IP filters are mirrored. Mirroring of IP filter entries must be explicitly defined.

The **no ip-filter** command, without the **entry** keyword, removes mirroring on all *entry-id*'s within the *ip-filter-id*.

When the **no** command is executed with the **entry** keyword and one or more *entry-id*'s, mirroring of that list of *entry-id*'s is terminated within the *ip-filter-id*. If an *entry-id* is listed that does not exist, an error will occur

and the command will not execute. If an *entry-id* is listed that is not currently being mirrored, no error will occur for that *entry-id* and the command will execute normally.

Parameters

ip-filter-id

Specifies the IP filter ID whose entries are mirrored. If the *ip-filter-id* does not exist, an error will occur and the command will not execute. Mirroring of packets will commence once the *ip-filter-id* is defined on a SAP or IP interface.

Values 1 to 65535

name, up to 64 characters

entry-id

Specifies the IP filter entries to use as match criteria for packet mirroring. The **entry** keyword begins a list of *entry-id*'s for mirroring. Multiple *entry-id* entries may be specified with a single command. Each *entry-id* must be separated by a space.

If an *entry-id* does not exist within the IP filter, an error occurs and the command will not execute.

If the filter's *entry-id* is renumbered within the IP filter definition, the old *entry-id* is removed but the new *entry-id* must be manually added to the configuration to include the new (renumbered) entry's criteria.

Values 1 to 2097151

Platforms

7705 SAR Gen 2

ip-filter

Syntax

ip-filter *ip-filter-id* **entry** *entry-id* [*entry-id*]

no ip-filter *ip-filter-id* [**entry** *entry-id*]

Context

[Tree] (debug>mirror-source ip-filter)

Full Context

debug mirror-source ip-filter

Description

This command enables mirroring of packets that match specific entries in an existing IP filter.

The **ip-filter** command directs packets which match the defined list of entry IDs to be mirrored to the mirror destination referenced by the *mirror-dest-service-id* of the **mirror-source**.

The IP filter must already exist in order for the command to execute. Filters are configured in the **config>filter** context. If the IP filter does not exist, an error will occur. If the filter exists but has not been associated with a SAP or IP interface, an error is not generated but mirroring will not be enabled (there are no packets to mirror). Once the IP filter is defined to a SAP or IP interface, mirroring is enabled.

If the IP filter is defined as ingress, only ingress packets are mirrored. Ingress mirrored packets are mirrored to the mirror destination prior to any ingress packet modifications.

If the IP filter is defined as egress, only egress packets are mirrored. Egress mirrored packets are mirrored to the mirror destination after all egress packet modifications.

An *entry-id* within an IP filter can only be mirrored to a single mirror destination. If the same *entry-id* is defined multiple times, an error occurs and only the first **mirror-source** definition is in effect.

By default, no packets matching any IP filters are mirrored. Mirroring of IP filter entries must be explicitly defined.

The **no ip-filter** command, without the **entry** keyword, removes mirroring on all *entry-id*'s within the *ip-filter-id*.

When the **no** command is executed with the **entry** keyword and one or more *entry-id*'s, mirroring of that list of *entry-id*'s is terminated within the *ip-filter-id*. If an *entry-id* is listed that does not exist, an error will occur and the command will not execute. If an *entry-id* is listed that is not currently being mirrored, no error will occur for that *entry-id* and the command will execute normally.

Parameters

ip-filter-id

The IP filter ID whose entries are mirrored. If the *ip-filter-id* does not exist, an error will occur and the command will not execute. Mirroring of packets will commence once the *ip-filter-id* is defined on a SAP or IP interface.

entry-id

The IP filter entries to use as match criteria for packet mirroring. The **entry** keyword begins a list of *entry-id*'s for mirroring. A maximum of eight *entry-id* entries may be specified with a single command. Each *entry-id* must be separated by a space.

If an *entry-id* does not exist within the IP filter, an error occurs and the command will not execute.

If the filter's *entry-id* is renumbered within the IP filter definition, the old *entry-id* is removed but the new *entry-id* must be manually added to the configuration to include the new (renumbered) entry's criteria.

Platforms

7705 SAR Gen 2

ip-filter

Syntax

ip-filter *filter-id* [**name**] [**create**]

no ip-filter {*filter-id* | *filter-name*}

Context

[Tree] (config>filter ip-filter)

Full Context

configure filter ip-filter

Description

Commands in this context configure the specified IPv4 filter policy.

The **no** form of the command deletes the IPv4 filter policy. A filter policy cannot be deleted until it is removed from all objects where it is applied.

Parameters

filter-id

Specifies the IPv4 filter policy ID expressed as a decimal integer.

Values 1 to 65535

name

Configures an optional filter name, up to 64 characters in length, to a given filter. This filter name can then be used in configuration references, display, and show commands throughout the system. A defined filter name can help the service provider or administrator to identify and manage filters within the SR OS platforms.

To create a filter, you must assign a filter ID, however, after it is created, either the filter ID or filter name can be used to identify and reference a filter.

If a name is not specified at creation time, then SR OS assigns a string version of the *filter-id* as the name.

Filter names may not begin with an integer (0 to 9).

filter-name

Specifies a string, up to 64 characters, uniquely identifying this IPv4 filter policy.

create

This keyword is required to create the configuration context. Once it is created, the context can be enabled with or without the **create** keyword.

Platforms

7705 SAR Gen 2

ip-filter

Syntax

[**no**] **ip-filter**

Context

[Tree] (config>system>security>mgmt-access-filter ip-filter)

Full Context

configure system security management-access-filter ip-filter

Description

Commands in this context configure management access IP filter parameters.

Platforms

7705 SAR Gen 2

ip-filter

Syntax

ip-filter *src-filter-id* [**src-entry** *src-entry-id*] **to** *dst-filter-id* [**dst-entry** *dst-entry-id*] [**overwrite**]

Context

[Tree] (config>filter>copy ip-filter)

Full Context

configure filter copy ip-filter

Description

This command copies an existing filter entry for a specific filter ID to another filter ID. The command is a configuration level maintenance tool used to create new entries using an existing filter policy. If **overwrite** is not specified, an error will occur if the destination filter entry exists.

Parameters

src-filter-id

Identifies the source filter policy from which the copy command will attempt to copy. The filter policy must exist within the context of the preceding keyword (**ip-filter**).

dst-filter-id

Identifies the destination filter policy to which the copy command will attempt to copy. If the **overwrite** keyword is not specified, the filter entry ID cannot already exist in the destination filter policy. If the **overwrite** keyword is present, the destination entry ID may or may not exist.

overwrite

Specifies that the destination filter entry may exist. If it does, everything in the existing destination filter entry will be completely overwritten with the contents of the source filter entry. If the destination filter entry exists, either **overwrite** must be specified or an error message will be returned. If **overwrite** is specified, the function of copying from source to destination occurs in a "break before make" manner and therefore should be handled with care.

Platforms

7705 SAR Gen 2

13.120 ip-helper-address

ip-helper-address

Syntax

ip-helper-address *gateway-address*

no ip-helper-address

Context

[Tree] (config>service>ies>if ip-helper-address)

Full Context

configure service ies interface ip-helper-address

Description

This command enables broadcast UDP packets received on the associated interface to be redirected to the specified gateway address and then forwarded on to the gateway.

The **no** form of this command removes the gateway address from the interface configuration and stops the UDP broadcast redirect function.

Parameters

gateway-address

Specifies the IPv4 address of the target UDP broadcast gateway.

Platforms

7705 SAR Gen 2

ip-helper-address

Syntax

ip-helper-address *gateway-address*

no ip-helper-address

Context

[Tree] (config>service>vprn>if ip-helper-address)

Full Context

configure service vprn interface ip-helper-address

Description

This command enables broadcast UDP packets received on the associated interface to be redirected to the specified gateway address and then forwarded on to the gateway.

The **no** form of this command removes the gateway address from the interface configuration and stops the UDP broadcast redirect function.

Parameters

gateway-address

Specifies the IPv4 address of the target UDP broadcast gateway.

Platforms

7705 SAR Gen 2

ip-helper-address

Syntax

ip-helper-address *gateway-address*

no ip-helper-address

Context

[Tree] (config>router>if ip-helper-address)

Full Context

configure router interface ip-helper-address

Description

This command enables broadcast UDP packets received on the associated interface to be redirected to the specified gateway address and then forwarded on to the gateway.

The **no** form of this command removes the gateway address from the interface configuration and stops the UDP broadcast redirect function.

Parameters

gateway-address

Specifies the IPv4 address of the target UDP broadcast gateway.

Platforms

7705 SAR Gen 2

13.121 ip-mirror

ip-mirror

Syntax

ip-mirror

Context

[Tree] (config>mirror>mirror-dest>sap>egress ip-mirror)

Full Context

configure mirror mirror-dest sap egress ip-mirror

Description

This command configures IP mirror information.

Platforms

7705 SAR Gen 2

13.122 ip-mirror-interface

ip-mirror-interface

Syntax

ip-mirror-interface *ip-int-name* [**create**]

no ip-mirror-interface *ip-int-name*

Context

[Tree] (config>service>vprn ip-mirror-interface)

Full Context

configure service vprn ip-mirror-interface

Description

This command is used for remote mirroring, where the mirror source is a separate system then the mirror destination. The mirror source can only be of IP type and is only supported for the following services: IES, VPRN, VPLS and Ipipe. The mirror destination on a remote system will configure an interface on a VPRN as **ip-mirror-interface**. This interface only supports spoke sdp termination. The IP mirror interface requires PBR to determine the next outgoing interface for the mirror packet to be delivered to.

The **no** form of this command removes the interface name from the configuration.

Parameters

ip-int-name

Specifies the name of the IP interface, up to 32 characters. An interface name cannot be in the form of an IP address.

create

Keyword used to create an IP mirror interface.

Platforms

7705 SAR Gen 2

13.123 ip-mtu

ip-mtu

Syntax

ip-mtu *octets*

no ip-mtu

Context

[Tree] (config>service>vprn>if ip-mtu)

Full Context

configure service vprn interface ip-mtu

Description

This command specifies the maximum size of IP packets on this group interface. Packets larger than this are fragmented.

The **ip-mtu** applies to all IPoE host types (DHCP, ARP, or static). For PPP/L2TP sessions, the **ip-mtu** is not considered for the MTU negotiation. The **ppp-mtu** in the PPP policy should be used instead.

The **no** form of this command reverts to the default.

Default

no ip-mtu

Parameters

octets

Specifies the largest frame size (in octets) that this interface can handle.

Values 512 to 9000

Platforms

7705 SAR Gen 2

ip-mtu

Syntax

ip-mtu *octets*

no ip-mtu

Context

[Tree] (config>service>ies>if ip-mtu)

[Tree] (config>service>ies>if>sap>ip-tunnel ip-mtu)

Full Context

configure service ies interface ip-mtu

configure service ies interface sap ip-tunnel ip-mtu

Description

This command configures the IP maximum transmit unit (packet) for this interface.

Because this connects a Layer 2 to a Layer 3 service, this parameter can be adjusted under the IES interface.

The MTU that is advertised from the IES size is:

MINIMUM((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))

By default (for Ethernet network interface) if no ip-mtu is configured it is (1568 - 14) = 1554.

The **no** form of this command returns the default value.

Default

no ip-mtu

Parameters

octets

Specifies the maximum number of octets that can be transmitted.

Values 512 to 9786 (for IES interface)
512 to 9000 (for ip-tunnel interface)

Platforms

7705 SAR Gen 2

ip-mtu

Syntax

ip-mtu *bytes*

no ip-mtu

Context

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel ip-mtu)

[Tree] (config>router>if>ipsec>ipsec-tunnel ip-mtu)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel ip-mtu)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel ip-mtu)

[Tree] (config>service>vprn>if>sap>ip-tunnel ip-mtu)

Full Context

configure service ies interface ipsec ipsec-tunnel ip-mtu

configure router interface ipsec ipsec-tunnel ip-mtu

configure service vprn interface sap ipsec-tunnel ip-mtu

configure service vprn interface ipsec ipsec-tunnel ip-mtu

configure service vprn interface sap ip-tunnel ip-mtu

Description

This command configures the IP maximum transmit unit (packet) for this interface.

Because this connects a Layer 2 to a Layer 3 service, this parameter can be adjusted under the IES interface.

The MTU that is advertised from the IES size is:

MINIMUM((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))

By default (for the Ethernet network interface), if no ip-mtu is configured it is (1568 - 14) equals 1554.

The **ip-mtu** command instructs the MS-ISA to perform IP packet fragmentation, prior to IPsec encryption and encapsulation, based on the configured MTU value. In particular:

If the length of a payload IP packet (including its header) exceeds the configured MTU value and the DF flag is clear (due to the presence of the clear-df-bit command or because the original DF value was 0) then the MS-ISA fragments the payload packet as efficiently as possible (i.e. it creates the minimum number of fragments each less than or equal to the configured MTU size); in each created fragment the DF bit shall be 0.

If the length of a payload IP packet (including its header) exceeds the configured MTU value and the DF flag is set (because the original DF value was 1 and the tunnel has no clear-df-bit in its configuration) then the MS-ISA discards the payload packet without sending an ICMP type 3/code 4 message back to the packet's source address.

The effective MTU for packets entering a tunnel is the minimum of the private tunnel SAP interface IP MTU value (used by the IOM) and the tunnel IP MTU value (configured using the above command and used

by the MS-ISA). To fragment IP packets larger than X bytes with DF set, rather than discarding them, the tunnel IP MTU should be set to X and the private tunnel SAP interface IP MTU should be set to a value larger than X.

The **no ip-mtu** command, corresponding to the default behavior, disables fragmentation of IP packets by the MS-ISA; all IP packets, regardless of size or DF bit setting, are allowed into the tunnel.

Default

no ip-mtu

Parameters

bytes

Specifies the IP maximum transmit unit (packet) for this interface.

Values 512 to 9000

Platforms

7705 SAR Gen 2

ip-mtu

Syntax

ip-mtu *octets*

no ip-mtu

Context

[Tree] (config>service>vprn>nw-if ip-mtu)

Full Context

configure service vprn network-interface ip-mtu

Description

This command configures the IP maximum transmit unit (packet) for the associated router IP interface.

The configured IP-MTU cannot be larger than the calculated IP MTU based on the port MTU configuration.

The MTU that is advertised from the IES size is:

MINIMUM((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))

The **no** form of this command returns the associated IP interfaces MTU to its default value, which is calculated based on the port MTU setting. For Ethernet ports this will typically be 1554.

Default

no ip-mtu

Parameters

octets

Specifies the octets.

Values 512 to 9786

Platforms

7705 SAR Gen 2

ip-mtu

Syntax

ip-mtu *octets*

no ip-mtu

Context

[Tree] (config>ipsec>tnl-temp ip-mtu)

Full Context

configure ipsec tunnel-template ip-mtu

Description

This command configures the template IP MTU.

Default

no ip-mtu

Parameters

octets

Specifies the maximum size in octets.

Values 512 to 9000

Platforms

7705 SAR Gen 2

ip-mtu

Syntax

ip-mtu *octets*

no ip-mtu

Context

[Tree] (config>router>if ip-mtu)

Full Context

configure router interface ip-mtu

Description

This command configures the IP maximum transmit unit (packet) for the associated router IP interface.

The operational IP MTU that is used for the interface is determined based on both the configured IP MTU and the port MTU of the port bound to this interface.

The MTU that is used is:

MINIMUM((Port_MTU - EthernetHeaderSize), (configured ip-mtu))

The **no** form of this command returns the associated IP interfaces MTU to its default value, which is calculated based on the port MTU setting. (For Ethernet ports the default IP MTU is 1500 octets.)

Default

no ip-mtu

Parameters

octets

Specifies the IP MTU value associated with the IP interface, specified in octets. If the interface supports IPv6 packets, the IP-MTU must be set to a value greater than or equal to (≥) 1280 in accordance with RFC 2460 *Internet Protocol, Version 6 (IPv6) Specification*.

Values 512 to 9786

Platforms

7705 SAR Gen 2

ip-mtu

Syntax

ip-mtu *octets*

no ip-mtu

Context

[Tree] (bof ip-mtu)

Full Context

bof ip-mtu

Description

This command configures the IP maximum transmit unit (packet) for the management router instance.

The operational IP MTU that is used for the interface is determined based on both the configured IP MTU and the port MTU of the port bound to this interface.

The MTU that is used is:

MINIMUM((Port_MTU - EthernetHeaderSize), (configured ip-mtu))

For the management port, the port MTU is fixed at 1514 and the EthernetHeaderSize is 14 so the first element of the equation above is 1500 octets.

The **no** form of this command returns the associated IP interfaces MTU to its default value, which is calculated based on the port MTU setting. (For the management port the default IP MTU is 1500 octets.)

Default

ip-mtu 1500

Parameters

octets

Specifies the IP MTU value associated with the IP interface, specified in octets. If the interface supports IPv6 packets, the IP-MTU must be set to a value greater than or equal to (≥) 1280 in accordance with RFC 2460 *Internet Protocol, Version 6 (IPv6) Specification*.

Values 512 to 9786

Platforms

7705 SAR Gen 2

13.124 ip-prefix-list

ip-prefix-list

Syntax

ip-prefix-list *ip-prefix-list-name* [**create**]

no ip-prefix-list *ip-prefix-list-name*

Context

[Tree] (config>qos>match-list ip-prefix-list)

Full Context

configure qos match-list ip-prefix-list

Description

This command creates a list of IPv4 prefixes for match criteria in QoS policies.

An IP prefix list must contain only IPv4 address prefixes created using the prefix command and cannot be deleted if it is referenced by a QoS policy.

The **no** form of this command deletes the specified list.

Parameters

ip-prefix-list-name

A string of up to 32 characters of printable ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. The name **default** (case insensitive) is reserved by the system.

Platforms

7705 SAR Gen 2

ip-prefix-list

Syntax

ip-prefix-list *ip-prefix-list-name* [**create**]

no ip-prefix-list *ip-prefix-list-name*

Context

[Tree] (config>filter>match-list ip-prefix-list)

Full Context

configure filter match-list ip-prefix-list

Description

This command creates a list of IPv4 prefixes for match criteria in IPv4 ACL and CPM filter policies.

The **no** form of this command deletes the specified list.

Operational Notes:

An **ip-prefix-list** must contain only IPv4 address prefixes.

An IPv4 prefix match list cannot be deleted if it is referenced by a filter policy.

See general description related to match-list usage in filter policies.

Parameters

ip-prefix-list-name

Specifies a string of up to 32 printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

13.125 ip-prefix-routes

ip-prefix-routes

Syntax

ip-prefix-routes

Context

[Tree] (config>service>system>bgp-evpn ip-prefix-routes)

Full Context

configure service system bgp-evpn ip-prefix-routes

Description

Commands in this context configure attribute uniform propagation and BGP path selection.

Platforms

7705 SAR Gen 2

13.126 ip-route-advertisement

ip-route-advertisement

Syntax

ip-route-advertisement [**incl-host**] [**domain-id** *global-field*:*local-field*]

no ip-route-advertisement

Context

[Tree] (config>service>vpls>bgp-evpn ip-route-advertisement)

Full Context

configure service vpls bgp-evpn ip-route-advertisement

Description

This command enables and disables the advertisement of IP prefixes in EVPN. If enabled, any active route in the R-VPLS VPRN route table are advertised in EVPN using the VPLS BGP configuration. The interface host addresses are not advertised in EVPN unless the **ip-route-advertisement incl-host** command is enabled.

The **no** form of this command disables IP prefixes advertisement in EVPN.

Default

no ip-route-advertisement

Parameters

incl-host

Specifies to advertise the interface host addresses in EVPN.

global-field:local-field

Specifies the domain ID.

Values *4byte-GlobalAdminValue:2byte-LocalAdminValue*

4byte-GlobalAdminValue:	0 to 4294967295
2byte-LocalAdminValue	0 to 65535

Platforms

7705 SAR Gen 2

13.127 ip-route-link-bandwidth

ip-route-link-bandwidth

Syntax

ip-route-link-bandwidth

Context

[Tree] (config>service>vpls>bgp-evpn ip-route-link-bandwidth)

Full Context

configure service vpls bgp-evpn ip-route-link-bandwidth

Description

Commands in this context configure the IP route link bandwidth.

Platforms

7705 SAR Gen 2

13.128 ip-tunnel

ip-tunnel

Syntax

ip-tunnel *name* [**create**]

no ip-tunnel *name*

Context

[Tree] (config>service>vprn>if>sap ip-tunnel)

[Tree] (config>service>ies>if>sap ip-tunnel)

Full Context

configure service vprn interface sap ip-tunnel

configure service ies interface sap ip-tunnel

Description

This command is used to configure an IP-GRE or IP-IP tunnel and associate it with a private tunnel SAP within an IES or VPRN service.

The **no** form of this command deletes the specified IP/GRE or IP-IP tunnel from the configuration. The tunnel must be administratively shutdown before issuing the **no ip-tunnel** command.

Default

no-ip tunnel *name*

Parameters

name

Specifies the name of the IP tunnel. Tunnel names can be from 1 to 32 alphanumeric characters. If the string contains special characters (for example, #, $, spaces), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

13.129 ipoe

ipoe

Syntax

ipoe

Context

[Tree] (config>subscr-mgmt>loc-user-db ipoe)

Full Context

configure subscriber-mgmt local-user-db ipoe

Description

Commands in this context configure IPoE host parameters.

Platforms

7705 SAR Gen 2

13.130 ipsec

ipsec

Syntax

ipsec

Context

[Tree] (admin ipsec)

Full Context

admin ipsec

Description

Commands in this context perform Internet Protocol Security (IPsec) operations. IPsec is a structure of open standards to ensure private, secure communications over Internet Protocol (IP) networks by using cryptographic security services.

Platforms

7705 SAR Gen 2

ipsec

Syntax

[**no**] **ipsec**

Context

[Tree] (config>redundancy>multi-chassis>peer>sync ipsec)

Full Context

configure redundancy multi-chassis peer sync ipsec

Description

This command enables multi-chassis synchronization of IPsec states on system level.

Default

no ipsec

Platforms

7705 SAR Gen 2

ipsec

Syntax

ipsec

Context

[Tree] (config ipsec)

Full Context

configure ipsec

Description

Commands in this context configure Internet Protocol Security (IPsec) parameters. IPsec is a structure of open standards to ensure private, secure communications over Internet Protocol (IP) networks by using cryptographic security services.

Platforms

7705 SAR Gen 2

ipsec

Syntax

ipsec [**tunnel-group** *ipsec-group-id*] [**public-sap** *public-sap*]

no ipsec

Context

[Tree] (config>router>if ipsec)

[Tree] (config>service>ies>if ipsec)

[Tree] (config>service>vprn ipsec)

Full Context

configure router interface ipsec

configure service ies interface ipsec

configure service vprn ipsec

Description

Commands in this context configure IPsec policies.

Parameters

ipsec-group-id

Specifies the IPsec group ID used for the IPsec tunnels configured under this context.

Values 1 to 16

public-sap

Specifies the public SAP ID used for the IPsec tunnels configured under this context.

Values 0 to 4096

Platforms

7705 SAR Gen 2

13.131 ipsec-domain

ipsec-domain

Syntax

ipsec-domain *ipsec-domain-id* [**create**]

no ipsec-domain *ipsec-domain-id*

Context

[Tree] (config>redundancy>multi-chassis ipsec-domain)

Full Context

configure redundancy multi-chassis ipsec-domain

Description

Commands in this context configure parameters for the multi-chassis IPsec domain configured on this system.

The **no** form of this command removes the ID from the configuration.

Parameters

ipsec-domain-id

Specifies IPsec domain ID.

Values 1 to 255

create

Keyword used to create the command instance.

Platforms

7705 SAR Gen 2

13.132 ipsec-gw

ipsec-gw

Syntax

ipsec-gw *name*

no ipsec-gw

Context

[Tree] (config>service>vprn>if>sap ipsec-gw)

[Tree] (config>service>ies>if>sap ipsec-gw)

Full Context

configure service vprn interface sap ipsec-gw

configure service ies interface sap ipsec-gw

Description

This command configures an IPsec gateway.

Platforms

7705 SAR Gen 2

13.133 ipsec-lifetime

ipsec-lifetime

Syntax

ipsec-lifetime *ipsec-lifetime*

no ipsec-lifetime

Context

[Tree] (config>ipsec>ike-policy ipsec-lifetime)

Full Context

configure ipsec ike-policy ipsec-lifetime

Description

This command specifies the lifetime of the Phase 2 IKE key.

The **no** form of this command reverts to the default, which is 3600 seconds.

Default

no ipsec-lifetime

Parameters

ipsec-lifetime

Specifies the Phase 2 lifetime for this IKE policy in seconds.

Values 1200 to 31536000

Platforms

7705 SAR Gen 2

ipsec-lifetime

Syntax

ipsec-lifetime *seconds*

ipsec-lifetime inherit

Context

[Tree] (config>ipsec>ipsec-transform ipsec-lifetime)

Full Context

configure ipsec ipsec-transform ipsec-lifetime

Description

This command specifies the CHILD_SA. If the **inherit** parameter is specified, then the system uses the IPsec lifetime configuration in the corresponding IKE policy configured in the same IPsec gateway or IPsec tunnel.

Default

ipsec-lifetime inherit

Parameters

seconds

Specifies the lifetime of the Phase 2 IKE key in seconds.

Values 1200 to 31536000

inherit

Specifies that the system uses the **ipsec-lifetime** configuration in the corresponding IKE policy that is configured for the same IPsec gateway or IPsec tunnel.

Platforms

7705 SAR Gen 2

13.134 ipsec-responder-only

ipsec-responder-only

Syntax

[**no**] **ipsec-responder-only**

Context

[Tree] (config>isa>tunnel-group ipsec-responder-only)

Full Context

configure isa tunnel-group ipsec-responder-only

Description

With this command configured, system will only act as IKE responder except for the automatic CHILD_SA re-key upon MC-IPsec switchover.

Default

no ipsec-responder-only

Platforms

7705 SAR Gen 2

13.135 ipsec-transform

ipsec-transform

Syntax

ipsec-transform *transform-id* [**create**]

no ipsec-transform *transform-id*

Context

[Tree] (config>ipsec ipsec-transform)

Full Context

configure ipsec ipsec-transform

Description

Commands in this context create an **ipsec-transform** policy. IPsec transforms policies can be shared. A change to the ipsec-transform is allowed at any time. The change will not impact tunnels that have been established until they are renegotiated. If the change is required immediately the tunnel must be cleared (reset) for force renegotiation.

IPsec transform policy assignments to a tunnel require the tunnel to be shutdown.

The **no** form of this command removes the ID from the configuration.

Parameters

transform-id

Specifies a policy ID value to identify the IPsec transform policy.

Values 1 to 2048

create

This keyword is mandatory when creating an ipsec-transform policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

13.136 ipsec-transport-mode-profile

ipsec-transport-mode-profile

Syntax

ipsec-transport-mode-profile *name* [**create**]

no ipsec-transport-mode-profile *name*

Context

[Tree] (config>ipsec ipsec-transport-mode-profile)

Full Context

configure ipsec ipsec-transport-mode-profile

Description

Commands in this context configure an IPsec transport mode profile.

The **no** form of this command removes the name from the configuration.

Parameters

name

Specifies the name of the IPsec transport mode profile, up to 32 characters.

create

Keyword used to create the IPsec transport mode profile instance.

Platforms

7705 SAR Gen 2

ipsec-transport-mode-profile

Syntax

ipsec-transport-mode-profile *name*

no ipsec-transport-mode-profile

Context

[Tree] (config>service>ies>if>sap>ip-tunnel ipsec-transport-mode-profile)

[Tree] (config>service>vprn>if>sap>ip-tunnel ipsec-transport-mode-profile)

Full Context

configure service ies interface sap ip-tunnel ipsec-transport-mode-profile

configure service vprn interface sap ip-tunnel ipsec-transport-mode-profile

Description

This command specifies an IPsec transport mode profile name to the SAP.

The **no** form of this command removes the profile name from the service configuration.

Parameters

name

Specifies the name of an existing IPsec transport mode profile, up to 32 characters

Platforms

7705 SAR Gen 2

13.137 ipsec-tunnel

ipsec-tunnel

Syntax

ipsec-tunnel *ipsec-tunnel-name*

no ipsec-tunnel [*ipsec-tunnel-name*]

Context

[Tree] (config>service>vprn>static-route-entry ipsec-tunnel)

Full Context

configure service vprn static-route-entry ipsec-tunnel

Description

This command creates a static route in a VPRN service context that points to the global routing context (base router). This is primarily used to allow traffic that ingress through a VPRN service to be routed out of the global routing context.

This **next-hop** type cannot be used in conjunction with any other next-hop types.

Default

no ipsec-tunnel

Parameters

ipsec-tunnel-name

IPsec tunnel name; maximum length up to 32 characters.

Platforms

7705 SAR Gen 2

ipsec-tunnel

Syntax

ipsec-tunnel *name* [**private-sap** [*0..4094*]] [**private-service-name** *private-service-name*] [**create**]

no ipsec-tunnel *ipsec-tunnel-name*

Context

[Tree] (config>service>vprn>if>sap ipsec-tunnel)

[Tree] (config>router>if>ipsec ipsec-tunnel)

[Tree] (config>service>ies>if>ipsec ipsec-tunnel)

Full Context

configure service vprn interface sap ipsec-tunnel

configure router interface ipsec ipsec-tunnel

configure service ies interface ipsec ipsec-tunnel

Description

This command configures a secured interface IPsec tunnel. If the **private-service-name** is not specified, the private service is the secured interface service.

The **no** form of this command removes the IPsec tunnel from the configuration.

Parameters

name

Specifies the name of the IPsec tunnel.

private-sap

Specifies the private SAP ID.

Values 0 to 4094

private-service-name

Specifies the private service name.

create

Keyword used to create the IPsec tunnel instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

13.138 ipv4

ipv4

Syntax

ipv4 *max-paths* [**ebgp** *ebgp-max-paths*] [**ibgp** *ibgp-max-paths*] [**restrict** { **same-neighbor-as** | **exact-as-path**}] [**unequal-cost**]

no ipv4

Context

[Tree] (config>service>vprn>bgp>multi-path ipv4)

Full Context

configure service vprn bgp multi-path ipv4

Description

This command sets ECMP multipath parameters that apply only to the (unlabeled) IPv4 unicast address family. These settings override the values set by the **maximum-paths** command.

When multipath is enabled, traffic to the destination is load-shared across a set of paths (BGP routes) that the BGP decision process considers equal to the best path. The actual distribution of traffic over the multiple paths may be equal or unequal (that is, based on weights derived from the Link Bandwidth Extended Community).

To qualify as a multipath, a non-best route must meet the following criteria (some criteria are controlled by this command):

- The multi-path route must be the same type of route as the best path (same AFI/SAFI and, in some cases, same next-hop resolution method).
- The multi-path route must be tied with the best path for all criteria of greater significance than next-hop cost, except for criteria that are configured to be ignored.
- If the best path selection reaches the next-hop cost comparison, the multi-path route must have the same next-hop cost as the best route unless the **unequal-cost** option is configured.
- The multi-path route must not have the same BGP next-hop as the best path or any other multi-path route.
- The multi-path route must not cause the ECMP limit of the routing instance to be exceeded (configured using the **ecmp** command with a value in the range 1 to 64).
- The multi-path route must not cause the applicable *max-paths* limit to be exceeded. If the best path is an EBGP learned route and the **ebgp** option is used, the *ebgp-max-paths* limit overrides the *max-paths* limit. If the best path is an IBGP-learned route and the **ibgp** option is used, the *ibgp-max-paths* limit overrides the *max-paths* limit. All path limits are configurable up to a maximum of 64. Multi-path is effectively disabled if a value is set to 1.
- The multi-path route must have the same neighbor AS in its AS path as the best path if the **restrict same-neighbor-as** option is configured. By default, any path with the same AS path length as the best path (regardless of neighbor AS) is eligible for multi-path.

- The route must have the same AS path as the best path if the **restrict exact-as-path** option is configured. By default, any path with the same AS path length as the best path (regardless of the actual AS numbers) is eligible for multi-path.

The **no** form of this command removes IPv4-specific overrides.

Default

no ipv4

Parameters

max-paths

Specifies the maximum number of multipaths per prefix/NLRI if *ebgp-max-paths* or *ibgp-max-paths* does not apply.

Values 1 to 64

egp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an EBGP learned route.

Values 1 to 64

ibgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an IBGP learned route.

Values 1 to 64

restrict same-neighbor-as

Specifies that the non-best path must have the same neighbor AS in its AS path as the best path.

restrict exact-as-path-as

Specifies that the non-best path must have the same AS path as the best path.

unequal-cost

Instructs BGP to ignore differences in the next-hop cost only when determining eligible multipaths.

Platforms

7705 SAR Gen 2

ipv4

Syntax

ipv4

Context

[Tree] (config>router>ldp>if-params>if ipv4)

Full Context
configure router ldp interface-parameters interface ipv4

Description
Commands in this context configure LDP interfaces and parameters applied to an IPv4 LDP interface.

Platforms
7705 SAR Gen 2

ipv4

Syntax
[**no**] **ipv4**

Context
[Tree] (config>router>ldp>if-params ipv4)

Full Context
configure router ldp interface-parameters ipv4

Description
Commands in this context configure IPv4 LDP parameters applied to the interface.

Platforms
7705 SAR Gen 2

ipv4

Syntax
ipv4

Context
[Tree] (config>router>ldp>targeted-session ipv4)

Full Context
configure router ldp targeted-session ipv4

Description
Commands in this context configure parameters applied to targeted sessions to all IPv4 LDP peers.

Platforms
7705 SAR Gen 2

ipv4

Syntax

ipv4

Context

[Tree] (config>router>ldp>targeted-session>auto-tx ipv4)

[Tree] (config>router>ldp>targeted-session>auto-rx ipv4)

Full Context

configure router ldp targeted-session auto-tx ipv4

configure router ldp targeted-session auto-rx ipv4

Description

Commands in this context configure IPv4 parameters of an automatic targeted LDP session.

Platforms

7705 SAR Gen 2

ipv4

Syntax

ipv4

Context

[Tree] (config>service>ies>if>sap>ipsec-gw>lcl-addr-assign ipv4)

[Tree] (config>service>vprn>if>sap>ipsec-gw>lcl-addr-assign ipv4)

Full Context

configure service ies interface sap ipsec-gw local-address-assignment ipv4

configure service vprn interface sap ipsec-gw local-address-assignment ipv4

Description

Commands in this context configure IPv4 local address assignment parameters for the IPsec gateway.

Platforms

7705 SAR Gen 2

ipv4

Syntax

ipv4

Context

[Tree] (bof>autoconfigure ipv4)

Full Context

bof autoconfigure ipv4

Description

Commands in this context autoconfigure the IPv4 DHCP client.

Platforms

7705 SAR Gen 2

ipv4

Syntax

ipv4 *ip-address*

ipv4 auto-generate [**vendor-id-value** *vendor-id-value*]

no ipv4

Context

[Tree] (config>system>ned>prof>neip ipv4)

Full Context

configure system network-element-discovery profile neip ipv4

Description

This command configures the IPv4 NEIP for this profile. The NEIP can be configured manually or set to be automatically generated using the NEID. If the NEID option is set, the first most significant byte of the IPv4 NEIP is set to 140 and the remaining 3 bytes are set to the NEID value. The NEID can be configured with a vendor ID value, in which case the first most significant byte of the IPv4 NEIP is set to this vendor ID value.

The **no** form of this command removes the IPv4 address association for this profile.

Default

no ipv4

Parameters

ip-address

Specifies the IPv4 address of the NEIP.

auto-generate

Specifies that the NEIP is automatically generated using the NEID.

vendor-id-value

Specifies the vendor ID value.

Values 1 to 255

Platforms

7705 SAR Gen 2

ipv4

Syntax

ipv4 send *send-limit* **receive** [**none**]

ipv4 send *send-limit*

no ipv4

Context

[Tree] (config>router>bgp>group>add-paths ipv4)

[Tree] (config>router>bgp>add-paths ipv4)

[Tree] (config>router>bgp>group>neighbor>add-paths ipv4)

Full Context

configure router bgp group add-paths ipv4

configure router bgp add-paths ipv4

configure router bgp group neighbor add-paths ipv4

Description

This command configures the add-paths capability for unlabeled IPv4 unicast routes. By default, add-paths is not enabled for unlabeled IPv4 unicast routes.

The maximum number of unlabeled unicast paths per IPv4 prefix to send is the configured send limit, which is a mandatory parameter. The capability to receive multiple unlabeled IPv4 unicast paths per prefix from a peer is configurable using the **receive** keyword, which is optional. If the **receive** keyword is not included in the command the receive capability is enabled by default.

The **no** form of this command disables add-paths support for unlabeled IPv4 unicast routes, causing sessions established using add-paths for unlabeled IPv4 unicast to go down and come back up without the add-paths capability.

Default

no ipv4

Parameters

send-limit

Specifies the maximum number of paths per unlabeled IPv4 unicast prefix that are allowed to be advertised to add-paths peers, the actual number of advertised routes may be less. If the value is **none**, the router does not negotiate the send capability with respect to IPv4 AFI/SAFI. If the value is **multipaths**, then BGP advertises all of the used BGP multipaths for each IPv4 NLRI if the peer has signaled support to receive multiple add-paths.

Values 1 to 16, none, multipaths

receive

Specifies that the router negotiates to receive multiple unlabeled unicast routes per IPv4 prefix.

none

Specifies that the router does not negotiate to receive multiple unlabeled unicast routes per IPv4 prefix.

Platforms

7705 SAR Gen 2

ipv4

Syntax

ipv4 *max-paths* [**ebgp** *ebgp-max-paths*] [**ibgp** *ibgp-max-paths*] [**restrict** { **same-neighbor-as** | **exact-as-path**}] [**unequal-cost**]

no ipv4

Context

[Tree] (config>router>bgp>multi-path ipv4)

Full Context

configure router bgp multi-path ipv4

Description

This command sets ECMP multipath parameters that apply only to the (unlabeled) IPv4 unicast address family. These settings override the values set by the **maximum-paths** command.

When multipath is enabled, traffic to the destination is load-shared across a set of paths (BGP routes) that the BGP decision process considers equal to the best path. The actual distribution of traffic over the multiple paths may be equal or unequal (that is, based on weights derived from the Link Bandwidth Extended Community).

The **no** form of this command removes IPv4-specific overrides.

Default

no ipv4

Parameters

max-paths

Specifies the maximum number of multipaths per prefix/NLRI if *ebgp-max-paths* or *ibgp-max-paths* does not apply.

Values 1 to 64

ebgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an EBGP learned route.

Values 1 to 64

ibgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an IBGP learned route.

Values 1 to 64

restrict same-neighbor-as

Specifies that the non-best path must have the same neighbor AS in its AS path as the best path.

restrict exact-as-path

Specifies that the non-best path must have the same AS path as the best path.

unequal-cost

Instructs BGP to ignore differences in the next-hop cost only when determining eligible multipaths.

Platforms

7705 SAR Gen 2

13.139 ipv4-adjacency-sid

ipv4-adjacency-sid

Syntax

ipv4-adjacency-sid label *value*

no ipv4-adjacency-sid

Context

[Tree] (config>router>isis>interface ipv4-adjacency-sid)

Full Context

configure router isis interface ipv4-adjacency-sid

Description

This command allows a static value to be assigned to an IPv4 adjacency SID in IS-IS segment routing.

The **label** option specifies that the value is assigned to an MPLS label.

The **no** form of this command removes the adjacency SID.

Parameters

value

Specifies the adjacency SID label.

Values 18432 to 5248 | 1048575 (FP4 or FP5 only)

Platforms

7705 SAR Gen 2

13.140 ipv4-multicast

ipv4-multicast

Syntax

[**no**] **ipv4-multicast**

Context

[Tree] (config>service>vprn>isis>multi-topology ipv4-multicast)

Full Context

configure service vprn isis multi-topology ipv4-multicast

Description

This command enables support for the IPv4 topology (MT3) within the associate IS-IS instance.

The **no** form of this command disables support for the IPv4 topology (MT3) within the associated IS-IS instance.

Default

no ipv4-multicast

Platforms

7705 SAR Gen 2

ipv4-multicast

Syntax

[**no**] **ipv4-multicast**

Context

[Tree] (config>router>isis>multi-topology ipv4-multicast)

Full Context

configure router isis multi-topology ipv4-multicast

Description

This command enables support for the IPv4 topology (MT3) within the associate IS-IS instance.

The **no** form of this command disables support for the IPv4 topology (MT3) within the associated IS-IS instance.

Default

no ipv4-multicast

Platforms

7705 SAR Gen 2

13.141 ipv4-multicast-disable

ipv4-multicast-disable

Syntax

[**no**] **ipv4-multicast-disable**

Context

[Tree] (config>service>vprn>isis>if ipv4-multicast-disable)

Full Context

configure service vprn isis interface ipv4-multicast-disable

Description

This command administratively disables/enables ISIS operation for IPv4.

Default

no ipv4-multicast-disable

Platforms

7705 SAR Gen 2

ipv4-multicast-disable

Syntax

[**no**] **ipv4-multicast-disable**

Context

[Tree] (config>service>vprn>pim ipv4-multicast-disable)

[Tree] (config>service>vprn>pim>if ipv4-multicast-disable)

Full Context

configure service vprn pim ipv4-multicast-disable

configure service vprn pim interface ipv4-multicast-disable

Description

This command administratively disables/enables PIM operation for IPv4.

Default

no ipv4-multicast-disable

Platforms

7705 SAR Gen 2

ipv4-multicast-disable

Syntax

[**no**] **ipv4-multicast-disable**

Context

[Tree] (config>router>pim ipv4-multicast-disable)

[Tree] (config>router>pim>interface ipv4-multicast-disable)

Full Context

configure router pim ipv4-multicast-disable

configure router pim interface ipv4-multicast-disable

Description

This command administratively enables PIM operation for IPv4.

IPv4 multicast must be enabled to enable MLDP in-band signaling for IPv4 PIM joins; see **config>router>pim>interface p2mp-ldp-tree-join**.

The **no** form of this command disables the PIM operation for IPv4.

Default

no ipv4-multicast-disable

Platforms

7705 SAR Gen 2

ipv4-multicast-disable

Syntax

[**no**] **ipv4-multicast-disable**

Context

[Tree] (config>router>isis>interface ipv4-multicast-disable)

Full Context

configure router isis interface ipv4-multicast-disable

Description

This command disables IS-IS IPv4 multicast routing for the interface.

The **no** form of this command enables IS-IS IPv4 multicast routing for the interface.

Platforms

7705 SAR Gen 2

13.142 ipv4-multicast-metric

ipv4-multicast-metric

Syntax

ipv4-multicast-metric *metric*

no ipv4-multicast-metric

Context

[Tree] (config>service>vprn>isis>if>level ipv4-multicast-metric)

Full Context

configure service vprn isis interface level ipv4-multicast-metric

Description

This command configures IS-IS interface metric for IPv4 multicast for the VPRN instance.

The **no** form of this command removes the metric from the configuration.

Parameters

metric

Specifies the IS-IS interface metric for IPv4 multicast.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

ipv4-multicast-metric

Syntax

ipv4-multicast-metric *metric*

no ipv4-multicast-metric

Context

[Tree] (config>router>isis>if>level ipv4-multicast-metric)

Full Context

configure router isis interface level ipv4-multicast-metric

Description

This command configures the IS-IS interface metric for IPv4 multicast.

The **no** form of this command removes the metric from the configuration.

Parameters

metric

Specifies the IS-IS interface metric for IPv4 multicast.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

13.143 ipv4-multicast-metric-offset

ipv4-multicast-metric-offset

Syntax

ipv4-multicast-metric-offset *offset-value*

no ipv4-multicast-metric-offset

Context

[Tree] (config>service>vprn>isis>link-group>level ipv4-multicast-metric-offset)

Full Context

configure service vprn isis link-group level ipv4-multicast-metric-offset

Description

This command sets the offset value for the IPv4 multicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric for the IPv4 multicast topology

The **no** form of this command reverts the offset value to 0.

Default

no ipv4-multicast-metric-offset

Parameters

offset-value

Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold.

Values 0 to 6777215

Platforms

7705 SAR Gen 2

ipv4-multicast-metric-offset

Syntax

ipv4-multicast-metric-offset *offset-value*

no ipv4-multicast-metric-offset

Context

[Tree] (config>router>isis>link-group>level ipv4-multicast-metric-offset)

Full Context

configure router isis link-group level ipv4-multicast-metric-offset

Description

This command sets the offset value for the IPv4 multicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric for the IPv4 multicast topology.

The **no** form of this command reverts the offset value to 0.

Default

no ipv4-multicast-metric-offset

Parameters

offset-value

Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold

Values 0 to 6777215

Platforms

7705 SAR Gen 2

13.144 ipv4-multicast-routing

ipv4-multicast-routing

Syntax

ipv4-multicast-routing {**native** | **mt**}

[**no**] **ipv4-multicast-routing**

Context

[Tree] (config>service>vprn>isis ipv4-multicast-routing)

Full Context

configure service vprn isis ipv4-multicast-routing

Description

The multicast RTM is used for Reverse Path Forwarding checks. This command controls which IS-IS topology is used to populate the IPv4 multicast RTM.

The **no** ipv4-multicast-routing form of this command results in none of the IS-IS routes being populated in the IPv4 multicast RTM and would be used if multicast is configured to use the unicast RTM for the RPF check.

Default

ipv4-multicast-routing native

Parameters

native

Causes IPv4 routes from the MT0 topology to be added to the multicast RTM for RPF checks.

mt

Causes IPv4 routes from the MT3 topology to be added to the multicast RTM for RPF checks.

Platforms

7705 SAR Gen 2

ipv4-multicast-routing

Syntax

ipv4-multicast-routing {**native** | **mt**}

[**no**] **ipv4-multicast-routing**

Context

[Tree] (config>router>isis ipv4-multicast-routing)

Full Context

configure router isis ipv4-multicast-routing

Description

The multicast RTM is used for Reverse Path Forwarding checks. This command controls which IS-IS topology is used to populate the IPv4 multicast RTM.

The **no** form of this command results in none of the IS-IS routes being populated in the IPv4 multicast RTM and would be used if multicast is configured to use the unicast RTM for the RPF check.

Default

ipv4-multicast-routing native

Parameters

native

Causes IPv4 routes from the MT0 topology to be added to the multicast RTM for RPF checks.

mt

Causes IPv4 routes from the MT3 topology to be added to the multicast RTM for RPF checks.

Platforms

7705 SAR Gen 2

13.145 ipv4-node-sid

ipv4-node-sid

Syntax

ipv4-node-sid index *index-value* [**clear-n-flag**]

ipv4-node-sid label *label-value* [**clear-n-flag**]

no ipv4-node-sid

Context

[Tree] (config>router>isis>interface ipv4-node-sid)

Full Context

configure router isis interface ipv4-node-sid

Description

This command assigns a node SID index or label value to the prefix representing the primary address of an IPv4 network interface of **type loopback**. Only a single node SID can be assigned to an interface. The secondary address of an IPv4 interface cannot be assigned a node SID index and does not inherit the SID of the primary IPv4 address.

The command fails if the network interface is not of **type loopback** or if the interface is defined in an IES or a VPRN context. Also, assigning the same SID index or label value to the same interface in two different IGP instances is not allowed within the same node.

The value of the label or index SID is taken from the range configured for this IGP instance. When using the global mode of operation, a new segment routing module checks that the same index or label value cannot be assigned to more than one loopback interface address. When using the per-instance mode of operation, this check is not required since the index and thus label ranges of the various IGP instance are not allowed to overlap.

The **clear-n-flag** option allows the user to clear the N-flag (node-sid flag) in an IS-IS prefix SID sub-TLV originated for the IPv4 prefix of a loopback interface on the system.

By default, the prefix SID sub-TLV for the prefix of a loopback interface is tagged as a node SID, meaning that it belongs to this node only. However, when the user wants to configure and advertise an anycast SID using the same loopback interface prefix on multiple nodes, you must clear the N-flag to assure interoperability with third party implementations, which may perform a strict check on the receiving end and drop duplicate prefix SID sub-TLVs when the N-flag is set.

The SR OS implementation is relaxed on the receiving end and accepts duplicate prefix SIDs with the N-flag set or cleared. SR OS will resolve to the closest owner, or owners if ECMP is configured, of the prefix SID according to its cost.

Default

no ipv4-node-sid

Parameters

index index-value

Specifies the index value.

Values 0 to 4294967295

label label-value

Specifies the label value.

Values 0 to 4294967295

clear-n-flag

Clears the node SID flag.

Default no clear-n-flag

Platforms

7705 SAR Gen 2

13.146 ipv4-prefix

ipv4-prefix

Syntax

[**no**] **ipv4-prefix**

Context

[Tree] (debug>router>rpki-session>packet ipv4-prefix)

Full Context

debug router rpki-session packet ipv4-prefix

Description

This command enables debugging for IPv4 prefix RPKI packets.

The **no** form of this command disables debugging for IPv4 prefix RPKI packets.

Platforms

7705 SAR Gen 2

13.147 ipv4-routing

ipv4-routing

Syntax

[**no**] **ipv4-routing**

Context

[Tree] (config>service>vprn>isis ipv4-routing)

Full Context

configure service vprn isis ipv4-routing

Description

This command specifies whether this IS-IS instance supports IPv4.

The **no** form of this command disables IPv4 on the IS-IS instance.

Default

ipv4-routing

Platforms

7705 SAR Gen 2

ipv4-routing

Syntax

[**no**] **ipv4-routing**

Context

[Tree] (config>router>isis ipv4-routing)

Full Context

configure router isis ipv4-routing

Description

This command specifies whether this IS-IS instance supports IPv4.

The **no** form of this command disables IPv4 on the IS-IS instance.

Default

ipv4-routing

Platforms

7705 SAR Gen 2

13.148 ipv4-sid

ipv4-sid

Syntax

ipv4-sid index *index-id*

ipv4-sid label *label-id*

no ipv4-sid

Context

[Tree] (config>router>segment-routing>sr-mpls>prefix-sids ipv4-sid)

Full Context

configure router segment-routing sr-mpls prefix-sids ipv4-sid

Description

This command is used to configure the IPv4 segment routing SID associated with the primary IPv4 address of the loopback or system interface.

The **no** form of this command removes the configuration of the IPv4 segment routing SID associated with the primary IPv4 interface address.

Default

no ipv4-sid

Parameters

index *index-id*

Specifies the node SID index for this interface.

Values 0 to 4294967295

label *label-id*

Specifies the label value for the node SID.

Values 32 to 1048575

Platforms

7705 SAR Gen 2

13.149 ipv4-source-address

ipv4-source-address

Syntax

ipv4-source-address *ipv4-address*

no ipv4-source-address

Context

[Tree] (config>service>vprn>dns ipv4-source-address)

Full Context

configure service vprn dns ipv4-source-address

Description

This command configures the IPv4 address of the default secondary DNS server for the subscribers using this interface. Subscribers that cannot obtain an IPv4 DNS server address by other means, can use this for DNS name resolution.

The ipv4-address value can only be set to a nonzero value if the value of VPRN type is set to **subscriber-split-horizon**.

The **no** form of this command reverts to the default.

Parameters

ipv4-address

Specifies the IPv4 address of the default secondary DNS server.

Values ipv4-address - a.b.c.d

Platforms

7705 SAR Gen 2

ipv4-source-address

Syntax

ipv4-source-address *ip-address*

no ipv4-source-address

Context

[Tree] (config>system>file-trans-prof ipv4-source-address)

Full Context

configure system file-transmission-profile ipv4-source-address

Description

This command specifies the IPv4 source address used for transport protocol.

The **no** form of this command uses the default source address which typically is the address of the egress interface.

Default

no ipv4-source-address

Parameters

ip-address

Specifies a unicast v4 address. This should be a local interface address.

Platforms

7705 SAR Gen 2

13.150 ipv4-unicast-metric-offset

ipv4-unicast-metric-offset

Syntax

ipv4-unicast-metric-offset *offset-value*

no ipv4-unicast-metric-offset

Context

[Tree] (config>service>vprn>isis>link-group>level ipv4-unicast-metric-offset)

Full Context

configure service vprn isis link-group level ipv4-unicast-metric-offset

Description

This command sets the offset value for the IPv4 unicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric.

The **no** form of this command reverts the offset value to 0.

Default

no ipv4-unicast-metric-offset

Parameters

offset-value

Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold.

Values 0 to 6777215

Platforms

7705 SAR Gen 2

ipv4-unicast-metric-offset

Syntax

ipv4-unicast-metric-offset *offset-value*

no ipv4-unicast-metric-offset

Context

[Tree] (config>router>isis>link-group>level ipv4-unicast-metric-offset)

Full Context

configure router isis link-group level ipv4-unicast-metric-offset

Description

This command sets the offset value for the IPv4 unicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric.

The **no** form of this command reverts the offset value to 0.

Default

no ipv4-unicast-metric-offset

Parameters

offset-value

Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold.

Values 0 to 6777215

Platforms
7705 SAR Gen 2

13.151 ipv6

ipv6

Syntax
[no] ipv6

Context
[Tree] (config>service>vprn>if ipv6)
[Tree] (config>service>ies>if ipv6)

Full Context
configure service vprn interface ipv6
configure service ies interface ipv6

Description
Commands in this context configure IPv6 parameters for the interface.

Platforms
7705 SAR Gen 2

ipv6

Syntax
ipv6 *name*
no ipv6

Context
[Tree] (config>service>template>vpls-sap-template>egress>filter-name ipv6)
[Tree] (config>service>template>vpls-sap-template>ingress>filter-name ipv6)

Full Context
configure service template vpls-sap-template egress filter-name ipv6
configure service template vpls-sap-template ingress filter-name ipv6

Description

This command associates an existing IP filter policy with the template.

Parameters

name

Specifies the IPv6 filter policy name, up to 64 characters.

Platforms

7705 SAR Gen 2

ipv6

Syntax

ipv6 max-paths [**ebgp** *ebgp-max-paths*] [**ibgp** *ibgp-max-paths*] [**restrict** { **same-neighbor-as** | **exact-as-path**}] [**unequal-cost**]

no ipv6

Context

[Tree] (config>service>vprn>bgp>multi-path ipv6)

Full Context

configure service vprn bgp multi-path ipv6

Description

This command sets ECMP multipath parameters that apply only to the (unlabeled) IPv6 unicast address family. These settings override the values set by the **maximum-paths** command.

When multipath is enabled, traffic to the destination is load-shared across a set of paths (BGP routes) that the BGP decision process considers equal to the best path. The actual distribution of traffic over the multiple paths may be equal or unequal (that is, based on weights derived from the Link Bandwidth Extended Community).

To qualify as a multipath, a non-best route must meet the following criteria (some criteria are controlled by this command):

- The multi-path route must be the same type of route as the best path (same AFI/SAFI and, in some cases, same next-hop resolution method).
- The multi-path route must be tied with the best path for all criteria of greater significance than next-hop cost, except for criteria that are configured to be ignored.
- If the best path selection reaches the next-hop cost comparison, the multi-path route must have the same next-hop cost as the best route unless the **unequal-cost** option is configured.
- The multi-path route must not have the same BGP next-hop as the best path or any other multi-path route.
- The multi-path route must not cause the ECMP limit of the routing instance to be exceeded (configured using the **ecmp** command with a value in the range 1 to 64)

- The multi-path route must not cause the applicable *max-paths* limit to be exceeded. If the best path is an EBGP learned route and the **ebgp** option is used, the *ebgp-max-paths* limit overrides the *max-paths* limit. If the best path is an IBGP-learned route and the **ibgp** option is used, the *ibgp-max-paths* limit overrides the *max-paths* limit. All path limits are configurable up to a maximum of 64. Multi-path is effectively disabled if a value is set to 1.
- The multi-path route must have the same neighbor AS in its AS path as the best path if the **restrict same-neighbor-as** option is configured. By default, any path with the same AS path length as the best path (regardless of neighbor AS) is eligible for multi-path.
- The route must have the same AS path as the best path if the **restrict exact-as-path** option is configured. By default, any path with the same AS path length as the best path (regardless of the actual AS numbers) is eligible for multi-path.

The **no** form of this command removes IPv6-specific overrides.

Default

no ipv6

Parameters

max-paths

Specifies the maximum number of multipaths per prefix/NLRI if *ebgp-max-paths* or *ibgp-max-paths* does not apply.

Values 1 to 64

egp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an EBGP learned route.

Values 1 to 64

ibgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an IBGP learned route.

Values 1 to 64

restrict same-neighbor-as

Specifies that the non-best path must have the same neighbor AS in its AS path as the best path.

restrict exact-as-path-as

Specifies that the non-best path must have the same AS path as the best path.

unequal-cost

Instructs BGP to ignore differences in the next-hop cost only when determining eligible multipaths.

When enabled, the Alc-App-Prof-Str VSA is ignored in a radius Accept that enables portal redirection using this redirect policy. AA functionality will be disabled during portal authentication.

The **no** version of this command allows an Alc-App-Prof-Str to be present and will enable Application Assurance during portal authentication. In this case redirection rules defined in this policy are bypassed and it is assumed the AA function is configured for portal redirection.

Platforms

7705 SAR Gen 2

ipv6

Syntax

ipv6

Context

[Tree] (config>service>vprn>pim>rp ipv6)

Full Context

configure service vprn pim rp ipv6

Description

This command enables access to the context to configure the rendezvous point (RP) of a PIM IPv6 protocol instance.

A Nokia IPv6 PIM router acting as an RP must respond to an IPv6 PIM register message specifying an SSM multicast group address by sending to the first hop router stop register message(s). It does not build an (S, G) shortest path tree toward the first hop router. An SSM multicast group address can be either from the SSM default range or from a multicast group address range that was explicitly configured for SSM.

Default

ipv6 RP enabled when IPv6 PIM is enabled.

Platforms

7705 SAR Gen 2

ipv6

Syntax

[**no**] **ipv6**

Context

[Tree] (config>router>ldp>if-params>if ipv6)

Full Context

configure router ldp interface-parameters interface ipv6

Description

Commands in this context configure IPv6 LDP parameters applied to the interface.

Platforms

7705 SAR Gen 2

ipv6

Syntax

ipv6

Context

[Tree] (config>router>ldp>if-params ipv6)

Full Context

configure router ldp interface-parameters ipv6

Description

Commands in this context configure LDP interfaces and parameters applied to an IPv6 LDP interface.

Platforms

7705 SAR Gen 2

ipv6

Syntax

ipv6

Context

[Tree] (config>router>ldp>targeted-session ipv6)

Full Context

configure router ldp targeted-session ipv6

Description

Commands in this context configure parameters applied to targeted sessions to all IPv6 LDP peers.

Platforms

7705 SAR Gen 2

ipv6

Syntax

ipv6

Context

[Tree] (config>service>vprn>if>sap>ipsec-gw>lcl-addr-assign ipv6)

[Tree] (config>service>ies>if>sap>ipsec-gw>lcl-addr-assign ipv6)

Full Context

configure service vprn interface sap ipsec-gw local-address-assignment ipv6

configure service ies interface sap ipsec-gw local-address-assignment ipv6

Description

Commands in this context configure IPv6 local address assignment parameters for the IPsec gateway.

Platforms

7705 SAR Gen 2

ipv6

Syntax

ipv6

Context

[Tree] (config>router>pim>rp ipv6)

Full Context

configure router pim rp ipv6

Description

Commands in this context configure IPv6 parameters.

Platforms

7705 SAR Gen 2

ipv6

Syntax

[**no**] **ipv6** *ipv6-filter-id*

Context

[Tree] (config>filter>system-filter ipv6)

Full Context

configure filter system-filter ipv6

Description

This command activates an IPv6 system filter policy. Once activated, all IPv6 ACL filter policies that chain to the system filter (**config>filter>ipv6-filter>chain-to-system-filter**) will automatically execute system filter policy rules first.

The **no** form of the command deactivates the system filter policy.

Parameters

ipv6-filter-id

Specifies the existing IPv6 filter policy with scope **system**. This parameter can either be expressed as a decimal integer, or as an ASCII string of up to 64 characters in length.

Values 1 to 65535 or the filter policy name

Platforms

7705 SAR Gen 2

ipv6

Syntax

ipv6

Context

[Tree] (config>router ipv6)

Full Context

configure router ipv6

Description

Commands in this context configure the IPv6 interface of the router.

Default

ipv6

Platforms

7705 SAR Gen 2

ipv6

Syntax

[**no**] **ipv6**

Context

[Tree] (config>router>if ipv6)

Full Context

configure router interface ipv6

Description

This command configures IPv6 for a router interface.

The **no** form of this command disables IPv6 on the interface.

Default

no ipv6

Platforms

7705 SAR Gen 2

ipv6

Syntax

ipv6

Context

[Tree] (bof>autoconfigure ipv6)

Full Context

bof autoconfigure ipv6

Description

Commands in this context autoconfigure the IPv6 DHCP client.

Platforms

7705 SAR Gen 2

ipv6

Syntax

ipv6 *ipv6-address*

ipv6 auto-generate [**vendor-id-value** *vendor-id*]

no ipv6

Context

[Tree] (config>system>ned>prof>neip ipv6)

Full Context

configure system network-element-discovery profile neip ipv6

Description

This command configures the IPv6 NEIP for this profile. The NEIP can be configured manually or set to be automatically generated. If the NEIP is set to be automatically generated, the NEID is used for the subnet and host portion of the IPv6 address and the vendor ID value is set to 140 by default. The vendor ID value can be configured.

The **no** form of this command removes the IPv6 address association for this profile.

Default

no ipv6

Parameters

ipv6-address

Specifies the IPv6 address of the NEIP.

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0..FFFF]H
	d: [0..255]D

auto-generate

Specifies that the NEIP is automatically generated using the NEID.

vendor-id-value

Specifies the vendor ID value.

Values 1 to 255

Platforms

7705 SAR Gen 2

ipv6

Syntax

ipv6 send *send-limit* **receive** [**none**]

ipv6 send *send-limit*

no ipv6

Context

[Tree] (config>router>bgp>group>neighbor>add-paths ipv6)

[Tree] (config>router>bgp>group>add-paths ipv6)

[Tree] (config>router>bgp>add-paths ipv6)

Full Context

configure router bgp group neighbor add-paths ipv6

configure router bgp group add-paths ipv6

configure router bgp add-paths ipv6

Description

This command configures the add-paths capability for unlabeled IPv6 unicast routes. By default, add-paths is not enabled for unlabeled IPv6 unicast routes.

The maximum number of unlabeled unicast paths per IPv6 prefix to send is the configured send limit, which is a mandatory parameter. The capability to receive multiple unlabeled IPv6 unicast paths per prefix from a peer is configurable using the **receive** keyword, which is optional. If the **receive** keyword is not included in the command the receive capability is enabled by default.

The **no** form of this command disables add-paths support for unlabeled IPv6 unicast routes, causing sessions established using add-paths for unlabeled IPv6 unicast to go down and come back up without the add-paths capability.

Default

no ipv6

Parameters

send *send-limit*

Specifies the maximum number of paths per unlabeled IPv6 unicast prefix that are allowed to be advertised to add-paths peers. (The actual number of advertised routes may be less.) If the value is **none**, the router does not negotiate the send capability with respect to IPv6 AFI/SAFI. If the value is **multipaths**, then BGP advertises all the used BGP multipaths for each IPv6 NLRI if the peer has signaled support to receive multiple add-paths.

Values 1 to 16, none, multipaths

receive

Specifies the router negotiates to receive multiple unlabeled unicast routes per IPv6 prefix.

none

Specifies the router does not negotiate to receive multiple unlabeled unicast routes per IPv6 prefix.

Platforms

7705 SAR Gen 2

ipv6

Syntax

ipv6 *max-paths* [**ebgp** *ebgp-max-paths*] [**ibgp** *ibgp-max-paths*] [**restrict** { **same-neighbor-as** | **exact-as-path**}] [**unequal-cost**]

no ipv6

Context

[Tree] (config>router>bgp>multi-path ipv6)

Full Context

configure router bgp multi-path ipv6

Description

This command sets ECMP multipath parameters that apply only to the (unlabeled) IPv6 unicast address family. These settings override the values set by the **maximum-paths** command.

When multipath is enabled, traffic to the destination is load-shared across a set of paths (BGP routes) that the BGP decision process considers equal to the best path. The actual distribution of traffic over the multiple paths may be equal or unequal (that is, based on weights derived from the Link Bandwidth Extended Community).

The **no** form of this command removes IPv6-specific overrides.

Default

no ipv6

Parameters

max-paths

Specifies the maximum number of multipaths per prefix/NLRI if *ebgp-max-paths* or *ibgp-max-paths* does not apply.

Values 1 to 64

ebgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an EBGP learned route.

Values 1 to 64

ibgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an IBGP learned route.

Values 1 to 64

restrict same-neighbor-as

Specifies that the non-best path must have the same neighbor AS in its AS path as the best path.

restrict exact-as-path

Specifies that the non-best path must have the same AS path as the best path.

unequal-cost

Instructs BGP to ignore differences in the next-hop cost only when determining eligible multipaths.

Platforms

7705 SAR Gen 2

ipv6

Syntax

[**no**] **ipv6**

Context

[Tree] (config>router>isis>traffic-engineering-options ipv6)

Full Context

configure router isis traffic-engineering-options ipv6

Description

This command enables the advertisement of IPv6 TE in the IS-IS instance. When this command is enabled, traffic engineering behavior with IPv6 TE links is enabled. This IS-IS instance automatically begins advertising the new RFC 6119 IPv6 and TE TLVs and sub-TLVs.

The **no** form of this command disables IPv6 TE in this ISIS instance.

Default

no ipv6

Platforms

7705 SAR Gen 2

ipv6

Syntax

ipv6

Context

[Tree] (config>test-oam>icmp ipv6)

Full Context

configure test-oam icmp ipv6

Description

Commands in this context configure IPv6 traceroute packet handling.

Platforms

7705 SAR Gen 2

13.152 ipv6-adjacency-sid

ipv6-adjacency-sid

Syntax

ipv6-adjacency-sid label *value*

no ipv6-adjacency-sid

Context

[Tree] (config>router>isis>interface ipv6-adjacency-sid)

Full Context

configure router isis interface ipv6-adjacency-sid

Description

This command allows a static value to be assigned to an IPv6 adjacency SID in IS-IS segment routing.

The **label** option specifies that the value is assigned to an MPLS label.

The **no** form of this command removes the adjacency SID.

Parameters

value

Specifies the adjacency SID label.

Values 18432 to 5248, 1048575 (FP4 or FP5 only)

Platforms

7705 SAR Gen 2

13.153 ipv6-criteria

ipv6-criteria

Syntax

[**no**] **ipv6-criteria**

Context

[Tree] (config>qos>sap-ingress ipv6-criteria)

[Tree] (config>qos>sap-egress ipv6-criteria)

Full Context

configure qos sap-ingress ipv6-criteria

configure qos sap-egress ipv6-criteria

Description

IPv6 criteria-based SAP egress or ingress policies are used to select the appropriate ingress or egress queue or policer and corresponding forwarding class and packet profile for matched traffic.

This command is used to enter the node to create or edit policy entries that specify IPv6 criteria such as IP quintuple lookup or DiffServ code point.

The OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason, entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. When ipv6-criteria entries are removed from a SAP ingress policy, the ipv6-criteria is removed from all services where that policy is applied.

Platforms

7705 SAR Gen 2

ipv6-criteria

Syntax

[**no**] **ipv6-criteria**

Context

[Tree] (config>qos>network>ingress ipv6-criteria)

[Tree] (config>qos>network>egress ipv6-criteria)

Full Context

configure qos network ingress ipv6-criteria

configure qos network egress ipv6-criteria

Description

IPv6 criteria-based network ingress and egress policies are used to select the appropriate ingress or egress queue or policer, and the corresponding forwarding class and packet profile for matched traffic. This command is used to enter the context to create or edit policy entries that specify IPv6 criteria such as IP quintuple lookup or DSCP.

The SR OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. Entries must be sequenced correctly from most to least explicit.

The ingress classification only applies to the outer IPv6 header of non-tunneled traffic.

Attempting to apply a network QoS policy containing an **ipv6-criteria** statement to any object except a network IP interface will result in an error.

The **no** form of this command deletes all entries specified under this node. When IP criteria entries are removed from a network policy, the IPv6 criteria are removed from all network interfaces to which that policy is applied.

Platforms

7705 SAR Gen 2

13.154 ipv6-eh

ipv6-eh

Syntax

ipv6-eh {**max** | **limited**}

no ipv6-eh

Context

[Tree] (config>system>ip ipv6-eh)

Full Context

configure system ip ipv6-eh

Description

This command defines the maximum number of IPv6 extension headers parsed in the line cards. The system parses up to six extension headers when **ipv6-eh max** is configured.

When the **ipv6-eh limited** command is configured, the system does not parse IPv6 extension headers and provides consistent ipv6-filter matches for the next-header value found in the IPv6 packet header. LAG and ECMP hashing of IPv6 packets with extension headers is limited to Layer 3 IP addresses. Layer 4 ports, TEID, and SPI values are not available for hashing. MLD snooping on Layer 2 services is also not supported in this mode.

The **no** form of this command reverts to the default value.

Default

ipv6-eh max

Parameters

max

Specifies that the maximum number of IPv6 extension headers is parsed in the line cards.

limited

Specifies that the system does not parse IPv6 extension headers and provides consistent ipv6-filter matches for the next-header value found in the IPv6 packet header.

Platforms

7705 SAR Gen 2

13.155 ipv6-error

ipv6-error

Syntax

[**no**] **ipv6-error**

Context

[Tree] (debug>router>ip>event ipv6-error)

Full Context

debug router ip event ipv6-error

Description

This command enables debugging for IPv6 error events.

The **no** form of this command disables debugging for IPv6 error events

Platforms

7705 SAR Gen 2

13.156 ipv6-exception

ipv6-exception

Syntax

ipv6-exception *exception*

no ipv6-exception

Context

[Tree] (config>service>vprn>if>ipsec ipv6-exception)

[Tree] (config>router>if>ipsec ipv6-exception)

[Tree] (config>service>ies>if>ipsec ipv6-exception)

Full Context

configure service vprn interface ipsec ipv6-exception

configure router interface ipsec ipv6-exception

configure service ies interface ipsec ipv6-exception

Description

This command configures the IPv6 filter exception for an IPsec-secured IPv6 interface. When an IPv6 filter exception is added, clear text packets that match the exception criteria in the IPv6 filter exception policy can ingress the interface, even when IPsec is enabled on that interface.

The **no** form of this command removes the IPv6 filter exception.

Default

no ipv6-exception

Parameters

exception

Specifies the IPv6 filter exception that is used to bypass encryption.

Values *exception-id*: 1 to 65535

exception-name: An existing IPv6 filter exception name up to 64 characters.

Platforms

7705 SAR Gen 2

ipv6-exception

Syntax

ipv6-exception *exception-id* [**name** *exception-name*] [**create**]

no ipv6-exception {*exception-id* | *exception-name*}

Context

[Tree] (config>filter ipv6-exception)

Full Context

configure filter ipv6-exception

Description

Commands in this context configure the specified IPv6 exception filter.

The **no** form of the command deletes the IPv6 exception filter.

Parameters

exception-id

Specifies the IPv6 filter exception ID expressed as a decimal integer.

Values 1 to 65535

name *exception-name*

Specifies the IPv6 filter exception as a name, up to 64 characters.

create

This keyword is required to create the configuration context. Once it is created, the context can be enabled with or without the **create** keyword.

Platforms

7705 SAR Gen 2

13.157 ipv6-filter

ipv6-filter

Syntax

ipv6-filter *ipv6-filter-id* **entry** *entry-id* [*entry-id*]

no ipv6-filter *ipv6-filter-id* [**entry** *entry-id*]

Context

[Tree] (config>mirror>mirror-source ipv6-filter)

Full Context

configure mirror mirror-source ipv6-filter

Description

This command enables mirroring of packets that match specific entries in an existing IPv6 filter.

The **ipv6-filter** command directs packets which match the defined list of entry IDs to be mirrored to the mirror destination referenced by the *mirror-dest-service-id* of the **mirror-source**.

The IPv6 filter must already exist in order for the command to execute. Filters are configured in the **config>filter** context. If the IPv6 filter does not exist, an error will occur. If the filter exists but has not been associated with a SAP or IPv6 interface, an error is not generated but mirroring will not be enabled (there are no packets to mirror). Once the IPv6 filter is defined to a SAP or IPv6 interface, mirroring is enabled.

If the IPv6 filter is defined as ingress, only ingress packets are mirrored. Ingress mirrored packets are mirrored to the mirror destination prior to any ingress packet modifications.

If the IPv6 filter is defined as egress, only egress packets are mirrored. Egress mirrored packets are mirrored to the mirror destination after all egress packet modifications.

An *entry-id* within an IPv6 filter can only be mirrored to a single mirror destination. If the same *entry-id* is defined multiple times, an error occurs and only the first **mirror-source** definition is in effect.

By default, no packets matching any IPv6 filters are mirrored. Mirroring of IPv6 filter entries must be explicitly defined.

The **no ipv6-filter** command, without the **entry** keyword, removes mirroring on all *entry-id*'s within the *ip-filter-id*.

When the **no** form of this command is executed with the **entry** keyword and one or more *entry-id*'s, mirroring that *entry-id* list is terminated within the *ip-filter-id*. If an *entry-id* is listed that does not exist, an error will occur and the command will not execute. If an *entry-id* is listed that is not currently being mirrored, no error will occur for that *entry-id* and the command will execute normally.

Parameters

ipv6-filter-id

Specifies the IPv6 filter ID whose entries are mirrored. If the *ipv6-filter-id* does not exist, an error will occur and the command will not execute. Mirroring of packets will commence once the *ipv6-filter-id* is defined on a SAP or IPv6 interface.

entry-id

Specifies the IP filter entries to use as match criteria for packet mirroring. The **entry** keyword begins a list of *entry-id*'s for mirroring. Multiple *entry-id* entries may be specified with a single command. Each *entry-id* must be separated by a space.

If an *entry-id* does not exist within the IP filter, an error occurs and the command will not execute.

If the filter's *entry-id* is renumbered within the IP filter definition, the old *entry-id* is removed but the new *entry-id* must be manually added to the configuration to include the new (renumbered) entry's criteria.

Platforms

7705 SAR Gen 2

ipv6-filter

Syntax

ipv6-filter *filter-id* [**name** *filter-name*] [**create**]

no ipv6-filter {*filter-id* | *filter-name*}

Context

[Tree] (config>filter ipv6-filter)

Full Context

configure filter ipv6-filter

Description

Commands in this context configure the specified IPv6 filter policy.

The **no** form of the command deletes the IPv6 filter policy. A filter policy cannot be deleted until it is removed from all objects where it is applied.

Parameters

filter-id

Specifies the IPv6 filter policy ID expressed as a decimal integer.

Values 1 to 65535

name

Configures an optional filter name, up to 64 characters in length, to a given filter. This filter name can then be used in configuration references, display, and show commands throughout the system. A defined filter name can help the service provider or administrator to identify and manage filters within the SR OS platforms.

To create a filter, you must assign a filter ID, however, after it is created, either the filter ID or filter name can be used to identify and reference a filter.

If a name is not specified at creation time, then SR OS assigns a string version of the *filter-id* as the name.

Filter names may not begin with an integer (0 to 9).

Values *name*: 64 characters maximum

filter-name

Specifies a string of up to 64 characters uniquely identifying this IPv6 filter policy.

create

This keyword is required to create the configuration context. Once it is created, the context can be enabled with or without the **create** keyword.

Platforms

7705 SAR Gen 2

ipv6-filter

Syntax

[**no**] **ipv6-filter**

Context

[Tree] (config>system>security>mgmt-access-filter ipv6-filter)

Full Context

configure system security management-access-filter ipv6-filter

Description

Commands in this context configure management access IPv6 filter parameters.

Platforms

7705 SAR Gen 2

ipv6-filter

Syntax

ipv6-filter *src-filter-id* [**src-entry** *src-entry-id*] **to** *dst-filter-id* [**dst-entry** *dst-entry-id*] [**overwrite**]

Context

[Tree] (config>filter>copy ipv6-filter)

Full Context

configure filter copy ipv6-filter

Description

This command copies an existing filter entry for a specific filter ID to another filter ID. The command is a configuration level maintenance tool used to create new entries using an existing filter policy. If **overwrite** is not specified, an error will occur if the destination filter entry exists.

Parameters

src-filter-id

Identifies the source filter policy from which the copy command will attempt to copy. The filter policy must exist within the context of the preceding keyword (**ipv6-filter**).

dst-filter-id

Identifies the destination filter policy to which the copy command will attempt to copy. If the **overwrite** keyword is not specified, the filter entry ID cannot already exist in the destination filter policy. If the **overwrite** keyword is present, the destination entry ID may or may not exist.

overwrite

Specifies that the destination filter entry may exist. If it does, everything in the existing destination filter entry will be completely overwritten with the contents of the source filter entry. If the destination filter entry exists, either **overwrite** must be specified or an error message will be returned. If **overwrite** is specified, the function of copying from source to destination occurs in a "break before make" manner and therefore should be handled with care.

Platforms

7705 SAR Gen 2

13.158 ipv6-multicast

ipv6-multicast

Syntax

[**no**] **ipv6-multicast**

Context

[Tree] (config>router>isis>multi-topology ipv6-multicast)

Full Context

configure router isis multi-topology ipv6-multicast

Description

This command enables support for the IPv6 topology (MT4) within the associate IS-IS instance.

The **no** form of this command disables support for the IPv6 topology (MT4) within the associated IS-IS instance.

Default

no ipv6-multicast

Platforms

7705 SAR Gen 2

13.159 ipv6-multicast-disable

ipv6-multicast-disable

Syntax

ipv6-multicast-disable

Context

[Tree] (config>service>vprn>pim>if ipv6-multicast-disable)

[Tree] (config>service>vprn>pim ipv6-multicast-disable)

Full Context

configure service vprn pim interface ipv6-multicast-disable

configure service vprn pim ipv6-multicast-disable

Description

This command administratively disables/enables PIM operation for IPv6.

Default

ipv6-multicast-disable (config>service>vprn>pim)

no ipv6-multicast-disable (config>service>vprn>pim>if)

Platforms

7705 SAR Gen 2

ipv6-multicast-disable

Syntax

[**no**] **ipv6-multicast-disable**

Context

[Tree] (config>router>pim ipv6-multicast-disable)

[Tree] (config>router>pim>interface ipv6-multicast-disable)

Full Context

configure router pim ipv6-multicast-disable

configure router pim interface ipv6-multicast-disable

Description

This command administratively enables PIM operation for IPv6.

IPv6 multicast must be enabled to enable MLDP in-band signaling for IPv6 PIM joins; see **config>router>pim>interface p2mp-ldp-tree-join**.

The **no** form of this command disables the PIM operation for IPv6.

Default

ipv6-multicast-disable

Platforms

7705 SAR Gen 2

ipv6-multicast-disable

Syntax

[**no**] **ipv6-multicast-disable**

Context

[Tree] (config>router>isis>interface ipv6-multicast-disable)

Full Context

configure router isis interface ipv6-multicast-disable

Description

This command disables IS-IS IPv6 multicast routing for the interface.

The **no** form of this command enables IS-IS IPv6 multicast routing for the interface.

Platforms

7705 SAR Gen 2

13.160 ipv6-multicast-metric

ipv6-multicast-metric

Syntax

ipv6-multicast-metric *metric*

no ipv6-multicast-metric

Context

[Tree] (config>router>isis>if>level ipv6-multicast-metric)

Full Context

configure router isis interface level ipv6-multicast-metric

Description

This command configures the IS-IS interface metric for IPv6 multicast.

The **no** form of this command removes the metric from the configuration.

Default

no ipv6-multicast-metric

Parameters

metric

Specifies the IS-IS interface metric for IPv6 multicast.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

13.161 ipv6-multicast-metric-offset

ipv6-multicast-metric-offset

Syntax

ipv6-multicast-metric-offset *offset-value*

no ipv6-multicast-metric-offset

Context

[Tree] (config>router>isis>link-group>level ipv6-multicast-metric-offset)

Full Context

configure router isis link-group level ipv6-multicast-metric-offset

Description

This command sets the offset value for the IPv6 multicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric for the IPv6 multicast topology.

The **no** form of this command reverts the offset value to 0.

Default

no ipv6-multicast-metric-offset

Parameters

offset-value

Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold

Values 0 to 6777215

Platforms

7705 SAR Gen 2

13.162 ipv6-multicast-routing

ipv6-multicast-routing

Syntax

ipv6-multicast-routing {**native** | **mt**}

[**no**] **ipv6-multicast-routing**

Context

[Tree] (config>router>isis ipv6-multicast-routing)

Full Context

configure router isis ipv6-multicast-routing

Description

The multicast RTM is used for Reverse Path Forwarding checks. This command controls which IS-IS topology is used to populate the IPv6 multicast RTM.

The **no** form of this command results in none of the IS-IS routes being populated in the IPv4 multicast RTM and would be used if multicast is configured to use the unicast RTM for the RPF check.

Default

ipv6-multicast-routing native

Parameters

native

Causes IPv6 routes from the MT0 topology to be added to the multicast RTM for RPF checks.

mt

Causes IPv6 routes from the MT3 topology to be added to the multicast RTM for RPF checks.

Platforms

7705 SAR Gen 2

13.163 ipv6-node-sid

ipv6-node-sid

Syntax

ipv6-node-sid index *index-value* [**clear-n-flag**]

ipv6-node-sid label *label-value* [**clear-n-flag**]

no ipv6-node-sid

Context

[Tree] (config>router>isis>interface ipv6-node-sid)

Full Context

configure router isis interface ipv6-node-sid

Description

This command assigns a node SID index or label value to the prefix representing the primary address of an IPv6 network interface of type loopback. Only a single node SID can be assigned to an IPv6 interface. When an IPv6 interface has multiple global addresses, the primary address is always the first one in the list, as displayed by the **interface info** command.

The command fails if the network interface is not of loopback type or if the interface is defined in an IES or a VPRN context. Assigning the same SID index/label value to the same interface in two different IGP instances is not allowed within the same node.

The value of the label or index SID is taken from the range configured for this IGP instance. When using the global mode of operation, a new segment routing module checks that the same index or label value cannot be assigned to more than one loopback interface address. When using the per-instance mode of operation, this check is not required since the index and thus label ranges of the various IGP instance are not allowed to overlap.

The **clear-n-flag** option allows the user to clear the N-flag (node-sid flag) in an IS-IS prefix SID sub-TLV originated for the IPv6 prefix of a loopback interface on the system.

By default, the prefix SID sub-TLV for the prefix of a loopback interface is tagged as a node SID, meaning that it belongs to this node only. However, when the user wants to configure and advertise an anycast SID using the same loopback interface prefix on multiple nodes, you must clear the N-flag to assure interoperability with third-party implementations, which may perform a strict check on the receiving end and drop duplicate prefix SID sub-TLVs when the N-flag is set.

The SR OS implementation is relaxed on the receiving end and accepts duplicate prefix SIDs with the N-flag set or cleared. SR OS will resolve to the closest owner, or owners if ECMP is configured, of the prefix SID according to its cost.

Default

no ipv6-node-sid

Parameters

index-value

Specifies the index value.

Values 0 to 4294967295

label-value

Specifies the label value.

Values 0 to 4294967295

clear-n-flag

Clears the node SID flag.

Default no clear-n-flag

Platforms

7705 SAR Gen 2

13.164 ipv6-prefix

ipv6-prefix

Syntax

[**no**] **ipv6-prefix**

Context

[Tree] (debug>router>rpki-session>packet ipv6-prefix)

Full Context

debug router rpki-session packet ipv6-prefix

Description

This command enables debugging for IPv6 prefix RPKI packets.

The **no** form of this command disables debugging for IPv6 prefix RPKI packets.

Platforms

7705 SAR Gen 2

13.165 ipv6-prefix-list

ipv6-prefix-list

Syntax

ipv6-prefix-list *ipv6-prefix-list-name* [**create**]

no ipv6-prefix-list *ipv6-prefix-list-name*

Context

[Tree] (config>qos>match-list ipv6-prefix-list)

Full Context

configure qos match-list ipv6-prefix-list

Description

This command creates a list of IPv6 prefixes for match criteria in QoS policies. An ipv6-prefix-list must contain only IPv6 address prefixes created using the **prefix** command and cannot be deleted if it is referenced by a QoS policy.

The **no** form of this command deletes the specified list.

Parameters

ipv6-prefix-list-name

A string of up to 32 characters of printable ASCII characters. If special characters are used (#, ?, space), the string must be enclosed within double quotes. The name **default** (case insensitive) is reserved by the system.

create

Creates IPv6 prefixes for match criteria in QoS policies.

Platforms

7705 SAR Gen 2

ipv6-prefix-list

Syntax

ipv6-prefix-list *ipv6-prefix-list-name* [**create**]

no ipv6-prefix-list *ipv6-prefix-list-name*

Context

[Tree] (config>filter>match-list ipv6-prefix-list)

Full Context

configure filter match-list ipv6-prefix-list

Description

This command creates a list of IPv6 prefixes for match criteria in ACL and CPM IPv6 filter policies.

The **no** form of this command deletes the specified list.

Operational Notes:

An IPv6 prefix list must contain only IPv6 address prefixes.

An IPv6 prefix list cannot be deleted if it is referenced by a filter policy.

See general description related to match-list usage in filter policies.

Parameters

ipv6-prefix-list-name

Specifies a string of up to 32 printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

13.166 ipv6-routing

ipv6-routing

Syntax

[**no**] **ipv6-routing** {**native** | **mt**}

Context

[Tree] (config>service>vprn>isis ipv6-routing)

Full Context

configure service vprn isis ipv6-routing

Description

This command enables IPv6 routing.

The **no** form of this command disables support for IS-IS IPv6 TLVs for IPv6 routing.

Default

no ipv6-routing

Parameters

native

Enables IS-IS IPv6 TLVs for IPv6 routing and enables support for native IPv6 TLVs.

mt

Enables IS-IS multi-topology TLVs for IPv6 routing. When this parameter is specified, the support for native IPv6 TLVs is disabled.

Platforms

7705 SAR Gen 2

ipv6-routing

Syntax

[**no**] **ipv6-routing** {**native** | **mt**}

Context

[Tree] (config>router>isis ipv6-routing)

Full Context

configure router isis ipv6-routing

Description

This command enables IPv6 routing.

The **no** form of this command disables support for IS-IS IPv6 TLVs for IPv6 routing.

Default

no ipv6-routing

Parameters

native

Enables IS-IS IPv6 TLVs for IPv6 routing and enables support for native IPv6 TLVs.

mt

Enables IS-IS multi-topology TLVs for IPv6 routing. When this parameter is specified, the support for native IPv6 TLVs is disabled.

Platforms

7705 SAR Gen 2

13.167 ipv6-sid

ipv6-sid

Syntax

ipv6-sid index *index-id*

ipv6-sid label *label-id*

no ipv6-sid

Context

[Tree] (config>router>segment-routing>sr-mpls>prefix-sids ipv6-sid)

Full Context

configure router segment-routing sr-mpls prefix-sids ipv6-sid

Description

This command is used to configure the IPv6 segment routing SID associated with the primary IPv6 address of the loopback or system interface.

The **no** form of this command removes the configuration of the IPv6 segment routing SID associated with the primary IPv6 interface address.

Default

no ipv6-sid

Parameters

index *index-id*

Specifies the node SID index for this interface.

Values 0 to 4294967295

label *label-id*

Specifies the label value for the node SID.

Values 32 to 1048575

Platforms

7705 SAR Gen 2

13.168 ipv6-source-address

ipv6-source-address

Syntax

ipv6-source-address *ipv6-address*

no ipv6-source-address

Context

[Tree] (config>aaa>radius-srv-plcy>servers ipv6-source-address)

Full Context

configure aaa radius-server-policy servers ipv6-source-address

Description

This command configures the source address of an IPv6 RADIUS packet.

When no ipv6-source-address is configured, the system IPv6 address (inband RADIUS server connection) or Boot Option File (BOF) IPv6 address (outband RADIUS server connection) must be configured in order for the RADIUS client to work with an IPv6 RADIUS server.

This address is also used in the NAS-IPv6-Address attribute.

The **no** form of this command reverts to the default value.

Parameters

ipv6-address

Specifies the source address of an IPv6 RADIUS packet.

Platforms

7705 SAR Gen 2

ipv6-source-address

Syntax

ipv6-source-address *ipv6-address*

no ipv6-source-address

Context

[Tree] (config>service>vprn>dns ipv6-source-address)

Full Context

configure service vprn dns ipv6-source-address

Description

This command configures the IPv6 address of the default secondary DNS server for the subscribers using this interface. Subscribers that cannot obtain an IPv6 DNS server address by other means, can use this for DNS name resolution.

The ipv6-address value can only be set to a nonzero value if the value of VPRN type is set to **subscriber-split-horizon**.

The **no** form of this command reverts to the default.

Parameters

ipv6-address

Specifies the IPv6 address of the default secondary DNS server.

Values ipv6-address - a.b.c.d

Platforms

7705 SAR Gen 2

ipv6-source-address

Syntax

ipv6-source-address *ipv6-address*

no ipv6-source-address

Context

[Tree] (config>system>file-trans-prof ipv6-source-address)

Full Context

configure system file-transmission-profile ipv6-source-address

Description

This command specifies the IPv6 source address used for transport protocol.

The **no** form of this command uses the default source address which typically is the address of egress interface.

Default

no ipv6-source-address

Parameters

Ipv6-address

Specifies a unicast v6 address. This should be a local interface address.

Platforms

7705 SAR Gen 2

13.169 ipv6-te-router-id

ipv6-te-router-id

Syntax

ipv6-te-router-id interface *interface-name*

no ipv6-te-router-id

Context

[Tree] (config>router ipv6-te-router-id)

Full Context

configure router ipv6-te-router-id

Description

This command configures the IPv6 TE Router ID. The IPv6 TE Router ID, when configured, uniquely identifies the router as being IPv6 TE capable to other routers in an IGP TE domain.

IS-IS advertises this information using the IPv6 TE Router ID TLV.

If this command is not configured, the IPv6 TE Router ID will use the global unicast address of the system interface by default. The user can specify the system interface using this command to achieve the same result. If a different interface is specified, the preferred primary global unicast address of that interface is used instead

The **no** form of this command reverts the IPv6 TE Router ID to the default value.

Parameters

interface interface-name

Specifies the name of the interface to be added or removed. Only system and loopback interfaces are accepted.

Platforms

7705 SAR Gen 2

13.170 ipv6-unicast

ipv6-unicast

Syntax

[**no**] **ipv6-unicast**

Context

[Tree] (config>service>vprn>isis>multi-topology ipv6-unicast)

Full Context

configure service vprn isis multi-topology ipv6-unicast

Description

This command enables multi-topology TLVs.

The **no** form of this command disables multi-topology TLVs.

Platforms

7705 SAR Gen 2

ipv6-unicast

Syntax

[**no**] **ipv6-unicast**

Context

[Tree] (config>router>isis>multi-topology ipv6-unicast)

Full Context

configure router isis multi-topology ipv6-unicast

Description

This command enables multi-topology TLVs.

The **no** form of this command disables multi-topology TLVs.

Default

no ipv6-unicast

Platforms

7705 SAR Gen 2

13.171 ipv6-unicast-disable

ipv6-unicast-disable

Syntax

[**no**] **ipv6-unicast-disable**

Context

[Tree] (config>service>vprn>isis>if ipv6-unicast-disable)

[Tree] (config>router>isis>if ipv6-unicast-disable)

Full Context

configure service vprn isis interface ipv6-unicast-disable

configure router isis interface ipv6-unicast-disable

Description

This command disables IS-IS IPv6 unicast routing for the interface.

By default IPv6 unicast on all interfaces is enabled. However, IPv6 unicast routing on IS-IS is in effect when the **config>router>isis>ipv6-routing mt** command is configured.

The **no** form of this command enables IS-IS IPv6 unicast routing for the interface.

Platforms

7705 SAR Gen 2

13.172 ipv6-unicast-metric

ipv6-unicast-metric

Syntax

ipv6-unicast-metric *metric*

no ipv6-unicast-metric

Context

[Tree] (config>service>vprn>isis>if>level ipv6-unicast-metric)

Full Context

configure service vprn isis interface level ipv6-unicast-metric

Description

This command configures IS-IS interface metric for IPv6 unicast.

The **no** form of this command removes the metric from the configuration.

Parameters

metric

Specifies the IS-IS interface metric for IPv6 unicast.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

ipv6-unicast-metric

Syntax

ipv6-unicast-metric *metric*

no ipv6-unicast-metric

Context

[Tree] (config>router>isis>if>level ipv6-unicast-metric)

Full Context

configure router isis interface level ipv6-unicast-metric

Description

This command configures the IS-IS interface metric for IPv6 unicast.

The **no** form of this command removes the metric from the configuration.

Default

no ipv6-unicast-metric

Parameters

metric

Specifies the IS-IS interface metric for IPv6 unicast.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

13.173 ipv6-unicast-metric-offset

ipv6-unicast-metric-offset

Syntax

ipv6-unicast-metric-offset *offset-value*

no ipv6-unicast-metric-offset

Context

[Tree] (config>service>vprn>isis>link-group>level ipv6-unicast-metric-offset)

Full Context

configure service vprn isis link-group level ipv6-unicast-metric-offset

Description

This command sets the offset value for the IPv6 unicast address family. If the number of operational links drops below the **oper-members** threshold, the configured offset is applied to the interface metric for the IPv6 topology.

The **no** form of this command reverts the offset value to 0.

Default

no ipv6-unicast-metric-offset

Parameters

offset-value

Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold.

Values 0 to 6777215

Platforms

7705 SAR Gen 2

ipv6-unicast-metric-offset

Syntax

ipv6-unicast-metric-offset *offset-value*

no ipv6-unicast-metric-offset

Context

[Tree] (config>router>isis>link-group>level ipv6-unicast-metric-offset)

Full Context

configure router isis link-group level ipv6-unicast-metric-offset

Description

This command sets the offset value for the IPv6 unicast address family. If the number of operational links drops below the **oper-members** threshold, the configured offset is applied to the interface metric for the IPv6 topology.

The **no** form of this command reverts the offset value to 0.

Default

no ipv6-unicast-metric-offset

Parameters

offset-value

Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold.

Values 0 to 6777215

Platforms

7705 SAR Gen 2

13.174 isa

isa

Syntax

isa

Context

[Tree] (config isa)

Full Context

configure isa

Description

Commands in this context configure Integrated Services Adapter (ISA) parameters.

Platforms
7705 SAR Gen 2

13.175 isa-dp-cpu-usage

isa-dp-cpu-usage

Syntax
[**no**] **isa-dp-cpu-usage**

Context
[Tree] (config>isa>tunnel-grp>stats-collection isa-dp-cpu-usage)

Full Context
configure isa tunnel-group stats-collection isa-dp-cpu-usage

Description
This command enables the system to collect statistics used to derive ISA CPU data plane usage. When enabled, this command impacts the ISA performance.

Platforms
7705 SAR Gen 2

13.176 isakmp-lifetime

isakmp-lifetime

Syntax
isakmp-lifetime *seconds*

Context
[Tree] (config>ipsec>ike-transform isakmp-lifetime)

Full Context
configure ipsec ike-transform isakmp-lifetime

Description
This command specifies the lifetime of the IKE SA.

Default

isakmp-lifetime 86400

Parameters

seconds

Specifies the Phase 1 life time for this IKE transform.

Values 1200 to 31536000

Platforms

7705 SAR Gen 2

13.177 isis

isis

Syntax

[**no**] **isis** *isis-instance*

Context

[Tree] (config>service>vprn isis)

Full Context

configure service vprn isis

Description

Commands in this context configure the Intermediate-System-to-Intermediate-System (IS-IS) protocol instance in the VPRN.

The IS-IS protocol instance is enabled with the **no shutdown** command in the **config>service>vprn>isis** context. Alternatively, the IS-IS protocol instance is disabled with the **shutdown** command in the **config>service>vprn>isis** context.

IS-IS instances are **shutdown** when created, so that all parameters can be configured prior to the instance being enabled.

The **no** form of this command disables the ISIS protocol instance from the given VPRN service.

Default

0

Parameters

isis-instance

Specifies the instance ID for an IS-IS instance.

Values 0 to 127

Platforms

7705 SAR Gen 2

isis

Syntax

[**no**] **isis** [*isis-instance*]

Context

[Tree] (config>router isis)

Full Context

configure router isis

Description

Commands in this context configure the Intermediate-System-to-Intermediate-System (IS-IS) protocol instance.

The IS-IS protocol instance is enabled with the **no shutdown** command in the **config>router>isis** context. Alternatively, the IS-IS protocol instance is disabled with the **shutdown** command in the **config>router>isis** context.

IS-IS instances are shutdown when created, so that all parameters can be configured prior to the instance being enabled.

The **no** form of this command deletes the IS-IS protocol instance. Deleting the protocol instance removes all configuration parameters for this IS-IS instance.

Parameters

isis-instance

Specifies the instance ID for an IS-IS instance.

Values 0 to 127

Platforms

7705 SAR Gen 2

isis

Syntax

isis [*isis-instance*]

Context

[Tree] (debug>router isis)

Full Context

debug router isis

Description

Commands in this context debug IS-IS protocol entities.

Parameters

isis-instance

Specifies the IS-IS instance.

Values 0 to 127

Platforms

7705 SAR Gen 2

Output

The following output is an example of the debugging information.

Output Example

```
*A:Dut-C# /tools dump router isis sr-database  prefix 10.20.1.5 detail
===============================================================================
Rtr Base ISIS Instance 0 SR Database
===============================================================================
103   474390 10.20.1.5       LfaNhops  1   0  15      1000    1    1
  1492   1500    1500   0 0  1  1   0100.2000.1005 SR_ERR_OK
        IP:10.10.5.5 gifId:3 ifId:4 protectId:7 numLabels:1 outLbl:474390 isAdv:1 is
LfaX:0
        IP:10.10.12.2 gifId:5 ifId:6 protectId:0 numLabels:2 outLbl1:474389 outLbl2:
474390 numLfaNhops:1 isAdv:0
-------------------------------------------------------------------------------
D = duplicate pending
xL = exclude from LFA
rL = remote LFA
Act = tunnel active
LDP = LDP FEC is the SID NH for SR-LDP stitching
===============================================================================
```

```
*A:Dut-C# /tools dump router isis sr-database nh-type ldp  detail
===============================================================================
Rtr Base ISIS Instance 0 SR Database
===============================================================================
SID   Label  Prefix          Last-act  Lev MT TnlPref Metric  IpNh SrNh
  Mtu    MtuPrim MtuBk  D xL rL Act AdvSystemId    SrErr
-------------------------------------------------------------------------------
1000  475287 10.20.1.4       AddTnl    1   0  15      0       1    1
  0      0       0      0 0  0  1   0100.2000.1004 SR_ERR_OK
        LDP: IP:10.20.1.4 tnlId:65546 tnlTyp:2
1001  475288 10.20.1.5       AddTnl    1   0  15      0       1    1
  0      0       0      0 0  0  1   0100.2000.1005 SR_ERR_OK
        LDP: IP:10.20.1.5 tnlId:65548 tnlTyp:2
```

```
1002  475289 10.20.1.6       AddTnl    1   0  15      0       1    1
  0      0       0      0 0  0  1   0100.2000.1006 SR_ERR_OK
         LDP: IP:10.20.1.6 tnlId:65549 tnlTyp:2
-------------------------------------------------------------------------------
D = duplicate pending
xL = exclude from LFA
rL = remote LFA
Act = tunnel active
LDP = LDP FEC is the SID NH for SR-LDP stitching
===============================================================================
```

14 j Commands

14.1 jitter-event

jitter-event

Syntax

jitter-event rising-threshold *threshold* [**falling-threshold** *threshold*] [*direction*]

no jitter-event

Context

[Tree] (config>saa>test jitter-event)

Full Context

configure saa test jitter-event

Description

This command specifies that at the termination of an SAA test probe, the calculated jitter value is evaluated against the configured rising and falling jitter thresholds. SAA threshold events are generated as required.

Once the threshold (rising/falling) is crossed, it is disabled from generating additional events until the opposite threshold is crossed. If a **falling-threshold** is not supplied, the **rising- threshold** is re-enabled when it falls below the threshold after the initial crossing that generated the event.

The configuration of jitter event thresholds is optional.

The **no** form of the command disables the jitter event.

Parameters

rising-threshold *threshold*

Specifies a rising threshold jitter value, in milliseconds. When the test run is completed, the calculated jitter value is compared to the configured jitter rising threshold. If the test run jitter value is greater than the configured rising threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Values 0 to 2147483

Default 0

falling-threshold *threshold*

Specifies a falling threshold jitter value, in milliseconds. When the test run is completed, the calculated jitter value is compared to the configured jitter falling threshold. If the test

run jitter value is greater than the configured falling threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Values 0 to 2147483

Default 0

direction

Specifies the direction for OAM ping responses received for an OAM ping test run.

Values **inbound** — Monitor the value of jitter calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run.

outbound — Monitor the value of jitter calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run.

roundtrip — Monitor the value of jitter calculated for the round trip, two-way, OAM ping requests and replies for an OAM ping test run.

Default roundtrip

Platforms

7705 SAR Gen 2

14.2 jp

jp

Syntax

jp [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]

no jp

Context

[Tree] (debug>router>pim jp)

Full Context

debug router pim jp

Description

This command enables debugging for PIM join and prune mechanisms.

The **no** form of this command disables PIM join and prune mechanisms debugging.

Parameters

grp-ip-address

Debugs information associated with the specified Join-Prune mechanism.

Values multicast group address (ipv4, ipv6) or zero

ip-address

Debugs information associated with the specified Join-Prune mechanism.

Values source address (ipv4, ipv6)

detail

Debugs detailed Join-Prune mechanism information.

Platforms

7705 SAR Gen 2

15 k Commands

15.1 kb-memory-use-alarm

kb-memory-use-alarm

Syntax

kb-memory-use-alarm rising-threshold *threshold* [**falling-threshold** *threshold*] **interval** *seconds* [*rmon-event-type*] [**startup-alarm** *alarm-type*]

no kb-memory-use-warn

Context

[Tree] (config>system>thresholds kb-memory-use-alarm)

Full Context

configure system thresholds kb-memory-use-alarm

Description

This command configures memory use, in kilobytes, alarm thresholds.

The **no** form of the command removes the parameters from the configuration.

Parameters

rising-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is greater than or equal to this threshold and the associated startup-alarm is equal to rising or either.

After a rising threshold crossing event is generated, another such event will not be generated until the sampled value falls below this threshold and reaches less than or equal the falling-threshold value.

The threshold value represents units of kilobytes.

Values -2147483648 to 2147483647

Default 0

falling-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater

than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is less than or equal to this threshold and the associated startup-alarm is equal to falling or either.

After a falling threshold crossing event is generated, another such event will not be generated until the sampled value rises above this threshold and reaches greater than or equal the rising-threshold value.

The threshold value represents units of kilobytes.

Values -2147483648 to 2147483647

Default 0

seconds

Specifies the polling period over which the data is sampled and compared with the rising and falling thresholds.

Values 1 to 2147483647

rmon-event-type

Specifies the type of notification action to be taken when this event occurs.

Values log — In the case of log, an entry is made in the RMON-MIB log table for each event occurrence. This does not create an SR OS logger entry. The RMON-MIB log table entries can be viewed using the **show>system>thresholds** CLI command.

trap — An SR OS logger event is generated. The SR OS logger utility then distributes the notification of this event to its configured log destinations which may be CONSOLE, telnet session, memory log, cflash file, syslog, or SNMP trap destinations logs.

both — Both an entry in the RMON-MIB logTable and an SR OS logger event are generated.

none — No action is taken.

Default both

startup-alarm *alarm-type*

Specifies the alarm that may be sent when this alarm is first created. If the first sample is greater than or equal to the rising threshold value and startup-alarm is equal to rising or either, then a single rising threshold crossing event is generated. If the first sample is less than or equal to the falling threshold value and startup-alarm is equal to falling or either, a single falling threshold crossing event is generated.

Values rising, falling, either

Default either

Platforms

7705 SAR Gen 2

15.2 kb-memory-use-warn

kb-memory-use-warn

Syntax

kb-memory-use-warn rising-threshold *threshold* [**falling-threshold** *threshold*] **interval** *seconds* [*rmon-event-type*] [**startup-alarm** *alarm-type*]

no kb-memory-use-warn

Context

[Tree] (config>system>thresholds kb-memory-use-warn)

Full Context

configure system thresholds kb-memory-use-warn

Description

This command configures memory usage, in kilobytes, for warning thresholds

Parameters

rising-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is greater than or equal to this threshold and the associated startup-alarm is equal to rising or either.

After a rising threshold crossing event is generated, another such event will not be generated until the sampled value falls below this threshold and reaches less than or equal the falling-threshold value.

The threshold value represents units of kilobytes.

Values -2147483648 to 2147483647

Default 0

falling-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is less than or equal to this threshold and the associated startup-alarm is equal to falling or either.

After a falling threshold crossing event is generated, another such event will not be generated until the sampled value rises above this threshold and reaches greater than or equal the rising-threshold value.

The threshold value represents units of kilobytes.

Values -2147483648 to 2147483647

Default 0

seconds

Specifies the polling period over which the data is sampled and compared with the rising and falling thresholds.

Values 1 to 2147483647

rmon-event-type

Specifies the type of notification action to be taken when this event occurs.

Values log — An entry is made in the RMON-MIB log table for each event occurrence. This does not create an SR OS logger entry. The RMON-MIB log table entries can be viewed using the show>system>thresholds CLI command.

trap — An SR OSS logger event is generated. The SR OS logger utility then distributes the notification of this event to its configured log destinations which may be CONSOLE, telnet session, memory log, cflash file, syslog, or SNMP trap destinations logs.

both — Both an entry in the RMON-MIB logTable and an SR OS logger event are generated.

none — No action is taken.

Default both

alarm-type

Specifies the alarm that may be sent when this alarm is first created. If the first sample is greater than or equal to the rising threshold value and startup-alarm is equal to rising or either, then a single rising threshold crossing event is generated. If the first sample is less than or equal to the falling threshold value and startup-alarm is equal to falling or either, a single falling threshold crossing event is generated.

Values rising, falling, either

Default either

Platforms

7705 SAR Gen 2

15.3 keep-alive

keep-alive

Syntax

keep-alive

Context

[Tree] (config>service>sdp keep-alive)

Full Context

configure service sdp keep-alive

Description

This command enables the context to configure SDP connectivity monitoring keepalive messages for the SDP ID.

SDP ID keepalive messages use SDP Echo Request and Reply messages to monitor SDP connectivity. The operating state of the SDP is affected by the keepalive state on the SDP ID. SDP Echo Request messages are only sent when the SDP ID is completely configured and administratively up. If the SDP ID is administratively down, keepalives for that SDP ID are disabled. SDP Echo Requests (when sent for keepalive messages) are always sent with the *originator-sdp-id*. All SDP ID keepalive SDP Echo Replies are sent using generic IP/GRE OAM encapsulation.

When a keepalive response is received that indicates an error condition, the SDP ID will immediately be brought operationally down. Once a response is received that indicates the error has cleared and the **hold-down-time** interval has expired, the SDP ID will be eligible to be put into the operationally up state. If no other condition prevents the operational change, the SDP ID will enter the operational state.

A set of event counters track the number of keepalive requests sent, the size of the message sent, non-error replies received and error replies received. A keepalive state value is kept indicating the last response event. A keepalive state timestamp value is kept indicating the time of the last event. With each keepalive event change, a log message is generated indicating the event type and the timestamp value.

Table 44: Keepalive Interpretation and Effect of SDP Echo Reply describes the keepalive interpretation of SDP echo reply response conditions and the effect on the SDP ID operational status.

Table 44: Keepalive Interpretation and Effect of SDP Echo Reply

Result of Request	Stored Response State	Operational State
keepalive request timeout without reply	Request Timeout	Down
keepalive request not sent due to non-existent *orig-sdp-id* (This condition should not occur)	Orig-SDP Non-Existent	Down

Result of Request	Stored Response State	Operational State
keepalive request not sent due to administratively down *orig-sdp-id*	Orig-SDP Admin-Down	Down
keepalive reply received, invalid origination-id	Far End: Originator-ID Invalid	Down
keepalive reply received, invalid responder-id	Far End: Responder-ID Error	Down
keepalive reply received, No Error	Success	Up (If no other condition prevents)

Platforms

7705 SAR Gen 2

15.4 keep-alive-interval

keep-alive-interval

Syntax

keep-alive-interval *interval*

no keep-alive-interval

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-lag keep-alive-interval)

Full Context

configure redundancy multi-chassis peer mc-lag keep-alive-interval

Description

This command sets the interval at which keep-alive messages are exchanged between two systems participating in MC-LAG. These keep-alive messages are used to determine remote-node failure and the interval is set in deciseconds.

The **no** form of this command sets the interval to default value.

Default

keep-alive-interval 10

Parameters

interval

The time interval expressed in tenths of a second.

Values 5 to 500

Platforms

7705 SAR Gen 2

keep-alive-interval

Syntax

keep-alive-interval *interval*

no keep-alive-interval

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ep keep-alive-interval)

Full Context

configure redundancy multi-chassis peer mc-endpoint keep-alive-interval

Description

This command sets the interval at which keep-alive messages are exchanged between two systems participating in MC-EP when bfd is not enabled or is down. These fast keep-alive messages are used to determine remote-node failure and the interval is set in deciseconds.

The **no** form of this command sets the interval to default value

Default

no keep-alive-interval

Parameters

interval

The time interval expressed in tenths of a second.

Values 5 to 500

Platforms

7705 SAR Gen 2

keep-alive-interval

Syntax

keep-alive-interval *interval*

no keep-alive-interval

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ipsec keep-alive-interval)

Full Context

configure redundancy multi-chassis peer mc-ipsec keep-alive-interval

Description

This command specifies the time interval of the mastership election protocol sending keep-alive packet.

The **no** form of this command reverts to the default.

Default

keep-alive-interval 10

Parameters

interval

Specifies the keep alive interval in tenths of seconds.

Values 5 to 500

Platforms

7705 SAR Gen 2

15.5 keep-multiplier

keep-multiplier

Syntax

[**no**] **keep-multiplier** *number*

no keep-multiplier

Context

[Tree] (config>router>rsvp keep-multiplier)

Full Context

configure router rsvp keep-multiplier

Description

The **keep-multiplier** *number* is an integer used by RSVP to declare that a reservation is down or the neighbor is down.

The **no** form of this command reverts to the default value.

Default

keep-multiplier 3

Parameters

number

Specifies the **keep-multiplier** value.

Values 1 to 255

Platforms

7705 SAR Gen 2

15.6 keepalive

keepalive

Syntax

keepalive *seconds*

no keepalive

Context

[Tree] (config>service>vprn>bgp>group>neighbor keepalive)

[Tree] (config>service>vprn>bgp>group keepalive)

[Tree] (config>service>vprn>bgp keepalive)

Full Context

configure service vprn bgp group neighbor keepalive

configure service vprn bgp group keepalive

configure service vprn bgp keepalive

Description

This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires. The **seconds** parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **keepalive** value is generally one-third of the **hold-time** interval. Even though the OS implementation allows the **keepalive** value and the **hold-time** interval to be independently set, under the following circumstances, the configured **keepalive** value is overridden by the **hold-time** value:

If the specified **keepalive** value is greater than the configured **hold-time**, then the specified value is ignored, and the **keepalive** is set to one third of the current **hold-time** value.

If the specified **hold-time** interval is less than the configured **hold-time** value, then the **keepalive** value is reset to one third of the specified **hold-time** interval.

If the **hold-time** interval is set to zero, then the configured value of the **keepalive** value is ignored. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

keepalive 30

Parameters

seconds

The keepalive timer in seconds, expressed as a decimal integer.

Values 0 to 21845

Platforms

7705 SAR Gen 2

keepalive

Syntax

keepalive *timeout factor*

no keepalive

Context

[Tree] (config>router>ldp>if-params>if>ipv4 keepalive)

[Tree] (config>router>ldp>if-params>if>ipv6 keepalive)

[Tree] (config>router>ldp>targ-session>peer keepalive)

[Tree] (config>router>ldp>targ-session>peer-template keepalive)

[Tree] (config>router>ldp>if-params>ipv6 keepalive)

[Tree] (config>router>ldp>targ-session>ipv6 keepalive)

[Tree] (config>router>ldp>targ-session>ipv4 keepalive)

[Tree] (config>router>ldp>if-params>ipv4 keepalive)

Full Context

configure router ldp interface-parameters interface ipv4 keepalive

configure router ldp interface-parameters interface ipv6 keepalive

configure router ldp targeted-session peer keepalive

configure router ldp targeted-session peer-template keepalive

configure router ldp interface-parameters ipv6 keepalive

configure router ldp targeted-session ipv6 keepalive

configure router ldp targeted-session ipv4 keepalive

configure router ldp interface-parameters ipv4 keepalive

Description

This command configures the time interval (in s), that LDP waits before tearing down the session. The **factor** parameter derives the keepalive interval.

If no LDP messages are exchanged for the configured time interval, the LDP session is torn down. Keepalive timeout is usually three times the keepalive interval. To maintain the session permanently, regardless of the activity, set the value to zero.

When LDP session is being set up, the keepalive timeout is negotiated to the lower of the two peers. Once an operational value is agreed upon, the keepalive factor is used to derive the value of the keepalive interval.

The **no** form of the command at the interface-parameters and targeted-session levels sets the **keepalive timeout** and the **keepalive factor** to the default value.

The **no** form of this command, at the interface level, sets the **keepalive timeout** and the **keepalive factor** to the value defined under the **interface-parameters** level.

The **no** form of this command, at the peer level, sets the **keepalive timeout** and the **keepalive factor** to the value defined under the **targeted-session** level.

The session must be flapped for the new settings to operate.

Default

Table 45: Timeout Factor Defaults lists the default values.

Table 45: Timeout Factor Defaults

Context	Timeout	Factor
config>router>ldp>if-params	30	3
config>router>ldp>targ-session	40	4
config>router>ldp>if-params>if	Inherits values from interface-parameters context.	
config>router>ldp>targ-session>peer	Inherits values from targeted-session context.	

Parameters

timeout

Configures the time interval, in seconds, that LDP waits before tearing down the session.

Values 1 to 65535

factor

Specifies the number of keepalive messages, expressed as a decimal integer, that should be sent on an idle LDP session in the keepalive timeout interval.

Values 1 to 255

Platforms

7705 SAR Gen 2

keepalive

Syntax

[**no**] **keepalive**

Context

[Tree] (debug>router>ldp>peer>packet keepalive)

Full Context

debug router ldp peer packet keepalive

Description

This command enables debugging for LDP Keepalive packets.

The **no** form of the command disables the debugging output.

Platforms

7705 SAR Gen 2

keepalive

Syntax

keepalive *seconds*

no keepalive

Context

[Tree] (config>router>pcep>pcc keepalive)

Full Context

configure router pcep pcc keepalive

Description

This command configures the PCEP session keep-alive value. A PCEP speaker (PCC or PCE) must send a keep-alive message if no other PCEP message is sent to the peer at the expiry of this timer. This timer is restarted every time a PCEP message or keep-alive message is sent.

The keep-alive mechanism is asymmetric, meaning that each peer can use a different keep-alive timer value at its end.

The **no** form of the command returns the keep-alive timer to the default value.

Default

keepalive 30

Parameters

seconds

the keep-alive value, in seconds

Values 1 to 255

Platforms

7705 SAR Gen 2

keepalive

Syntax

keepalive *seconds*

no keepalive

Context

[Tree] (config>router>bgp>group keepalive)

[Tree] (config>router>bgp>group>neighbor keepalive)

[Tree] (config>router>bgp keepalive)

Full Context

configure router bgp group keepalive

configure router bgp group neighbor keepalive

configure router bgp keepalive

Description

This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires.

The **keepalive** parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **keepalive** value is generally one-third of the **hold-time** interval. Even though the implementation allows the **keepalive** value and the **hold-time** interval to be independently set, under the following circumstances, the configured **keepalive** value is overridden by the **hold-time** value:

- If the specified **keepalive** value is greater than the configured **hold-time**, then the specified value is ignored and the **keepalive** is set to one third of the current **hold-time** value.
- If the specified **hold-time** interval is less than the configured **keepalive** value, then the **keepalive** value is reset to one third of the specified **hold-time** interval.
- If the **hold-time** interval is set to zero, then the configured value of the **keepalive** value is ignored. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

keepalive 30

Parameters

seconds

Specifies the keepalive timer, in seconds, expressed as a decimal integer.

Values 0 to 21845

Platforms

7705 SAR Gen 2

keepalive

Syntax

keepalive [**neighbor** *ip-addr* | **group** *name*]

no keepalive

Context

[Tree] (debug>router>bgp keepalive)

Full Context

debug router bgp keepalive

Description

This command decodes and logs all sent and received keepalive messages in the debug log.

The **no** form of this command disables the debugging.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x [-interface] (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H
- d: [0 to 255]D
- interface: up to 32 characters for link local addresses

group *name*

Debugs only events affecting the specified peer group name, up to 64 characters, and associated neighbors.

Platforms

7705 SAR Gen 2

15.7 kernel

kernel

Syntax

kernel -password *password*

no kernel

Context

[Tree] (environment kernel)

Full Context

environment kernel

Description

This command allows Nokia technical support to access the **kernel** commands. **kernel** commands are used only by Nokia technical support for troubleshooting.

The **no** form of this command disables the **kernel** commands.

Parameters

password

Specifies the password to access the **kernel** commands, up to 256 characters.

Platforms

7705 SAR Gen 2

15.8 kex

kex

Syntax

kex *index* **name** *kex-name*

no kex *index*

Context

[Tree] (config>system>security>ssh>server-kex-list kex)

[Tree] (config>system>security>ssh>client-kex-list kex)

Full Context

configure system security ssh server-kex-list kex

configure system security ssh client-kex-list kex

Description

This command configures phase 1 SSH v2 KEX algorithms for SR OS as an SSH server or an SSH client. By default, the client and server lists are empty. If the user configures this list, SSH uses the hard-coded list with the first-listed algorithm having the highest priority and so on. An empty server or client list is the default list and contains the following algorithms:

ecdh-sha2-nistp512

ecdh-sha2-nistp384

ecdh-sha2-nistp256

diffie-hellman-group16-sha512

diffie-hellman-group14-sha256

diffie-hellman-group14-sha1

diffie-hellman-group-exchange-sha1

diffie-hellman-group1-sha1

The **no** form of this command removes the specified KEX index. If all KEX indexes are removed, the default list is used again.

Parameters

index

Specifies the index of the algorithm in the list. The lowest index in the list is negotiated first on the SSH negotiation list, while the highest index is at the bottom of the SSH negotiation list.

Values 1 to 255

kex-name

Specifies the KEX algorithm for computing the shared secret key.

Values diffie-hellman-group16-sha512, diffie-hellman-group14-sha256, diffie-hellman-group14-sha1, diffie-hellman-group-exchange-sha1, diffie-hellman-group1-sha1

Platforms

7705 SAR Gen 2

15.9 key

key

Syntax

key *key-file-name*

Context

[Tree] (config>system>security>pki>cert-auto-upd>cert key)

Full Context

configure system security pki certificate-auto-update cert key

Description

This command configures the filename of the key corresponding to the certificate.

Parameters

key-file-name

Specifies the filename of the key.

Platforms

7705 SAR Gen 2

key

Syntax

key *key-filename*

no key

Context

[Tree] (config>ipsec>cert-profile>entry key)

Full Context

configure ipsec cert-profile entry key

Description

This command specifies the filename of an imported key for the **cert-profile entry**.

The **no** form of this command removes the key filename from the entry configuration.

Default

no key

Parameters

key-filename

Specifies the filename of an imported key.

Platforms

7705 SAR Gen 2

key

Syntax

key *password* [**hash** | **hash2** | **custom**] **reference** *reference-number*

no key reference *reference-number*

Context

[Tree] (config>system>security>pki>ca-profile>cmpv2>key-list key)

Full Context

configure system security pki ca-profile cmpv2 key-list key

Description

This command specifies a pre-shared key used for CMPv2 initial registration. Multiples of key commands are allowed to be configured under this context.

The password and reference-number is distributed by the CA via out-of-band means.

The configured password is stored in configuration file in an encrypted form by using SR OS hash2 algorithm.

The **no** form of this command removes the parameters from the configuration.

Parameters

password

Specifies a printable ASCII string, up to 64 characters.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

reference *reference-number*

Specifies a printable ASCII string, up to 64 characters in length.

Platforms

7705 SAR Gen 2

key

Syntax

key *key-filename*

no key

Context

[Tree] (config>system>security>tls>cert-profile>entry key)

Full Context

configure system security tls cert-profile entry key

Description

This command specifies the file name of an imported key for the **cert-profile** entry.

The **no** form of the command removes the key.

Default

no key

Parameters

key-filename

Specifies the file name of the key.

Platforms

7705 SAR Gen 2

15.10 key-generation

key-generation

Syntax

key-generation dsa size *bits*

key-generation ecdsa curve *curve*

key-generation rsa size *bits*

key-generation same-as-existing-key

Context

[Tree] (config>system>security>pki>cert-upd-prof key-generation)

Full Context

configure system security pki certificate-update-profile key-generation

Description

This command configures the key generation algorithm and behavior.

Default

key-generation same-as-existing-key

Parameters

bits

Specifies the size in bits..

Values 512 to 8192

Default 2048

curve

Specifies the elliptic curve for key generation.

Values secp256r1, secp384r1, secp521r1

Default secp256r1

same-as-existing-key
Specifies to use the same algorithm and key or size curve as the existing key.

Platforms
7705 SAR Gen 2

15.11 key-list

key-list

Syntax
key-list

Context
[Tree] (config>system>security>pki>ca-profile>cmpv2 key-list)

Full Context
configure system security pki ca-profile cmpv2 key-list

Description
This command enables the context to configure pre-shared key list parameters.

Platforms
7705 SAR Gen 2

15.12 key-re-exchange

key-re-exchange

Syntax
key-re-exchange

Context
[Tree] (config>system>security>ssh key-re-exchange)

Full Context

configure system security ssh key-re-exchange

Description

This command enables the key re-exchange context.

Platforms

7705 SAR Gen 2

15.13 key-rollover-interval

key-rollover-interval

Syntax

key-rollover-interval *key-rollover-interval*

Context

[Tree] (config>service>vprn>ospf3>area key-rollover-interval)

Full Context

configure service vprn ospf3 area key-rollover-interval

Description

This command configures the key rollover interval.

The **no** form of this command reverts to the default.

Default

key-rollover-interval 10

Parameters

key-rollover-interval

Specifies the time, in seconds, after which a key rollover will start.

Values 10 to 300

Platforms

7705 SAR Gen 2

key-rollover-interval

Syntax

key-rollover-interval *seconds*

Context

[Tree] (config>router>ospf3>area key-rollover-interval)

Full Context

configure router ospf3 area key-rollover-interval

Description

This command configures the key rollover interval.

Default

key-rollover-interval 10

Parameters

seconds

Specifies the time, in seconds, after which a key rollover will start.

Values 10 to 300

Platforms

7705 SAR Gen 2

15.14 key-update

key-update

Syntax

key-update ca *ca-profile-name* **newkey** *key-filename* **oldkey** *key-filename* **oldcert** *cert-filename* [**hash-alg** *hash-algorithm*] **save-as** *save-path-of-result-cert*

Context

[Tree] (admin>certificate>cmpv2 key-update)

Full Context

admin certificate cmpv2 key-update

Description

This command requests a new certificate from the CA to update an existing certificate due to reasons such as **key refresh** or **replacing compromised key**.

In some cases, the CA may not return certificate immediately, due to reasons such as request processing need manual intervention. In such cases, the admin certificate cmpv2 poll command can be used to poll the status of the request.

Parameters

ca-profile-name

Specifies a ca-profile name which includes CMP server information, up to 32 characters.

newkey ***key-filename***

Specifies the key file of the requesting certificate, up to 95 characters.

oldkey ***key-filename***

Specifies the key to be replaced, up to 95 characters.

cert-filename

Specifies the file name of an imported certificate to be replaced, up to 95 characters.

hash-algorithm

Specifies the hash algorithm for RSA key.

Values md5,sha1,sha224,sha256,sha384,sha512

save-path-of-result-cert

Specifies the save full path name of saving the result certificate, up to 200 characters.

Platforms

7705 SAR Gen 2

15.15 key-value

key-value

Syntax

key-value *public-key-value*

no key-value

Context

[Tree] (config>system>security>user>public-keys>rsa>rsa-key key-value)

[Tree] (config>system>security>user>public-keys>ecdsa>ecdsa-key key-value)

Full Context

configure system security user public-keys rsa rsa-key key-value

configure system security user public-keys ecdsa ecdsa-key key-value

Description

This command configures a value for the RSA or ECDSA public key. The public key must be enclosed in quotation marks. For RSA, the key is between 768 and 4096 bits. For ECDSA, the key is between 1 and 1024 bits.

Default

no key-value

Parameters

public-key-value

Specifies the public key value, up to 800 characters for RSA and up to 255 characters for ECDSA.

Platforms

7705 SAR Gen 2

15.16 keychain

keychain

Syntax

[**no**] **keychain** *keychain-name*

Context

[Tree] (config>system>security keychain)

Full Context

configure system security keychain

Description

This command enables the context to configure keychain parameters. A keychain must be configured on the system before it can be applied to a session.

The **no** form of this command removes the keychain nodal context and everything under it from the configuration. If the keychain to be removed is in use when the no keychain command is entered, the command will not be accepted and an error indicating that the keychain is in use will be printed.

Parameters

keychain-name

Specifies a keychain name which identifies this particular keychain entry.

Values An ASCII string up to 32 characters.

Platforms
7705 SAR Gen 2

15.17 keygroup-name

keygroup-name

Syntax
keygroup-name *keygroup-name*
no keygroup-name

Context
[Tree] (config>grp-encryp>encryp-keygrp keygroup-name)

Full Context
configure group-encryption encryption-keygroup keygroup-name

Description
This command is used to name the key group. The key group name can be used to reference a key group when configuring services or displaying information.

The **no** form of the command reverts to the default value.

Parameters
keygroup-name
The name of the key group, up to 64 characters.

Platforms
7705 SAR Gen 2

15.18 kill-session

kill-session

Syntax
[**no**] **kill-session**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization kill-session)

Full Context

configure system security profile netconf base-op-authorization kill-session

Description

This command authorizes a user associated with the profile to send a NETCONF <kill-session> RPC. This kill session operation allows a NETCONF client to kill another NETCONF session, but not the session in which the operation is requested.

The **no** form of the command denies the user from requesting a kill-session.

Default

no kill-session

Platforms

7705 SAR Gen 2

16 I Commands

16.1 l2pt-termination

l2pt-termination

Syntax

l2pt-termination [**cdp**] [**dtp**] [**pagp**] [**stp**] [**udld**] [**vtp**]

no l2pt-termination

Context

[Tree] (config>service>template>vpls-sap-template l2pt-termination)

[Tree] (config>service>vpls>spoke-sdp l2pt-termination)

[Tree] (config>service>vpls>sap l2pt-termination)

Full Context

configure service template vpls-sap-template l2pt-termination

configure service vpls spoke-sdp l2pt-termination

configure service vpls sap l2pt-termination

Description

This command enables Layer 2 Protocol Tunneling (L2PT) termination on a specified SAP or spoke-SDP. L2PT termination is supported only for STP BPDUs. PDUs of other protocols are discarded.

This feature can be enabled only if STP is disabled in the context of the specified VPLS service.

The **no** form of this command reverts to the default.

Default

no l2pt-termination

Parameters

cdp

Specifies the Cisco discovery protocol

dtp

Specifies the dynamic trunking protocol

pagp

Specifies the port aggregation protocol

stp

Specifies all spanning tree protocols: stp, rstp, mstp, pvst (default)

udld

Specifies unidirectional link detection

vtp

Specifies the virtual trunk protocol

Platforms

7705 SAR Gen 2

l2pt-termination

Syntax

l2pt-termination [**cdp**] [**dtp**] [**pagp**] [**stp**] [**udld**] [**vtp**]

no l2pt-termination

Context

[Tree] (config>service>pw-template l2pt-termination)

Full Context

configure service pw-template l2pt-termination

Description

This command enables Layer 2 Protocol Tunneling (L2PT) termination on a given SAP or spoke SDP. L2PT termination will be supported only for STP BPDUs. PDUs of other protocols will be discarded.

This feature can be enabled only if STP is disabled in the context of the given VPLS service.

Default

no l2pt-termination

Parameters

cdp

Specifies the Cisco discovery protocol.

dtp

Specifies the dynamic trunking protocol.

pagp

Specifies the port aggregation protocol.

stp

Specifies all spanning tree protocols: stp, rstp, mstp, pvst (default).

udld

Specifies unidirectional link detection.

vtp

Specifies the virtual trunk protocol.

Platforms

7705 SAR Gen 2

16.2 l2tp

l2tp

Syntax

[**no**] **l2tp**

Context

[Tree] (config>redundancy>multi-chassis>peer>sync l2tp)

Full Context

configure redundancy multi-chassis peer sync l2tp

Description

This command enables L2TP.

The **no** form of this command disables L2TP.

Platforms

7705 SAR Gen 2

16.3 l3-ring

l3-ring

Syntax

l3-ring *name* [**create**]

no l3-ring name

Context

[Tree] (config>redundancy>mc>peer>mcr l3-ring)

Full Context

configure redundancy multi-chassis peer mc-ring l3-ring

Description

This command configures a Layer 3 multi-chassis ring.

The **no** form of this command reverts to the default.

Platforms

7705 SAR Gen 2

16.4 l4-load-balancing

l4-load-balancing

Syntax

[**no**] **l4-load-balancing**

Context

[Tree] (config>system>load-balancing l4-load-balancing)

Full Context

configure system load-balancing l4-load-balancing

Description

This command configures system-wide Layer 4 load balancing. The configuration at the system level can enable or disable load balancing based on Layer 4 fields. If enabled, the Layer 4 source and destination port fields will be included in hashing calculation for TCP/UDP packets.

The hashing algorithm addresses finer spraying granularity where many hosts are connected to the network.

To address more efficient traffic distribution between network links (forming a LAG group), a hashing algorithm extension takes into account L4 information (that is, src/dst L4-protocol port).

The hashing index can be calculated according to the following algorithm:

Example:

```
– If [(TCP or UDP traffic) & enabled]
    – hash (TCP/UDP ports, IP addresses)
– else if (IP traffic)
    – hash (IP addresses)
– else
    – hash (MAC addresses)
– endif
```

This algorithm will be used in all cases where IP information in per-packet hashing is included (refer to "Traffic Load Balancing Options" in the *7705 SAR Gen 2 Interface Configuration Guide*). However, the Layer 4 information (TCP/UDP ports) will not be used for fragmented packets.

Default

no l4-load-balancing

Platforms

7705 SAR Gen 2

16.5 l4-src-port

l4-src-port

Syntax

l4-src-port *port* [*mask*]

no l4-src-port

Context

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter>entry l4-src-port)

[Tree] (config>system>security>mgmt-access-filter>ip-filter>entry l4-src-port)

Full Context

configure system security management-access-filter ipv6-filter entry l4-src-port

configure system security management-access-filter ip-filter entry l4-src-port

Description

This command configures a destination TCP or UDP port number or port range for a management access filter match criterion.

The **no** form of this command reverts to the default values.

Default

no l4-src-port

Parameters

port

Specifies the destination TCP or UDP port number as a match criterion.

Values 1 to 65535

Default 6 (exact match)

mask

Specifies the mask used to select a range of source port numbers. Table 46: Format Styles to Configure Mask lists the format styles to configure the 16-bit mask.

Table 46: Format Styles to Configure Mask

Format Style	Format Syntax	Example
Decimal	DDDDD	63488
Hexadecimal	0xHHHH	0xF800
Binary	0bBBBBBBBBBBBBBBB	0b1111100000000000

To select a range from 1024 up to 2047, specify 1024 and 0xFC00 for port and mask, respectively.

Values 1 to 65535 (decimal)

Default 65535 (exact match)

Platforms

7705 SAR Gen 2

16.6 label

label

Syntax

label [**detail**]

no label

Context

[Tree] (debug>router>ldp>peer>packet label)

Full Context

debug router ldp peer packet label

Description

This command enables debugging for LDP Label packets.

The **no** form of the command disables the debugging output.

Parameters

detail

Displays detailed information.

Platforms

7705 SAR Gen 2

16.7 label-allocation

label-allocation

Syntax

label-allocation

Context

[Tree] (config>router>bgp label-allocation)

Full Context

configure router bgp label-allocation

Description

This commands enables the context to configure the allocation of MPLS labels to specific BGP routes.

Platforms

7705 SAR Gen 2

16.8 label-ipv4

label-ipv4

Syntax

label-ipv4 *max-paths* [**ebgp** *ebgp-max-paths*] [**ibgp** *ibgp-max-paths*] [**restrict** {**same-neighbor-as** | **exact-as-path**}] [**unequal-cost**]

no label-ipv4

Context

[Tree] (config>service>vprn>bgp>multi-path label-ipv4)

Full Context

configure service vprn bgp multi-path label-ipv4

Description

This command sets ECMP multipath parameters that apply only to the label unicast IPv4 address family.

When multipath is enabled, traffic to the destination is load-shared across a set of paths (BGP routes) that the BGP decision process considers equal to the best path. The distribution of traffic over the multiple paths may or may not be equal. The distribution is based on weights derived from the Link Bandwidth Extended Community.

For more information about the criteria a non-best route must meet to qualify as a multipath, see "BGP route installation in the route table" in the *7705 SAR Gen 2 Unicast Routing Protocols Guide*.

The **no** form of this command removes label-IPv4-specific overrides.

Default

no label-ipv4

Parameters

max-paths

Specifies the maximum number of multipaths per prefix or NLRI. Setting this value to 1 disables multipath. This limit only applies if neither the *ebgp-max-paths* limit nor the *ibgp-max-paths* limit apply.

Values 1 to 64

ebgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an EBGP learned route. If the **ebgp** option is configured, this value overrides the *max-paths* limit. If the best path is an EBGP learned route and this value is set to 1, multipath is disabled.

Values 1 to 64

ibgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an IBGP learned route. If the **ibgp** option is configured, this value overrides the *max-paths* limit. If the best path is an IBGP learned route and this value is set to 1, multipath is disabled.

Values 1 to 64

restrict same-neighbor-as

Specifies that the non-best path must have the same neighbor AS in its AS path as the best path.

restrict exact-as-path-as

Specifies that the non-best path must have the same AS path as the best path.

unequal-cost

Instructs BGP to ignore differences in the next-hop cost only when determining eligible multipaths.

Platforms

7705 SAR Gen 2

label-ipv4

Syntax

label-ipv4 send *send-limit*

label-ipv4 send *send-limit* **receive** [**none**]

no label-ipv4

Context

[Tree] (config>router>bgp>add-paths label-ipv4)

[Tree] (config>router>bgp>group>add-paths label-ipv4)

[Tree] (config>router>bgp>group>neighbor>add-paths label-ipv4)

Full Context

configure router bgp add-paths label-ipv4

configure router bgp group add-paths label-ipv4

configure router bgp group neighbor add-paths label-ipv4

Description

This command configures the add-paths capability for labeled-unicast IPv4 routes. By default, add-paths is not enabled for labeled-unicast IPv4 routes.

The maximum number of labeled-unicast paths per IPv4 prefix to send is the configured *send-limit*, which is a mandatory parameter. The capability to receive multiple labeled-unicast paths per prefix from a peer is configurable using the **receive** keyword, which is optional. If the **receive** keyword is not included in the command, receive capability is enabled by default.

The **no** form of this command disables add-paths support for labeled-unicast IPv4 routes, causing sessions established using add-paths for labeled-unicast IPv4 to go down and come back up without the add-paths capability.

Default

no label-ipv4

Parameters

send-limit

Specifies the maximum number of paths per labeled-unicast IPv4 prefix that are allowed to be advertised to add-paths peers. (The actual number of advertised routes may be

less.) If the value is none, the router does not negotiate the send capability with respect to label-IPv4 AFI/SAFI. If the value is **multipaths**, then BGP advertises all the used BGP multipaths for each IPv4 NLRI if the peer has signaled support to receive multiple add paths.

Values 1 to 16, none, multipaths

receive

Specifies the router negotiates to receive multiple labeled-unicast routes per IPv4 prefix.

none

Specifies that the router does not negotiate to receive multiple labeled-unicast routes per IPv4 prefix.

Platforms

7705 SAR Gen 2

label-ipv4

Syntax

label-ipv4 *max-paths* [**ebgp** *ebgp-max-paths*] [**ibgp** *ibgp-max-paths*] [**restrict** {**same-neighbor-as** | **exact-as-path**}] [**unequal-cost**]

no label-ipv4

Context

[Tree] (config>router>bgp>multi-path label-ipv4)

Full Context

configure router bgp multi-path label-ipv4

Description

This command sets ECMP multipath parameters that apply only to the label IPv4 unicast address family. These settings override the values set by the **maximum-paths** command.

When multipath is enabled, traffic to the destination is load-shared across a set of paths (BGP routes) that the BGP decision process considers equal to the best path. The actual distribution of traffic over the multiple paths may be equal or unequal (that is, based on weights derived from the Link Bandwidth Extended Community).

The **no** form of this command removes label-IPv4-specific overrides.

Default

no label-ipv4

Parameters

max-paths

Specifies the maximum number of multipaths per prefix/NLRI if *ebgp-max-paths* or *ibgp-max-paths* does not apply.

Values 1 to 64

ebgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an EBGP learned route.

Values 1 to 64

ibgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an IBGP learned route.

Values 1 to 64

restrict same-neighbor-as

Specifies that the non-best path must have the same neighbor AS in its AS path as the best path.

restrict exact-as-path

Specifies that the non-best path must have the same AS path as the best path.

unequal-cost

Instructs BGP to ignore differences in the next-hop cost only when determining eligible multipaths.

Platforms

7705 SAR Gen 2

16.9 label-ipv6

label-ipv6

Syntax

label-ipv6 *max-paths* [**ebgp** *ebgp-max-paths*] [**ibgp** *ibgp-max-paths*] [**restrict** {**same-neighbor-as** | **exact-as-path**}] [**unequal-cost**]

no label-ipv6

Context

[Tree] (config>service>vprn>bgp>multi-path label-ipv6)

Full Context

configure service vprn bgp multi-path label-ipv6

Description

This command sets ECMP multipath parameters that apply only to the label unicast IPv6 address family.

When multipath is enabled, traffic to the destination is load-shared across a set of paths (BGP routes) that the BGP decision process considers equal to the best path. The distribution of traffic over the multiple paths may or may not be equal. The distribution is based on weights derived from the Link Bandwidth Extended Community.

For more information about the criteria a non-best route must meet to qualify as a multipath, see "BGP route installation in the route table" in the *7705 SAR Gen 2 Unicast Routing Protocols Guide*.

The **no** form of this command removes label-IPv6-specific overrides.

Default

no label-ipv6

Parameters

max-paths

Specifies the maximum number of multipaths per prefix or NLRI. Setting this value to 1 disables multipath. This limit only applies if neither the *ebgp-max-paths* limit nor the *ibgp-max-paths* limit apply.

Values 1 to 64

ebgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an EBGP learned route. If the **ebgp** option is configured, this value overrides the *max-paths* limit. If the best path is an EBGP learned route and this value is set to 1, multipath is disabled.

Values 1 to 64

ibgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an IBGP learned route. If the **ibgp** option is configured, this value overrides the *max-paths* limit. If the best path is an IBGP learned route and this value is set to 1, multipath is disabled.

Values 1 to 64

restrict same-neighbor-as

Specifies that the non-best path must have the same neighbor AS in its AS path as the best path.

restrict exact-as-path-as

Specifies that the non-best path must have the same AS path as the best path.

unequal-cost

Instructs BGP to ignore differences in the next-hop cost only when determining eligible multipaths.

Platforms

7705 SAR Gen 2

label-ipv6

Syntax

label-ipv6 send *send-limit*

label-ipv6 send *send-limit* **receive** [**none**]

no label-ipv6

Context

[Tree] (config>router>bgp>group>neighbor>add-paths label-ipv6)

[Tree] (config>router>bgp>group>add-paths label-ipv6)

[Tree] (config>router>bgp>add-paths label-ipv6)

Full Context

configure router bgp group neighbor add-paths label-ipv6

configure router bgp group add-paths label-ipv6

configure router bgp add-paths label-ipv6

Description

This command configures the add-paths capability for labeled-unicast IPv6 routes. By default, add-paths is not enabled for labeled-unicast IPv6 routes.

The maximum number of labeled-unicast paths per IPv6 prefix to send is the configured *send-limit*, which is a mandatory parameter. The capability to receive multiple labeled-unicast paths per prefix from a peer is configurable using the **receive** keyword, which is optional. If the **receive** keyword is not included in the command, receive capability is enabled by default.

The **no** form of this command disables add-paths support for labeled-unicast IPv6 routes, causing sessions established using add-paths for labeled-unicast IPv6 to go down and come back up without the add-paths capability.

Default

no label-ipv6

Parameters

send-limit

Specifies the maximum number of paths per labeled-unicast IPv6 prefix that are allowed to be advertised to add-paths peers. (The actual number of advertised routes may be

less.) If the value is none, the router does not negotiate the send capability with respect to label-IPv6 AFI/SAFI. If the value is **multipaths**, then BGP advertises all the used BGP multipaths for each IPv6 NLRI if the peer has signaled support to receive multiple add paths.

Values 1 to 16, none, multipaths

receive

Specifies that the router negotiates to receive multiple labeled-unicast routes per IPv6 prefix.

none

Specifies that the router does not negotiate to receive multiple labeled-unicast routes per IPv6 prefix.

Platforms

7705 SAR Gen 2

label-ipv6

Syntax

label-ipv6 *max-paths* [**ebgp** *ebgp-max-paths*] [**ibgp** *ibgp-max-paths*] [**restrict** {**same-neighbor-as** | **exact-as-path**}] [**unequal-cost**]

no label-ipv6

Context

[Tree] (config>router>bgp>multi-path label-ipv6)

Full Context

configure router bgp multi-path label-ipv6

Description

This command sets ECMP multipath parameters that apply only to the label unicast IPv6 address family. These settings override the values set by the **maximum-paths** command.

When multipath is enabled, traffic to the destination is load-shared across a set of paths (BGP routes) that the BGP decision process considers equal to the best path. The actual distribution of traffic over the multiple paths may be equal or unequal (that is, based on weights derived from the Link Bandwidth Extended Community).

The **no** form of this command removes label-IPv6-specific overrides.

Default

no label-ipv6

Parameters

max-paths

Specifies the maximum number of multipaths per prefix/NLRI if *ebgp-max-paths* or *ibgp-max-paths* does not apply.

Values 1 to 64

ebgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an EBGP learned route.

Values 1 to 64

ibgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an IBGP learned route.

Values 1 to 64

restrict same-neighbor-as

Specifies that the non-best path must have the same neighbor AS in its AS path as the best path.

restrict exact-as-path

Specifies that the non-best path must have the same AS path as the best path.

unequal-cost

Instructs BGP to ignore differences in the next-hop cost only when determining eligible multipaths.

Platforms

7705 SAR Gen 2

label-ipv6

Syntax

label-ipv6

Context

[Tree] (config>router>bgp>label-allocation label-ipv6)

Full Context

configure router bgp label-allocation label-ipv6

Description

Commands in this context configure advertised label IPv6 programming rules.

Platforms

7705 SAR Gen 2

label-ipv6

Syntax

label-ipv6

Context

[Tree] (config>service>vprn>bgp>rib-management label-ipv6)

Full Context

configure service vprn bgp rib-management label-ipv6

Description

Commands in this context configure labeled IPv6 RIB.

Platforms

7705 SAR Gen 2

16.10 label-ipv6-explicit-null

label-ipv6-explicit-null

Syntax

[**no**] **label-ipv6-explicit-null**

Context

[Tree] (config>router>bgp>next-hop-res>lbl-routes>use-bgp-routes label-ipv6-explicit-null)

Full Context

configure router bgp next-hop-resolution labeled-routes use-bgp-routes label-ipv6-explicit-null

Description

This command allows a labelled IPv6 route with the explicit-null label to be resolved by other labelled IPv6 routes with the explicit-null label, and also by unlabeled IPv4 routes and unlabeled IPv6 routes that are resolved by static routes, interface routes, or tunnels. Up to four levels of recursive resolution are supported when the top route is a labelled IPv6 route with an explicit-null label.

Regardless of setting, a labelled IPv6 route with a regular label (other than explicit-null) is never resolved by other labelled IPv6 routes.

The **no** form of this command disables the label-ipv6-explicit-null functionality. When disabled, a labeled IPv6 route cannot be resolved by other labeled IPv6 routes.

Default

no label-ipv6-explicit-null

Platforms

7705 SAR Gen 2

16.11 label-map

label-map

Syntax

[**no**] **label-map** *in-label*

Context

[Tree] (config>router>mpls>interface label-map)

Full Context

configure router mpls interface label-map

Description

This command is used on transit routers when a static LSP is defined. The static LSP on the ingress router is initiated using the **config router mpls static-lsp** *lsp-name* command. An *in-label* can be associated with either a **pop** or a **swap** action, but not both. If both actions are specified, the last action specified takes effect.

The **no** form of this command deletes the static LSP configuration associated with the *in-label*.

Parameters

in-label

Specifies the incoming MPLS label on which to match.

Values 32 to 1023

Platforms

7705 SAR Gen 2

16.12 label-mode

label-mode

Syntax

label-mode {**vrf** | **next-hop**}

no label-mode

Context

[Tree] (config>service>vprn label-mode)

Full Context

configure service vprn label-mode

Description

This command controls the method by which service labels are allocated to routes exported by the VPRN as BGP-VPN routes. The **vrf** option selects service label per VRF mode while the **next-hop** option selects service label per next-hop mode.

The **no** form of this command sets the mode to the default mode of service label per VRF.

Default

no label-mode

Parameters

vrf

Selects service label per VRF mode.

next-hop

Selects service label per next-hop mode.

Platforms

7705 SAR Gen 2

16.13 label-preference

label-preference

Syntax

label-preference *value*

no label-preference

Context

[Tree] (config>service>vprn>bgp label-preference)

[Tree] (config>service>vprn>bgp>group label-preference)

[Tree] (config>service>vprn>bgp>group>neighbor label-preference)

Full Context

configure service vprn bgp label-preference

configure service vprn bgp group label-preference

configure service vprn bgp group neighbor label-preference

Description

This command configures the route preference for routes learned from labeled-unicast peers.

This command can be configured at three levels:

- Global level — applies to all peers
- Group level — applies to all peers in the peer-group
- Neighbor level — applies only to the specified peer

The most specific value is used.

The lower the preference, the higher the chance of the route being the active route.

The **no** form of this command used at the global level reverts to the default *value* of 170.

The **no** form of this command used at the group level reverts to the *value* defined at the global level.

The **no** form of this command used at the neighbor level reverts to the *value* defined at the group level.

Default

no label-preference

Parameters

value

Specifies the route preference value.

Values 1 to 255

Platforms

7705 SAR Gen 2

label-preference

Syntax

label-preference *value*

no label-preference

Context

[Tree] (config>router>bgp label-preference)

[Tree] (config>router>bgp>group label-preference)

[Tree] (config>router>bgp>group>neighbor label-preference)

Full Context

configure router bgp label-preference

configure router bgp group label-preference

configure router bgp group neighbor label-preference

Description

This command configures the route preference for routes learned from labeled-unicast peers.

This command can be configured at three levels:

- Global level — applies to all peers
- Group level — applies to all peers in the peer-group
- Neighbor level — applies only to the specified peer

The most specific value is used.

The lower the preference, the higher the chance of the route being the active route.

The **no** form of this command used at the global level reverts to the default *value* of 170.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no label-preference

Parameters

value

Specifies the route preference value.

Values 1 to 255

Platforms

7705 SAR Gen 2

16.14 label-route-local

label-route-local

Syntax

label-route-local [{**none** | **all**}]

Context

[Tree] (config>router>ttl-propagate label-route-local)

Full Context

configure router ttl-propagate label-route-local

Description

This command configures the TTL propagation for locally generated packets which are forwarded over a BGP label route in the Global Routing Table (GRT) context.

For IPv4 and IPv6 packets forwarded using an RFC 8277 label route in the global routing instance, including 6PE, the **all** value of the command enables TTL propagation from the IP header into all labels in the transport label stack. The none value reverts to the default mode which disables TTL propagation from the IP header to the labels in the transport label stack. This command does not have a no version.

The TTL of the IP packet is always propagated into the RFC 8277 label itself, and this command only controls the propagation into the transport labels, for example, labels of the RSVP or LDP LSP to which the BGP label route resolves and which are pushed on top of the BGP label.

If the BGP peer advertised the implicit-null label value for the BGP label route, the TTL propagation will not follow the configuration described, but will follow the configuration to which the BGP label route resolves:

RSVP LSP shortcut:

- configure router mpls shortcut-local-ttl-propagate

LDP LSP shortcut:

- configure router ldp shortcut-local-ttl-propagate

This feature does not impact packets forwarded over BGP shortcuts. The ingress LER operates in uniform mode by default and can be changed into pipe mode using the configuration of TTL propagation for RSVP or LDP LSP shortcut listed.

Default

label-route-local none

Parameters

none

Specifies that the TTL of the IP packet is not propagated into the transport label stack.

all

Specifies that the TTL of the IP packet is propagated into all labels of the transport label stack.

Platforms

7705 SAR Gen 2

16.15 label-route-transit

label-route-transit

Syntax

label-route-transit [{**none** | **all**}]

Context

[Tree] (config>router>ttl-propagate label-route-transit)

Full Context

configure router ttl-propagate label-route-transit

Description

This command configures the TTL propagation for transit packets which are forwarded over a BGP label route in the Global Routing Table (GRT) context.

For IPv4 and IPv6 packets forwarded using a RFC 8277 label route in the global routing instance, including 6PE, the **all** value of the command enables TTL propagation from the IP header into all labels in the transport label stack. The none value reverts to the default mode which disables TTL propagation from the IP header to the labels in the transport label stack. This command does not have a no version.

The TTL of the IP packet is always propagated into the RFC 8277 label itself, and this command only controls the propagation into the transport labels, for example, labels of the RSVP or LDP LSP to which the BGP label route resolves and which are pushed on top of the BGP label.

If the BGP peer advertised the implicit-null label value for the BGP label route, the TTL propagation will not follow the configuration described, but will follow the configuration to which the BGP label route resolves.

RSVP LSP shortcut:

- configure router mpls shortcut-transit-ttl-propagate

LDP LSP shortcut:

- configure router ldp shortcut-transit-ttl-propagate

This feature does not impact packets forwarded over BGP shortcuts. The ingress LER operates in uniform mode by default and can be changed into pipe mode using the configuration of TTL propagation for the listed RSVP or LDP LSP shortcut.

Default

label-route-transit none

Parameters

none

Specifies that the TTL of the IP packet is not propagated into the transport label stack.

all

Specifies that the TTL of the IP packet is propagated into all labels of the transport label stack.

Platforms

7705 SAR Gen 2

16.16 label-stack-reduction

label-stack-reduction

Syntax

[**no**] **label-stack-reduction**

Context

[Tree] (config>router>mpls>lsp-template label-stack-reduction)

[Tree] (config>router>mpls>lsp label-stack-reduction)

Full Context

configure router mpls lsp-template label-stack-reduction

configure router mpls lsp label-stack-reduction

Description

This command enables the label stack size reduction for a SR-TE LSP or SR-TE LSP template.

At a high level, the label stack reduction algorithm attempts to replace a segment of a computed SR-TE LSP path with the farthest node SID on that path that results in using ECMP paths with links which still comply to the TE constraints of the LSP path.

The **no** form of this command returns the command to its default value.

Default

no label-stack-reduction

Platforms

7705 SAR Gen 2

16.17 label-withdrawal-delay

label-withdrawal-delay

Syntax

label-withdrawal-delay *seconds*

no label-withdrawal-delay

Context

[Tree] (config>router>ldp label-withdrawal-delay)

Full Context

configure router ldp label-withdrawal-delay

Description

This command specifies configures the time interval (in s), LDP will delay for the withdrawal of FEC-label binding it distributed to its neighbors when FEC is de-activated. When the timer expires, LDP then sends a label withdrawal for the FEC to all its neighbors. This is applicable only to LDP IPv4 prefix FECs and is not applicable to pseudowires (service FECs).

When there is an upper layer (user of LDP) which depends on LDP control plane for failover detection then label withdrawal delay and tunnel-down-damp-time options must be set to 0.

An example is PW redundancy where the primary PW doesn’t have its own fast failover detection mechanism and the node depends on LDP tunnel down event to activate the standby PW.

Default

no label-withdrawal-delay

Parameters

seconds

Specifies the time that LDP delays the withdrawal of FEC-label binding it distributed to its neighbors when FEC is de-activated.

Values 3 to 120

Platforms

7705 SAR Gen 2

16.18 labeled-routes

labeled-routes

Syntax

labeled-routes

Context

[Tree] (config>router>bgp>next-hop-res labeled-routes)

Full Context

configure router bgp next-hop-resolution labeled-routes

Description

Commands in this context configure labeled route options for next-hop resolution.

Platforms

7705 SAR Gen 2

16.19 lacp

lacp

Syntax

lacp [**mode**] [**administrative-key** *admin-key*] [**system-id** *system-id*] [**system-priority** *priority*]

no lacp

Context

[Tree] (config>lag lacp)

Full Context

configure lag lacp

Description

This command enables the LACP protocol. Per the IEEE 802.1ax standard, the Link Aggregation Control Protocol (LACP) provides a standardized means for exchanging information between Partner Systems on a link to allow their Link Aggregation Control instances to reach agreement on the identity of the Link Aggregation Group to which the link belongs, move the link to that Link Aggregation Group, and enable its transmission and reception functions in an orderly manner.

If any of the parameters are omitted, the existing configuration is preserved. The default parameter values are used if a parameter is never explicitly configured.

Default

no lacp

Parameters

mode

Specifies the mode in which LACP will operate.

Values **passive** — Starts transmitting LACP packets only after receiving packets.

active — Initiates the transmission of LACP packets.

admin-key

Specifies an administrative key value to identify the channel group on each port configured to use LACP. A random key is assigned by default if a value is not specified when using classic CLI only.

Values 1 to 65535

system-id

Specifies the 48-bit system ID in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff, where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

Values 1 to 65535

Default 32768

priority

Specifies the system priority.

Values 0 to 65535

Default 32768

Platforms

7705 SAR Gen 2

16.20 lacp-fallback

lacp-fallback

Syntax

lacp-fallback mode {**static**} [**timeout** *timeout*]

no lacp-fallback

Context

[Tree] (config>lag lacp-fallback)

Full Context

configure lag lacp-fallback

Description

This command configures the LACP fallback mode. LACP fallback allows one or more designated links of an LACP-controlled LAG to go into forwarding mode if LACP is not yet operational after a configured timeout period. Links capable of forwarding traffic assume the Ethernet and IP characteristics configured for the LAG.

Note: In the saved configuration file, the **lacp-xmit-interval** command must precede the **lacp-fallback** command. As well, to change the **lacp-xmit-interval** from **fast** to **slow**, the **lacp-fallback** command must first be changed to a **timeout** value of 90.

The **no** form of this command disables LACP fallback.

Default

no lacp-fallback

Parameters

mode

Keyword to specify the LACP fallback mode.

Values **static** — Keyword to specify that if LACP is not operational after the timeout period, a single designated LAG member goes into forwarding mode.

individual — The individual mode is not supported.

timeout

Specifies the timeout period in seconds. The LACP link becomes active if no LACP PDUs are received within this timeout period.

Note: Timeout range depends on the configuration of the **configure lag lacp-xmit-limit** command.

Values	fast: 4 to 3600 slow: 90 to 3600
Default	fast: 4

Platforms

7705 SAR Gen 2

16.21 lacp-mux-control

lacp-mux-control

Syntax

lacp-mux-control {**coupled** | **independent**}

no lacp-mux-control

Context

[Tree] (config>lag lacp-mux-control)

Full Context

configure lag lacp-mux-control

Description

This command configures the type of multiplexing machine control to be used in a LAG with LACP in active/passive modes.

The **no** form of this command disables multiplexing machine control.

Default

lacp-mux-control coupled

Parameters

coupled

Specifies that TX and RX activate together.

independent

Specifies that RX activates independent of TX.

Platforms

7705 SAR Gen 2

16.22 lacp-system-priority

lacp-system-priority

Syntax

lacp-system-priority *lacp-system-priority*

no lacp-system-priority

Context

[Tree] (config>system lacp-system-priority)

Full Context

configure system lacp-system-priority

Description

This command configures the Link Aggregation Control Protocol (LACP) system priority on aggregated Ethernet interfaces. LACP allows the operator to aggregate multiple physical interfaces to form one logical interface.

Default

lacp-system-priority 32768

Parameters

lacp-system-priority

Specifies the LACP system priority.

Values 1 to 65535

Platforms

7705 SAR Gen 2

16.23 lacp-tunnel

lacp-tunnel

Syntax

[**no**] **lacp-tunnel**

Context

[Tree] (config>port>ethernet lacp-tunnel)

Full Context

configure port ethernet lacp-tunnel

Description

This command enables LACP packet tunneling for the Ethernet port. When tunneling is enabled, the port does not process any LACP packets but tunnels them instead. The port cannot be added as a member to a LAG group.

In this context, the **lacp-tunnel** command is supported for Epipe and VPLS services only.

The **no** form of this command disables LACP packet tunneling for the Ethernet port.

Default

no lacp-tunnel

Platforms

7705 SAR Gen 2

16.24 lacp-xmit-interval

lacp-xmit-interval

Syntax

lacp-xmit-interval {**slow** | **fast**}

no lacp-xmit-interval

Context

[Tree] (config>lag lacp-xmit-interval)

Full Context

configure lag lacp-xmit-interval

Description

This command specifies the interval signaled to the peer and tells the peer at which rate it should transmit.

Default

lacp-xmit-interval fast

Parameters

slow

Transmits packets every 30 seconds.

fast

Transmits packets every second.

Platforms

7705 SAR Gen 2

16.25 lacp-xmit-stdby

lacp-xmit-stdby

Syntax

[**no**] **lacp-xmit-stdby**

Context

[Tree] (config>lag lacp-xmit-stdby)

Full Context

configure lag lacp-xmit-stdby

Description

This command enables LACP message transmission on standby links.

The **no** form of this command disables LACP message transmission. This command should be disabled for compatibility when using active/standby groups. This forces a timeout of the standby links by the peer. Use the **no** form if the peer does not implement the correct behavior regarding the lacp sync bit.

Default

lacp-xmit-stdby

Platforms

7705 SAR Gen 2

16.26 lag

lag

Syntax

lag *lag-id* **lacp-key** *admin-key* **system-id** *system-id* [**remote-lag** *remote-lag-id*] **system-priority** *system-priority* **source-bmac-lsb use-lacp-key**

lag *lag-id* **lacp-key** *admin-key* **system-id** *system-id* [**remote-lag** *remote-lag-id*] **system-priority** *system-priority* **source-bmac-lsb** *MAC-Lsb*

lag *lag-id* **lacp-key** *admin-key* **system-id** *system-id* [**remote-lag** *remote-lag-id*] **system-priority** *system-priority*

lag *lag-id* [**remote-lag** *remote-lag-id*]

no lag *lag-id*

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-lag lag)

Full Context

configure redundancy multi-chassis peer mc-lag lag

Description

This command defines a LAG which is forming a redundant-pair for MC-LAG with a LAG configured on the given peer. The same LAG group can be defined only in the scope of 1 peer. In order MC-LAG to become operational, all parameters (**lacp-key**, **system-id**, **system-priority**) must be configured the same on both nodes of the same redundant pair.

The partner system (the system connected to all links forming MC-LAG) will consider all ports using the same **lacp-key**, **system-id**, **system-priority** as the part of the same LAG. In order to achieve this in MC operation, both redundant-pair nodes have to be configured with the same values. In case of the mismatch, MC-LAG is kept in oper-down status.

Note that the correct CLI command to enable MC LAG for a LAG in **standby-signaling power-off mode** is **lag** *lag-id* [**remote-lag** *remote-lag-id*]. In the CLI help output, the first three forms are used to enable MC LAG for a LAG in LACP mode. MC LAG is disabled (regardless of the mode) for a given LAG with **no lag** *lag-id.*

Parameters

lag-id

The LAG identifier, expressed as an integer. Specifying the *lag-id* allows the mismatch between lag-id on redundant-pair. If no **lag-id** is specified it is assumed that neighbor system uses the same *lag-id* as a part of the specific MC-LAG. If no matching MC-LAG group can be found between neighbor systems, the individual LAGs operates as usual (no MC-LAG operation is established).

Values 1 to 800

admin-key

Specifies a 16 bit key that needs to be configured in the same manner on both sides of the MC-LAG in order for the MC-LAG to come up.

Values 1 to 65535

system-id

Specifies a 6 byte value expressed in the same notation as MAC address.

Values xx:xx:xx:xx:xx:xx - xx [00 to FF]

remote-lag-id

Specifies the LAG ID on the remote system.

Values 1 to 800

system-priority

Specifies the system priority to be used in the context of the MC-LAG. The partner system will consider all ports using the same **lacp-key**, **system-id**, and **system-priority** as part of the same LAG.

Values 1 to 65535

MAC-Lsb

Configures the last 16 bit of the MAC address to be used for all traffic ingressing the MC-LAG link(s) or if use-lacp-key option is used, it will only copy the value of lacp-key (redundancy multi-chassis mc-lag lag lacp-key admin-key). The command will fail if the *value* is the same with any of the following configured attributes:

- Source-bmac-lsb assigned to other MC-LAG ports.
- lsb 16 bits value for the source-bmac configured at chassis or BVPLS level

  The first 32 bits will be copied from the source B-MAC of the BVPLS associated with the IVPLS for a specific IVPLS SAP mapped to the MC-LAG. The BVPLS source B-MAC can be provisioned for each BVPLS or can be inherited from the chassis PBB configuration.

  Values 1 to 65535 or xx-xx or xx:xx

Platforms

7705 SAR Gen 2

lag

Syntax

lag *lag-id* [**name** *lag-name*]

no lag *lag-id*

Context

[Tree] (config lag)

Full Context

configure lag

Description

Commands in this context configure Link Aggregation Group (LAG) attributes.

A LAG is used to group multiple ports into one logical link. The aggregation of multiple physical links allows for load sharing and offers seamless redundancy. If one link fails, traffic is redistributed over the remaining links.

Note:
For all ports in a LAG group, autonegotiation must be set to "limited" or "off".

There are three possible settings for autonegotiation, as follows:

- "on" or enabled with full port capabilities advertised
- "off" or disabled where there is no autonegotiation advertisements
- "limited" where a single speed/duplex is advertised.

When autonegotiation is enabled on a port, the link attempts to automatically negotiate the link speed and duplex parameters; the configured duplex and speed parameters are ignored.

When autonegotiation is disabled on a port, the port does not attempt to autonegotiate and will only operate at the **speed** and **duplex** settings configured for the port.

Note:
Disabling autonegotiation on gigabit ports is not allowed. This is in accordance with the IEEE 802.3 specification for gigabit Ethernet, which requires gigabyte to be enabled for far end fault indication.

If the **config**>**port**>**ethernet autonegotiate limited** keyword option is specified, the port will autonegotiate but only advertise the **speed** and **duplex** settings configured for the port. Use the **limited** mode on multi-speed gigabit ports to force gigabit operation while keeping autonegotiation is enabled for compliance with IEEE 801.3.

The system requires autonegotiation to be disabled or limited for ports in a LAG to guarantee a specific port speed.

The **no** form of this command deletes the LAG from the configuration. A LAG can only be deleted while the LAG is administratively shut down. Any dependencies, such as IP-Interface configurations, must be removed from the configuration before the **no lag** command is issued.

Parameters

lag-id

Specifies the LAG identifier, expressed as an integer.

The LAG ID ranging from 1 to 64 supports up to 64 LAG members and LAG ID above 64 supports 32 LAG members.

Values 1 to 200

lag-name

Specifies an optional LAG name, up to 27 characters.

In model-driven interfaces, the LAG name is used for configuration references and **show** commands. A service provider or administrator can use the defined LAG name to identify and manage LAGs within the SR OS platforms.

In the classic CLI interface, the user must assign a LAG ID to create the LAG. The LAG name is optional and, if specified, must always start with "lag-". If a name is not specified, SR OS automatically assigns a string version of the LAG ID as "lag-<lag-id>".

Values lag-<23 chars max>

Platforms

7705 SAR Gen 2

lag

Syntax

lag [**lag-id** *lag-id* [**port** *port-id*]] [**all**]

lag [**lag-id** *lag-id* [**port** *port-id*]] [**sm**] [**pkt**] [**cfg**] [**red**] [**iom-upd**] [**port-state**] [**timers**] [**sel-logic**] [**mc**] [**mc-pkt**]

no lag [**lag-id** *lag-id*]

Context

[Tree] (debug lag)

Full Context

debug lag

Description

This command enables debugging for LAG.

Parameters

lag-id

Specifies the link aggregation group ID.

Values 1 to 200

port-id

Specifies the physical port ID.

Values *slot*/*mda*/*port*

all

Specifies to display all LAG information.

sm

Specifies to display trace LACP state machine.

pkt

Specifies to display trace LACP packets.

cfg

Specifies to display trace LAG configuration.

red

Specifies to display trace LAG high availability.

iom-upd

Specifies to display trace LAG IOM updates.

port-state

Specifies to display trace LAG port state transitions.

timers

Specifies to display trace LAG timers.

sel-logic

Specifies to display trace LACP selection logic.

mc

Specifies to display multi-chassis parameters.

mc-packet

Specifies to display the MC-LAG control packets with valid authentication were received on this system.

Platforms

7705 SAR Gen 2

lag

Syntax

lag *lag-id*[**:***encap-val*]

no lag

Context

[Tree] (config>service>vprn>nw-if lag)

Full Context

configure service vprn network-interface lag

Description

This command binds the interface to a Link Aggregation Group (LAG)

The **no** form of this command removes the LAG id from the configuration.

Parameters

lag-id[:encap-val]

Specifies the LAG ID.

Values

lag-id	1 to 800
encap-val	0 (for null)
	0 to 4094 (for dot1q)

Platforms

7705 SAR Gen 2

16.27 lag-port-down

lag-port-down

Syntax

[**no**] **lag-port-down** *lag-id*

Context

[Tree] (config>vrrp>policy>priority-event lag-port-down)

Full Context

configure vrrp policy priority-event lag-port-down

Description

This command creates the context to configure Link Aggregation Group (LAG) priority control events that monitor the operational state of the links in the LAG.

The **lag-port-down** command configures a priority control event. The event monitors the operational state of each port in the specified LAG. When one or more of the ports enter the operational down state, the event is considered to be set. When all the ports enter the operational up state, the event is considered to be clear. As ports enter the operational up state, any previous set threshold that represents more down ports is considered cleared, while the event is considered to be set.

Multiple unique **lag-port-down** event nodes can be configured within the **priority-event** node up to the maximum of 32 events.

The **lag-port-down** command can reference an arbitrary LAG. The *lag-id* does have to already exist within the system. The operational state of the **lag-port-down** event will indicate:

- Set – non-existent
- Set – one port down
- Set – two ports down

- Set – three ports down
- Set – four ports down
- Set – five ports down
- Set – six ports down
- Set – seven ports down
- Set – eight ports down
- Cleared – all ports up

When the *lag-id* is created, or a port in *lag-id* becomes operationally up or down, the event operational state must be updated appropriately.

When one or more of the LAG composite ports enters the operationally down state or the *lag-id* is deleted or does not exist, the event is considered to be set. When an event transitions from clear to set, the set is processed immediately and must be reflected in the associated virtual router instances in-use priority value. As the event transitions from clear to set, a hold-set timer is loaded with the value configured by the events **hold-set** command. This timer prevents the event from clearing until it expires, damping the effect of event flapping. If the event clears and becomes set again before the hold-set timer expires, the timer is reset to the **hold-set** value, extending the time before another clear can take effect.

The **lag-port-down** event is considered to have a tiered event set state. While the priority impact per number of ports down is totally configurable, as more ports go down, the effect on the associated virtual router instances in-use priority is expected to increase (lowering the priority). When each configured threshold is crossed, any higher thresholds are considered further event sets and are processed immediately with the hold-set timer reset to the configured value of the **hold-set** command. As the thresholds are crossed in the opposite direction (fewer ports down then previously), the priority effect of the event is not processed until the hold-set timer expires. If the number of ports down threshold again increases before the hold-set timer expires, the timer is only reset to the **hold-set** value if the number of ports down is equal to or greater than the threshold that set the timer.

The event contains **number-down** nodes that define the priority delta or explicit value to be used based on the number of LAG composite ports that are in the operationally down state. These nodes represent the event set thresholds. Not all port down thresholds must be configured. As the number of down ports increase, the **number-down** *ports-down* node that expresses a value equal to or less than the number of down ports describes the delta or explicit priority value to be applied.

The **no** form of the command deletes the specific LAG monitoring event. The event can be removed at anytime. When the event is removed, the in-use priority of all associated virtual router instances must be reevaluated. The events **hold-set** timer has no effect on the removal procedure.

Default

no lag-port-down — No LAG priority control events are created.

Parameters

lag-id

The LAG ID that the specific event is to monitor expressed as a decimal integer. The *lag-id* can only be monitored by a single event in this policy. The LAG may be monitored by multiple VRRP priority control policies. A port within the LAG and the LAG ID itself are considered to be separate entities. A composite port may be monitored with the **port-down** event while the *lag-id* the port is in is monitored by a **lag-port-down** event in the same policy.

Values 1 to 200

Platforms

7705 SAR Gen 2

16.28 last-member-query-interval

last-member-query-interval

Syntax

last-member-query-interval *tenths-of-seconds*

no last-member-query-interval

Context

[Tree] (config>service>vpls>sap>igmp-snooping last-member-query-interval)

[Tree] (config>service>vpls>mesh-sdp>mld-snooping last-member-query-interval)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping last-member-query-interval)

[Tree] (config>service>vpls>spoke-sdp>mld-snooping last-member-query-interval)

[Tree] (config>service>vpls>sap>mld-snooping last-member-query-interval)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping last-member-query-interval)

Full Context

configure service vpls sap igmp-snooping last-member-query-interval

configure service vpls mesh-sdp mld-snooping last-member-query-interval

configure service vpls spoke-sdp igmp-snooping last-member-query-interval

configure service vpls spoke-sdp mld-snooping last-member-query-interval

configure service vpls sap mld-snooping last-member-query-interval

configure service vpls mesh-sdp igmp-snooping last-member-query-interval

Description

This command configures the maximum response time used in group-specific queries sent in response to 'leave' messages, and is also the amount of time between two consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.

The configured **last-member-query-interval** is ignored when fast leave is enabled on the SAP or SDP.

The **no** form of this command reverts to the default value.

Default

last-member-query-interval 10

Parameters

tenths-of-seconds

Specifies the frequency, in tenths of a second, at which query messages are sent.

Values 1 to 50

Platforms

7705 SAR Gen 2

last-member-query-interval

Syntax

last-member-query-interval *interval*

no last-member-query-interval

Context

[Tree] (config>service>pw-template>igmp-snooping last-member-query-interval)

Full Context

configure service pw-template igmp-snooping last-member-query-interval

Description

This command configures the maximum response time used in group-specific queries sent in response to 'leave' messages, and is also the amount of time between 2 consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.

The configured last-member-query-interval is ignored when fast-leave is enabled on the SAP or SDP.

Default

last-member-query-interval 10

Parameters

interval

Specifies the frequency, in tenths of seconds, at which query messages are sent.

Values 1 to 50

Platforms

7705 SAR Gen 2

16.29 latency-event

latency-event

Syntax

latency-event rising-threshold *threshold* [**falling-threshold** *threshold*] [*direction*]

no latency-event

Context

[Tree] (config>saa>test latency-event)

Full Context

configure saa test latency-event

Description

Specifies that at the termination of an SAA test probe, the calculated latency event value is evaluated against the configured rising and falling latency event thresholds. SAA threshold events are generated as required.

Once the threshold (rising/falling) is crossed, it is disabled from generating additional events until the opposite threshold is crossed. If a **falling-threshold** is not supplied, the **rising- threshold** is re-enabled when it falls below the threshold after the initial crossing that generated the event.

The configuration of latency event thresholds is optional.

The **no** form of this command disables the latency event.

Parameters

rising-threshold *threshold*

Specifies a rising threshold latency value, in milliseconds. When the test run is completed, the calculated latency value is compared to the configured latency rising threshold. If the test run latency value is greater than the configured rising threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Values 0 to 2147483

Default 0

falling-threshold *threshold*

Specifies a falling threshold latency value, in milliseconds. When the test run is completed, the calculated latency value is compared to the configured latency falling threshold. If the test run latency value is greater than the configured falling threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Values 0 to 2147483

Default 0

direction

Specifies the direction for OAM ping responses received for an OAM ping test run.

Values **inbound** — Monitor the value of jitter calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run.

outbound — Monitor the value of jitter calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run.

roundtrip — Monitor the value of jitter calculated for the round trip, two-way, OAM ping requests and replies for an OAM ping test run.

Default roundtrip

Platforms

7705 SAR Gen 2

16.30 ldap

ldap

Syntax

[**no**] **ldap**

Context

[Tree] (config>system>security ldap)

Full Context

configure system security ldap

Description

This command configures LDAP authentication parameters for the system.

The **no** form of this command de-configures the LDAP client from the SR OS.

Platforms

7705 SAR Gen 2

16.31 ldap-server

ldap-server

Syntax

ldap-server *server-name*

no ldap-server

Context

[Tree] (config>system>security>ldap>server ldap-server)

Full Context

configure system security ldap server ldap-server

Description

This command enables the LDAP server name or description.

The **no** form of this command disables the LDAP server name.

Parameters

server-name

Specifies the name of the server, up to 32 characters.

Platforms

7705 SAR Gen 2

16.32 ldp

ldp

Syntax

[**no**] **ldp**

Context

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter ldp)

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter ldp)

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel>resolution-filter ldp)

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter ldp)

Full Context

configure service vprn bgp-evpn mpls auto-bind-tunnel resolution-filter ldp

configure service vpls bgp-evpn mpls auto-bind-tunnel resolution-filter ldp

configure service vprn bgp-ipvpn mpls auto-bind-tunnel resolution-filter ldp

configure service epipe bgp-evpn mpls auto-bind-tunnel resolution-filter ldp

Description

This command enables LDP for the auto-bind tunnel resolution filter.

This command instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next hop.

The **no** form of this command removes the configuration.

Default

no ldp

Platforms

7705 SAR Gen 2

ldp

Syntax

[**no**] **ldp**

Context

[Tree] (config>router ldp)

Full Context

configure router ldp

Description

Commands in this context configure an LDP parameters.

To suspend the LDP protocol, use the **shutdown** command. Configuration parameters are not affected.

The **no** form of the command deletes the LDP protocol instance, removing all associated configuration parameters. The LDP instance must first be disabled with the **shutdown** command before being deleted.

Platforms

7705 SAR Gen 2

ldp

Syntax

[**no**] **ldp**

Context

[Tree] (debug>router ldp)

Full Context

debug router ldp

Description

Use this command to configure LDP debugging.

Platforms

7705 SAR Gen 2

ldp

Syntax

[**no**] **ldp**

Context

[Tree] (config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter ldp)

Full Context

configure router static-route-entry indirect tunnel-next-hop resolution-filter ldp

Description

This command enables the use of LDP-sourced tunnel entries in the TTM to resolve the associated static route next-hop.

The **no** form of this command disables the use of LDP-sourced tunnel entries to resolve static route next hops.

Default

no ldp

Platforms

7705 SAR Gen 2

ldp

Syntax

[**no**] **ldp**

Context

[Tree] (config>service>sdp ldp)

Full Context

configure service sdp ldp

Description

This command enables LDP-signaled LSPs on MPLS-encapsulated SDPs.

In MPLS SDP configurations *either* one or more LSP names can be specified *or* LDP can be enabled. The SDP **ldp** and **lsp** commands are mutually exclusive except if the mixed-lsp-mode option is also enabled. If an LSP is specified on an MPLS SDP, then LDP cannot be enabled on the SDP. To enable LDP on the SDP when an LSP is already specified, the LSP must be removed from the configuration using the **no lsp** *lsp-name* command or the mixed-lsp-mode option is also enabled.

Alternatively, if LDP is already enabled on an MPLS SDP, then an LSP cannot be specified on the SDP. To specify an LSP on the SDP, the LDP must be disabled. The LSP must have already been created in the **config>router>mpls** context with a valid far-end IP address. The above rules are relaxed when the **mixed-lsp** option is enabled on the SDP.

Default

no ldp (disabled)

Platforms

7705 SAR Gen 2

ldp

Syntax

[**no**] **ldp**

Context

[Tree] (config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter ldp)

[Tree] (config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter ldp)

Full Context

configure router bgp next-hop-resolution shortcut-tunnel family resolution-filter ldp

configure router bgp next-hop-resolution labeled-routes transport-tunnel family resolution-filter ldp

Description

This command enables LDP tunneling for next-hop resolution and specifies the LDP tunnels in the tunnel table corresponding to /32 IPv4 FECs and /128 IPv6 FECs.

The **no** form of this command disables LDP tunneling for next-hop resolution.

Platforms

7705 SAR Gen 2

ldp

Syntax

ldp

Context

[Tree] (config>service>vprn>auto-bind-tunnel>res-filter ldp)

Full Context

configure service vprn auto-bind-tunnel resolution-filter ldp

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

16.33 ldp-over-rsvp

ldp-over-rsvp

Syntax

ldp-over-rsvp [**include** | **exclude**]

Context

[Tree] (config>router>mpls>lsp ldp-over-rsvp)

[Tree] (config>router>mpls>lsp-template ldp-over-rsvp)

Full Context

configure router mpls lsp ldp-over-rsvp

configure router mpls lsp-template ldp-over-rsvp

Description

This command configures an LSP so that it can be used by the IGP to calculate its SPF tree.

When the **ldp-over-rsvp** option is also enabled in ISIS or OSPF, the IGP provides LDP with all ECMP IP next-hops and tunnel endpoints that it considers to be the lowest cost path to its destination.

IGP provides only the endpoints which are the closest to the destination in terms of IGP cost for each IP next-hop of a prefix. If this results in more endpoints than the ECMP value configured on the router, it will further prune the endpoints based on the lowest router-id and for the same router-id, it will select lowest interface-index first.

LDP then looks up the tunnel table to select the actual tunnels to the endpoint provided by IGP and further limits the endpoint selection to the ones which are the closest to destination across all the IP next-hops provided by IGP for a prefix. For each remaining endpoint, LDP selects a tunnel in a round-robin fashion until the router ECMP value is reached. For each endpoint, only tunnels with the same lowest metric are candidates. If more than one tunnel qualifies, the selection begins with the lowest tunnel-id.

Default

ldp-over-rsvp include

Platforms

7705 SAR Gen 2

ldp-over-rsvp

Syntax

[**no**] **ldp-over-rsvp**

Context

[Tree] (config>router>isis ldp-over-rsvp)

Full Context

configure router isis ldp-over-rsvp

Description

This command allows LDP over RSVP processing in IS-IS.

The **no** form of this command disables LDP over RSVP processing.

Default

no ldp-over-rsvp

Platforms

7705 SAR Gen 2

ldp-over-rsvp

Syntax

[**no**] **ldp-over-rsvp**

Context

[Tree] (config>router>ospf ldp-over-rsvp)

Full Context

configure router ospf ldp-over-rsvp

Description

This command allows LDP-over-RSVP processing in this OSPF instance.

Default

no ldp-over-rsvp

Platforms

7705 SAR Gen 2

16.34 ldp-shortcut

ldp-shortcut

Syntax

[**no**] **ldp-shortcut**

Context

[Tree] (config>router ldp-shortcut)

Full Context

configure router ldp-shortcut

Description

This command enables the resolution of IGP routes using LDP LSP across all network interfaces participating in the IS-IS and OSPF routing protocol in the system.

When LDP shortcut is enabled, LDP populates the routing table with next-hop entries corresponding to all prefixes for which it activated an LDP FEC. For a given prefix, two route entries are populated in the system routing table. One corresponds to the LDP shortcut next-hop and has an owner of LDP. The other one is the regular IP next-hop. The LDP shortcut next-hop always has preference over the regular IP next-

hop for forwarding user packets and specified control packets over a given outgoing interface to the route next-hop.

All user and specified control packets for which the longest prefix match in RTM yields the FEC prefix will be forwarded over the LDP LSP.

When an IPv4 packet is received on an ingress network interface, a subscriber IES interface, or a regular IES interface, the lookup of the packet by the ingress forwarding engine will result in the packet being sent labeled with the label stack corresponding to the NHLFE of the LDP LSP when the preferred RTM entry corresponds to an LDP shortcut.

If the preferred RTM entry corresponds to an IP next-hop, the IPv4 packet is forwarded without a label.

When ECMP is enabled and multiple equal-cost next-hops exit for the IGP route, the ingress forwarding engine will spray the packets for this route based on hashing routine currently supported for IPv4 packets. When the preferred RTM entry corresponds to an LDP shortcut route, spraying will be performed across the multiple next-hops for the LDP FEC. The FEC next-hops can either be direct link LDP neighbors or T-LDP neighbors reachable over RSVP LSPs in the case of LDP-over-RSVP but not both.

When the preferred RTM entry corresponds to a regular IP route, spraying will be performed across regular IP next-hops for the prefix.

The **no** form of this command disables the resolution of IGP routes using LDP shortcuts.

Default

no ldp-shortcut

Platforms

7705 SAR Gen 2

16.35 ldp-sync

ldp-sync

Syntax

[**no**] **ldp-sync**

Context

[Tree] (config>router>static-route-entry>next-hop ldp-sync)

Full Context

configure router static-route-entry next-hop ldp-sync

Description

This command extends the LDP synchronization feature to a static route. When an interface comes back up, it is possible that a preferred static route using the interface as next-hop for a given prefix is enabled before the LDP adjacency to the peer LSR comes up on this interface. In this case, traffic on an SDP that

uses the static route for the far-end address would be black-holed until the LDP session comes up and the FECs exchanged.

This option when enabled delays the activation of the static route until the LDP session comes up over the interface and the ldp-sync-timer configured on that interface has expired

Default

no ldp-sync

Platforms

7705 SAR Gen 2

16.36 ldp-sync-timer

ldp-sync-timer

Syntax

ldp-sync-timer *seconds* [**end-of-lib**]

no ldp-sync-timer

Context

[Tree] (config>router>if ldp-sync-timer)

Full Context

configure router interface ldp-sync-timer

Description

This command enables synchronization of an IGP and LDP. When a link is restored after a failure, the IGP sets the link cost to infinity and advertises it. The actual value advertised in OSPF is 0xFFFF (65535). The actual value advertised in IS-IS regular metric is 0x3F (63) and in IS-IS wide-metric is 0xFFFFFE (16777214). This feature is not supported on RIP interfaces.

If an interface belongs to both IS-IS and OSPF, a physical failure will cause both IGPs to advertise an infinite metric and to follow the IGP-LDP synchronization procedures. If only one IGP bounces on this interface or on the system, then only the affected IGP advertises the infinite metric and follows the IGP-LDP synchronization procedures.

Next, an LDP Hello adjacency is brought up with the neighbor. The LDP synchronization timer is started by the IGP when the LDP session to the neighbor is up over the interface. This is to allow time for the label-FEC bindings to be exchanged.

When the LDP synchronization timer expires, the link cost is restored and is readvertised. The IGP will announce a new best next hop and LDP will use it if the label binding for the neighbor’s FEC is available.

If the user changes the cost of an interface, the new value is advertised at the next flooding of link attributes by the IGP. However, if the LDP synchronization timer is still running, the new cost value will only

be advertised after the timer expires. The new cost value will also be advertised after the user executes any of the following commands:

- **tools>perform>router>isis>ldp-sync-exit**
- **tools>perform>router>ospf>ldp-sync-exit**
- **config>router>if>no ldp-sync-timer**
- **config>router>ospf>disable-ldp-sync**
- **router>isis>disable-ldp-sync**

If the user changes the value of the LDP synchronization timer parameter, the new value will take effect at the next synchronization event. If the timer is still running, it will continue to use the previous value.

If parallel links exist to the same neighbor, then the bindings and services should remain up as long as there is one interface that is up. However, the user-configured LDP synchronization timer still applies on the interface that failed and was restored. In this case, the router will only consider this interface for forwarding after the IGP re-advertises its actual cost value.

The LDP Sync Timer State is not always synchronized across to the standby CPM. Therefore, after an activity switch, the timer state might not be same as it was on the previously active CPM.

If the **end-of-lib** option is configured, then the system will start the LDP synchronization timer as usual. If the LDP End of LIB Typed Wildcard FEC messages are received for every FEC type negotiated for a given session to an LDP peer for that IGP interface, the **ldp-sync-timer** is terminated early and the IGP link cost is restored. If the **ldp-sync-timer** expires before the LDP End of LIB messages are received for every negotiated FEC type, then the system will restore the IGP link cost. The **end-of-lib** option is disabled by default.

The **no** form of this command disables IGP-LDP synchronization and deletes the configuration.

Default

no ldp-sync-timer

Parameters

seconds

Specifies the time interval for the IGP-LDP synchronization timer.

Values 1 to 1800

end-of-lib

Specifies that the system should terminate the **ldp-sync-timer** early if the LDP End of LIB Typed Wildcard FEC messages are received for every FEC type negotiated for a given session to an LDP peer for that IGP interface.

Platforms

7705 SAR Gen 2

16.37 ldra

ldra

Syntax

ldra

no ldra

Context

[Tree] (config>service>vpls>sap>dhcp6 ldra)

Full Context

configure service vpls sap dhcp6 ldra

Description

This command enables LDRA.

The **no** form of this command disables LDRA.

Default

no ldra

Platforms

7705 SAR Gen 2

16.38 leak

leak

Syntax

leak [*ip-address*]

no leak

Context

[Tree] (debug>router>isis leak)

Full Context

debug router isis leak

Description

This command enables debugging for IS-IS leaks.

The **no** form of the command disables debugging.

Parameters

ip-address

When specified, only the specified address is debugged for IS-IS leaks.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

leak

Syntax

leak [*ip-address*]

no leak

Context

[Tree] (debug>router>ospf3 leak)

[Tree] (debug>router>ospf leak)

Full Context

debug router ospf3 leak

debug router ospf leak

Description

This command enables debugging for OSPF leaks.

Parameters

ip-address

Specifies the IPv4 or IPv6 address to debug OSPF leaks.

Values ipv4-address:

- a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

16.39 leak-destination

leak-destination

Syntax

leak-destination

Context

[Tree] (config>router>static-route-entry leak-destination)

Full Context

configure router static-route-entry leak-destination

Description

Commands in this context configure a list of VPRNs that receive a leaked copy of the static route.

Platforms

7705 SAR Gen 2

16.40 leak-export

leak-export

Syntax

leak-export *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]

no leak-export

Context

[Tree] (config>router leak-export)

Full Context

configure router leak-export

Description

This command associates up to four policies to control the leaking of GRT routes into the associated VPRN.

If a route is evaluated and the action is accepted, that route is subject leaking into an associated VPRN instance, assuming the route is fully resolved and active.

This process creates the pool of routes that can be leaked. Within each VPRN, a corresponding **import-grt** policy must be configured to import select routes into that specific VPRN instance.

The **no** form of this command removes all route leaking policy associations and effectively disables the leaking of GRT routes into associated VPRNs.

Parameters

plcy-or-long-expr

Specifies the route policy name, up to 64 characters or a policy logical expression, up to 255 characters.

Values *plcy-or-long-expr*: *policy-name* | *long-expr*

policy-name: up to 64 characters

long-expr: up to 255 characters

plcy-or-expr

Specifies the route policy name, up to 64 characters or a policy logical expression, up to 64 characters long. A maximum of four policy names or policy logical expressions can be specified in a single statement.

Values *plcy-or-expr*: *policy-name* | expr

policy-name: up to 64 characters

expr: up to 64 characters

Platforms

7705 SAR Gen 2

16.41 leak-export-limit

leak-export-limit

Syntax

[**no**] **leak-export-limit** [*value*]

Context

[Tree] (config>router leak-export-limit)

Full Context

configure router leak-export-limit

Description

This command sets a maximum limit on the number of GRT routes that can be leaked into VPRN instances.

The **no** form of this command resets the **leak-export-limit** to its default value of 5.

Default

leak-export-limit 5

Parameters

value

Specifies the maximum number of eligible GRT routes that can be leaked into VPRN instances.

Values 1 to 10000

Platforms

7705 SAR Gen 2

16.42 leak-import

leak-import

Syntax

leak-import *plcy-or-long-expr* [*plcy-or-expr*]

no leak-import

Context

[Tree] (config>service>vprn>bgp>rib-management>ipv4 leak-import)

[Tree] (config>service>vprn>bgp>rib-management>label-ipv6 leak-import)

[Tree] (config>service>vprn>bgp>rib-management>label-ipv4 leak-import)

[Tree] (config>service>vprn>bgp>rib-management>ipv6 leak-import)

Full Context

configure service vprn bgp rib-management ipv4 leak-import

configure service vprn bgp rib-management label-ipv6 leak-import

configure service vprn bgp rib-management label-ipv4 leak-import

configure service vprn bgp rib-management ipv6 leak-import

Description

This command configures route policies that control the importation of leak-eligible routes from the BGP RIB of another routing instance into the unlabeled-IPv4, unlabeled-IPv6, labeled-IPv4, or labeled-IPv6 RIB of the VPRN instance. To leak a route from one routing instance to another, the origin and destination RIB types must be the same; for example, it is not possible to leak a route from an unlabeled-IPv4 RIB of a VPRN into the labeled-IPv4 RIB of the base router.

The **leak-import** command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine the final action to accept or reject the route.

Only one of the 15 objects referenced by the **leak-import** command can be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters while the remaining 14 objects have a maximum length of 64 characters each.

When a **leak-import** policy is not specified, no BGP routes from other routing instances are leaked into the VPRN BGP RIB.

The **no** form of this command removes the policy association.

Default

no leak-import

Parameters

plcy-or-long-expr

Specifies the route policy name (up to 64 characters) or a policy logical expression (up to 255 characters). Allowed values are any string of characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

plcy-or-expr

Specifies the route policy name (up to 64 characters) or a policy logical expression (up to 255 characters). Allowed values are any string of characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

leak-import

Syntax

leak-import *plcy-or-long-expr* [*plcy-or-expr*]

no leak-import

Context

[Tree] (config>router>bgp>rib-management>ipv6 leak-import)

[Tree] (config>router>bgp>rib-management>label-ipv4 leak-import)

[Tree] (config>router>bgp>rib-management>ipv4 leak-import)

Full Context

configure router bgp rib-management ipv6 leak-import

configure router bgp rib-management label-ipv4 leak-import

configure router bgp rib-management ipv4 leak-import

Description

This command configures the router to specify route policies that control the importation of leak-eligible routes from the BGP RIB of another routing instance into the unlabeled-IPv4, unlabeled-IPv6, or labeled-IPv4 RIB of the base router. To leak a route from one routing instance to another, the origin and destination RIB types must be the same; for example, it is not possible to leak a route from an unlabeled-IPv4 RIB of a VPRN into the labeled-IPv4 RIB of the base router.

The **leak-import** command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine final action to accept or reject the route.

Only one of the 15 objects referenced by the **leak-import** command can be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters while the remaining 14 objects have a maximum length of 64 characters each.

When a **leak-import** policy is not specified, no BGP routes from other routing instances are leaked into the base router BGP RIB.

The **no** form of this command removes the policy association.

Default

no leak-import

Parameters

plcy-or-long-expr

Specifies up to 14 route policy names (up to 64 characters long) or a policy logical expression (up to 255 characters long). Allowed values are any string of characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

plcy-or-expr

The route policy name (up to 64 characters long) or a policy logical expression (up to 64 characters long). Allowed values are any string of characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

16.43 learn-dynamic

learn-dynamic

Syntax

[**no**] **learn-dynamic**

Context

[Tree] (config>service>ies>if>vpls>evpn>arp learn-dynamic)

[Tree] (config>service>ies>if>vpls>evpn>nd learn-dynamic)

[Tree] (config>service>vprn>if>vpls>evpn>nd learn-dynamic)

[Tree] (config>service>vprn>if>vpls>evpn>arp learn-dynamic)

Full Context

configure service ies interface vpls evpn arp learn-dynamic

configure service ies interface vpls evpn nd learn-dynamic

configure service vprn interface vpls evpn nd learn-dynamic

configure service vprn interface vpls evpn arp learn-dynamic

Description

This command controls whether the ARP or ND frames received on EVPN binds are used to learn dynamic ARP and ND entries in the ARP/ND table.

The **no** form of the command reverts to the default.

Default

learn-dynamic

Platforms

7705 SAR Gen 2

16.44 lease-hold-time

lease-hold-time

Syntax

lease-hold-time [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no lease-hold-time

Context

[Tree] (config>router>dhcp>server lease-hold-time)

[Tree] (config>service>vprn>dhcp6>server lease-hold-time)

[Tree] (config>service>vprn>dhcp>server lease-hold-time)

[Tree] (config>router>dhcp6>server lease-hold-time)

Full Context

configure router dhcp local-dhcp-server lease-hold-time

configure service vprn dhcp6 local-dhcp-server lease-hold-time

configure service vprn dhcp local-dhcp-server lease-hold-time

configure router dhcp6 local-dhcp-server lease-hold-time

Description

This command configures the time to remember this lease and is applicable for unsolicited release conditions such as lease timeout if the **lease-hold-time-for** command is set to the default value **no solicited-release** and is additionally applicable for normal solicited releases from DHCP clients if the **lease-hold-time-for** command is set to **solicited-release**.

The **no** form of this command reverts to the default.

Default

lease-hold-time sec 0

Parameters

lease-hold-time

Specifies the amount of time to remember the lease.

Values	*days*	0 to 7305
	hours	0 to 23
	minutes	0 to 59
	seconds	0 to 59

Platforms

7705 SAR Gen 2

16.45 lease-hold-time-for

lease-hold-time-for

Syntax

[**no**] **lease-hold-time-for**

Context

[Tree] (config>router>dhcp6>server lease-hold-time-for)

[Tree] (config>service>vprn>dhcp6>server lease-hold-time-for)

[Tree] (config>service>vprn>dhcp>server lease-hold-time-for)

[Tree] (config>router>dhcp>server lease-hold-time-for)

Full Context

configure router dhcp6 local-dhcp-server lease-hold-time-for

configure service vprn dhcp6 local-dhcp-server lease-hold-time-for

configure service vprn dhcp local-dhcp-server lease-hold-time-for

configure router dhcp local-dhcp-server lease-hold-time-for

Description

Commands in this context configure **lease-hold-time-for** parameters which define additional types of lease or triggers that cause system to hold up leases.

The **no** form of this command reverts to the default.

Platforms

7705 SAR Gen 2

16.46 lease-populate

lease-populate

Syntax

lease-populate [*nbr-of-leases*]

lease-populate [*nbr-of-leases*] **l2-header** [**mac** *ieee-address*]

no lease-populate

Context

[Tree] (config>service>vprn>if>dhcp lease-populate)

[Tree] (config>service>vpls>sap>dhcp lease-populate)

[Tree] (config>service>ies>if>dhcp lease-populate)

Full Context

configure service vprn interface dhcp lease-populate

configure service vpls sap dhcp lease-populate

configure service ies interface dhcp lease-populate

Description

Commands in this context configure IPoE host parameters.

For VPLS, DHCP snooping must be explicitly enabled (using the **snoop** command) at all points where DHCP messages requiring snooping enter the VPLS instance (both from the DHCP server and from the subscribers). Lease state information is extracted from snooped DHCP ACK messages to populate lease state table entries for the SAP.

The optional *nbr-of-leases* parameter defines the number lease state table entries allowed.

- for this SAP in case of a VPLS service
- for this interface in case of an IES or VPRN interface
- for each SAP in case of an IES or VPRN group-interface
- for this interface in case of an IES or VPRN retail subscriber-interface

If the *nbr-of-leases* parameter is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP ACK messages are discarded.

The retained lease state information representing dynamic hosts may be used to:

- Populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering. If the system is unable to populate the dynamic host information in the anti-spoof filter table on the SAP, the DHCP ACK message must be discarded without adding new lease state entry or updating an existing lease state entry.

- Populate the system's ARP cache based on the arp-populate configuration. Applicable to IES and VPRN interfaces or group-interfaces.
- Populate managed entries into a VPLS forwarding database. VPLS forwarding database population is an implicit feature that automatically places the dynamic host's MAC address into the VPLS FDB. When a dynamic host's MAC address is placed in the lease state table, it will automatically be populated into the VPLS forwarding database associated with the SAP on which the host is learned. The dynamic host MAC address will override any static MAC entries using the same MAC and prevent dynamic learning of the MAC on another interface. Existing static MAC entries with the same MAC address as the dynamic host are marked as inactive but not deleted. If all entries in the lease state table associated with the MAC address are removed, the static MAC may be populated. New static MAC definitions for the VPLS instance may be created while a dynamic host exists associated with the static MAC address.
- Generate dynamic ARP replies if **arp-reply-agent** is enabled. Applicable to VPLS service SAPs

The **no** form of this command reverts to the default.

Parameters

nbr-of-leases

Specifies the number of DHCPv4 leases allowed.

l2-header

Indicates a mode of operation where anti-spoof entry associated with the given DHCP state is created based on the *src-mac* address from the Layer 2 header of the DHCP request message. The Layer 2 header flag is not set by default. This parameter is only applicable for group interfaces.

mac

Specifies that the provisioned *ieee-address* is used in the anti-spoofing entries for this SAP. The parameter may be changed mid-session. Existing sessions will not be re-programmed unless a **tools>perform>subscriber-mgmt>remap-lease-state** command is issued for the lease. This parameter is only applicable for group interfaces.

Platforms

7705 SAR Gen 2

lease-populate

Syntax

lease-populate [*nbr-of-leases*]

lease-populate [*nbr-of-leases*] **route-populate** [**pd**] **na** [**ta**]

lease-populate [*nbr-of-leases*] **route-populate pd** [**na**] [**ta**] [**exclude**]

lease-populate [*nbr-of-leases*] **route-populate** [**pd**] [**na**] **ta**

no lease-populate

Context

[Tree] (config>service>ies>if>ipv6>dhcp6-relay lease-populate)

Full Context

configure service ies interface ipv6 dhcp6-relay lease-populate

Description

This command specifies the maximum number of DHCPv6 lease states allocated by the DHCPv6 relay function, allowed on this interface.

Optionally, by specifying **route-populate** parameter, system could:

- Create routes based on the IA_PD/IA_NA/IA_TA prefix option in relay-reply message.
- Create black hole routes based on OPTION_PD_EXCLUDE in IA_PD in relay-reply message.

These routes could be redistributed into IGP/BGP by using route-policy, following protocol types that could be used in "from protocol”:

- dhcpv6-pd
- dhcpv6-na
- dhcpv6-ta
- dhcpv6-pd-excl

Parameters

nbr-of-leases

Defines the number lease state table entries allowed for this interface. If this parameter is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCPv6 REPLY messages are discarded.

Values 1 to 8000

route-populate

Specifies the route populate parameter.

Values pd/na/ta — Create route based on specified option.

exclude — Create blackhole route based on OPTION_PD_EXCLUDE.

Platforms

7705 SAR Gen 2

lease-populate

Syntax

lease-populate [*nbr-of-leases*]

lease-populate [*nbr-of-leases*] **route-populate** [**pd**] **na** [**ta**]

lease-populate [*nbr-of-leases*] **route-populate pd** [**na**] [**ta**] [**exclude**]

lease-populate [*nbr-of-leases*] **route-populate** [**pd**] [**na**] **ta**

no lease-populate

Context

[Tree] (config>service>ies>if>ipv6>dhcp6-relay lease-populate)

Full Context

configure service ies interface ipv6 dhcp6-relay lease-populate

Description

This command specifies the maximum number of DHCPv6 lease states allocated by the DHCPv6 relay function, allowed on this interface.

Optionally, by specifying "route-populate” parameter, system could:

- Create routes based on the IA_PD/IA_NA/IA_TA prefix option in relay-reply message.
- Create black hole routes based on OPTION_PD_EXCLUDE in IA_PD in relay-reply message.

These routes could be redistributed into IGP/BGP by using route-policy, following protocol types that could be used in "from protocol”:

- dhcpv6-pd
- dhcpv6-na
- dhcpv6-ta
- dhcpv6-pd-excl

Parameters

nbr-of-entries

Defines the number lease state table entries allowed for this interface. If this parameter is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCPv6 ACK messages are discarded.

Values 1 to 8000

route-populate

Specifies the route populate parameter.

Values pd/na/ta — Create route based on specified option.

exclude — Create blackhole route based on OPTION_PD_EXCLUDE.

Platforms

7705 SAR Gen 2

16.47 lease-rebind-time

lease-rebind-time

Syntax

lease-rebind-time [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no lease-rebind-time

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options lease-rebind-time)

[Tree] (config>router>dhcp>server>pool>options lease-rebind-time)

[Tree] (config>service>vprn>dhcp>server>pool>options lease-rebind-time)

Full Context

configure subscriber-mgmt local-user-db ipoe host options lease-rebind-time

configure router dhcp local-dhcp-server pool options lease-rebind-time

configure service vprn dhcp local-dhcp-server pool options lease-rebind-time

Description

This command configures the time the client transitions to a rebinding state for a DHCP client.

The **no** form of this command removes the time from the configuration.

Parameters

lease-rebind-time

Specifies the lease rebind time.

Values

days:	0 to 3650
hours:	0 to 23
minutes:	0 to 59
seconds	0 to 59

Platforms

7705 SAR Gen 2

16.48 lease-renew-time

lease-renew-time

Syntax

lease-renew-time [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no lease-renew-time

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options lease-renew-time)

[Tree] (config>service>vprn>dhcp>server>pool>options lease-renew-time)

[Tree] (config>router>dhcp>server>pool>options lease-renew-time)

Full Context

configure subscriber-mgmt local-user-db ipoe host options lease-renew-time

configure service vprn dhcp local-dhcp-server pool options lease-renew-time

configure router dhcp local-dhcp-server pool options lease-renew-time

Description

This command configures the time the client transitions to a renew state for a DHCP client.

The **no** form of this command removes the time from the configuration.

Parameters

lease-renew-time

Specifies the lease renew time.

Values

days:	0 to 3650
hours:	0 to 23
minutes:	0 to 59
seconds	0 to 59

Platforms

7705 SAR Gen 2

16.49 lease-time

lease-time

Syntax

lease-time [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no lease-time

Context

[Tree] (config>router>dhcp>server>pool>options lease-time)

[Tree] (config>service>vprn>dhcp>server>pool>options lease-time)

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options lease-time)

Full Context

configure router dhcp local-dhcp-server pool options lease-time

configure service vprn dhcp local-dhcp-server pool options lease-time

configure subscriber-mgmt local-user-db ipoe host options lease-time

Description

This command configures the amount of time that the DHCP server grants to the DHCP client permission to use a specific IP address.

The **no** form of this command removes the lease time parameters from the configuration.

Parameters

days

Specifies the number of days that the given IP address is valid.

Values 0 to 3650

hours

Specifies the number of hours that the given IP address is valid.

Values 0 to 23

minutes

Specifies the number of minutes that the given IP address is valid.

Values 0 to 59

seconds

Specifies the number of seconds that the given IP address is valid.

Values 0 to 59

Platforms

7705 SAR Gen 2

lease-time

Syntax

lease-time [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*] [**override**]

no lease-time

Context

[Tree] (config>service>vprn>if>dhcp>proxy lease-time)

[Tree] (config>service>vpls>sap>dhcp>proxy-server lease-time)

[Tree] (config>service>ies>if>dhcp>proxy-server lease-time)

Full Context

configure service vprn interface dhcp proxy-server lease-time

configure service vpls sap dhcp proxy-server lease-time

configure service ies interface dhcp proxy-server lease-time

Description

This command defines the length of lease-time that is provided to DHCP clients. By default, the local-proxy-server always makes use of the lease time information provide by either a RADIUS or DHCP server.

The **no** form of this command disables the use of the lease-time command. The local-proxy-server will use the lease-time offered by either a RADIUS or DHCP server.

Default

lease-time days 7

Parameters

override

Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients

days

Specifies the number of days that the given IP address is valid.

Values 0 to 3650

hours

Specifies the number of hours that the given IP address is valid.

Values 0 to 23

minutes

Specifies the number of minutes that the given IP address is valid.

Values 0 to 59

seconds

Specifies the number of seconds that the given IP address is valid.

Values 0 to 59

Platforms

7705 SAR Gen 2

16.50 least-fill

least-fill

Syntax

[**no**] **least-fill**

Context

[Tree] (config>router>mpls>lsp-template least-fill)

[Tree] (config>router>mpls>lsp least-fill)

Full Context

configure router mpls lsp-template least-fill

configure router mpls lsp least-fill

Description

This command enables the use of the least-fill path selection method for the computation of the path of this LSP.

When MPLS requests the computation of a path for this LSP, CSPF will find all equal cost shortest paths which satisfy the constraints of this path. Then, CSPF identifies the single link in each of these paths which has the least available bandwidth as a percentage of its maximum reservable bandwidth. It then selects the path which has the largest value of this percentage least available bandwidth figure. CSPF identifies the least available bandwidth link in each equal cost path after it has accounted for the bandwidth of the new requested path of this LSP.

CSPF applies the least-fill path selection method to all requests for a path, primary and secondary, of an LSP for which this option is enabled. The bandwidth of the path can be any value, including zero.

CSPF applies the least-fill criterion separately to each preemption priority in the base TE. A higher setup priority path can preemptively lower holding priority paths.

CSPF also applies the least-fill criterion separately to each Diff-Serv TE class if Diff-Serv TE is enabled on this node. A higher setup priority path can preemptively lower holding priority paths within a Class Type.

MPLS will re-signal and move the LSP to the new path in the following cases:

- Initial LSP path signaling.
- Re-try of an LSP path after failure.
- Make-before-break (MBB) due to pending soft preemption of the LSP path.
- MBB due to LSP path configuration change, that is, a user change to bandwidth parameter of primary or secondary path, or a user enabling of fast-reroute option for the LSP.
- MBB of secondary path due to an update to primary path SRLG.
- MBB due to FRR Global Revertive procedures on the primary path.
- Manual re-signaling of an LSP path or of all LSP paths by the user.

During a manual re-signaling of an LSP path, MPLS will always re-signal the path regardless of whether the new path is exactly the same or different than the current path and regardless of whether the metric of the new path is different or not from that of the current path.

During a timer-based re-signaling of an LSP path which has the least-fill option enabled, MPLS will only re-signal the path if the metric of the new path is different than the one of the current path.

The **no** form of this command deletes a specific node entry in this database.

Default

no least-fill. The path of an LSP is randomly chosen among a set of equal cost paths.

Platforms

7705 SAR Gen 2

16.51 least-fill-min-thd

least-fill-min-thd

Syntax

least-fill-min-thd *percent*

no least-fill-min-thd

Context

[Tree] (config>router>mpls least-fill-min-thd)

Full Context

configure router mpls least-fill-min-thd

Description

This parameter is used in the least-fill path selection process. When comparing the percentage of least available link bandwidth across the sorted paths, whenever two percentages differ by less than the value

configured as the least-fill-min-thresh, CSPF will consider them equal and will apply a random number generator to select the path among these paths

The **no** form of this command resets this parameter to its default value.

Default

least-fill-min-thd 5

Parameters

percentage

Specifies the least fill minimum threshold value as a percentage.

Values 1 to 100%

Platforms

7705 SAR Gen 2

16.52 least-fill-reoptim-thd

least-fill-reoptim-thd

Syntax

least-fill-reoptim-thd *percent*

no least-fill-reoptim-thd

Context

[Tree] (config>router>mpls least-fill-reoptim-thd)

Full Context

configure router mpls least-fill-reoptim-thd

Description

This parameter is used in the least-fill path selection method. During a timer-based re-signaling of an LSP path which has the least-fill option enabled, CSPF will first update the least-available bandwidth figure for the current path of this LSP. It then applies the least-fill path selection method to select a new path for this LSP. If the new computed path has the same cost as the current path, it will compare the least-available bandwidth figures of the two paths and if the difference exceeds the user configured optimization threshold, MPLS will generate a trap to indicate that a better least-fill path is available for this LSP. This trap can be used by an external SNMP based device to trigger a manual re-signaling of the LSP path since the timer-based re-signaling will not re-signal the path in this case. MPLS will generate a path update trap at the first MBB event which results in the re-signaling of the LSP path. This should clear the eligibility status of the path at the SNMP device.

The **no** form of this command resets this parameter to its default value.

Default

least-fill-reoptim-thd 10

Parameters

percentage

Specifies the least fill reoptimization threshold value as a percentage.

Values 1 to 100%

Platforms

7705 SAR Gen 2

16.53 legacy

legacy

Syntax

[**no**] **legacy**

Context

[Tree] (config>router>isis>te>application-link-attributes legacy)

Full Context

configure router isis traffic-engineering-options application-link-attributes legacy

Description

This command enables legacy mode of advertising TE attributes.

The **no** form of this command disables legacy mode, but enables the per-application TE attribute advertisement for RSVP-TE.

Default

legacy

Platforms

7705 SAR Gen 2

16.54 legacy-ipv4-lsr-interop

legacy-ipv4-lsr-interop

Syntax

[**no**] **legacy-ipv4-lsr-interop**

Context

[Tree] (config>router>ldp legacy-ipv4-lsr-interop)

Full Context

configure router ldp legacy-ipv4-lsr-interop

Description

This command provides for a global LDP knob to allow interoperability with legacy IPv4 LSR implementations which do not comply with the processing of Hello TLVs with the U-bit set. Specifically, this feature disables the following Hello TLVs:

- The Nokia proprietary Interface Info TLV (0x3E05) in the Hello message sent to the peer. This also results in the non-generation of the Nokia proprietary Hello Adjacency Status TLV (0x3E06) since the Interface Info TLV is not sent.

  This is performed in SR OS releases 12 and higher.

- The RFC 7552 standard dual-stack capability TLV (0x701) and the Nokia proprietary Adjacency capability TLV (0x3E07) in SR OS releases 13 and higher.

Platforms

7705 SAR Gen 2

16.55 legacy-mode

legacy-mode

Syntax

[**no**] **legacy-mode**

Context

[Tree] (config>router>bgp>error-handling legacy-mode)

[Tree] (config>service>vprn>bgp>error-handling legacy-mode)

Full Context

configure router bgp error-handling legacy-mode

configure service vprn bgp error-handling legacy-mode

Description

This command configures the BGP instance to handle BGP update error messages based on the configured **update-fault-tolerance** commands.

Note: If the **update-fault-tolerance** commands are not explicitly configured, BGP error handling follows the legacy procedures described in RFC 4271, which can result in disruptive session resets.

The **no** form of this command configures the BGP instance to ignore the configured **update-fault-tolerance** commands and apply the new error-handling procedures described in RFC 7606 on all sessions.

Default

no legacy-mode

Platforms

7705 SAR Gen 2

16.56 length

length

Syntax

length *lines*

Context

[Tree] (environment>terminal length)

Full Context

environment terminal length

Description

This command sets the number of lines on a screen.

Parameters

lines

Specifies the number of lines for the terminal screen length, expressed as a decimal integer.

Values 1 to 512

Default 24 — terminal dimensions are set to 24 lines long by 80 characters wide

Platforms
7705 SAR Gen 2

length

Syntax
length *lines*

Context
[Tree] (config>system>management-interface>cli>md-cli>environment>console length)

Full Context
configure system management-interface cli md-cli environment console length

Description
This command configures the set number of lines displayed on the console.

Default
length 24

Parameters
lines
Specifies the number of lines displayed in the console window.

Values 24 to 512

Platforms
7705 SAR Gen 2

16.57 length-field

length-field

Syntax
[**no**] **length-field**

Context

[Tree] (config>test-oam>icmp>ipv6 length-field)

Full Context

configure test-oam icmp ipv6 length-field

Description

This command enables the setting of the length field when building an RFC 4884, *Extended ICMP to Support Multi-Part Messages*, *ICMPv6 Destination Unreachable* message or *ICMPv6 Time Exceeded* message.

The **no** form of this command disables the length field modification.

Default

no length-field

Platforms

7705 SAR Gen 2

16.58 ler-use-dscp

ler-use-dscp

Syntax

[**no**] **ler-use-dscp**

Context

[Tree] (config>qos>network>ingress ler-use-dscp)

Full Context

configure qos network ingress ler-use-dscp

Description

This command is used to enable tunnel QoS mapping on all ingress network IP interfaces that the network-qos-policy-id is associated with. The command may be defined at any time after the network QoS policy has been created. Any network IP interfaces currently associated with the policy will immediately start to use the internal IP ToS field of any tunnel terminated IP routed packet received on the interface, ignoring any QoS markings in the tunnel portion of the packet.

This attribute provides the ability to ignore the network ingress QoS mapping of a terminated tunnel containing an IP packet that is to be routed to a base router or VPRN destination. This is advantageous when the mapping for the tunnel QoS marking does not accurately or completely reflect the required QoS handling for the IP routed packet. When the mechanism is enabled on an ingress network IP interface, the IP interface will ignore the tunnel's QoS mapping and derive the internal forwarding class and profile based

on the precedence or DiffServ Code Point (DSCP) values within the routed IP header ToS field compared to the Network QoS policy defined on the IP interface.

The default state is not to enforce tunnel termination IP routed QoS override within the network QoS policy.

The **no** form of this command removes tunnel termination IP routed QoS override from the network QoS policy and all ingress network IP interfaces associated with the policy.

Default

no ler-use-dscp

Platforms

7705 SAR Gen 2

16.59 less-specific

less-specific

Syntax

less-specific [**allow-default**]

no less-specific

Context

[Tree] (config>vrrp>policy>priority-event>route-unknown less-specific)

Full Context

configure vrrp policy priority-event route-unknown less-specific

Description

This command allows a CIDR shortest match hit on a route prefix that contains the IP route prefix associated with the route unknown priority event.

The **less-specific** command modifies the search parameters for the IP route prefix specified in the **route-unknown** priority event. Specifying **less-specific** allows a CIDR shortest match hit on a route prefix that contains the IP route prefix.

The **less-specific** command eases the RTM lookup criteria when searching for the *prefix/mask-length*. When the **route-unknown** priority event sends the prefix to the RTM (as if it was a destination lookup), the result route table prefix (if a result is found) is checked to see if it is an exact match or a less specific match. The **less-specific** command enables a less specific route table prefix to match the configured prefix. When **less-specific** is not specified, a less specific route table prefix fails to match the configured prefix. The **allow-default** optional parameter extends the **less-specific** match to include the default route (0.0.0.0).

The **no** form of the command prevents RTM lookup results that are less specific than the route prefix from matching.

Default

no less-specific — The route unknown priority events requires an exact prefix/mask match.

Parameters

allow-default

When the **allow-default** parameter is specified with the **less-specific** command, an RTM return of 0.0.0.0 matches the IP prefix. If **less-specific** is entered without the **allow-default** parameter, a return of 0.0.0.0 will not match the IP prefix. To disable **allow-default**, but continue to allow **less-specific** match operation, only enter the **less-specific** command (without the **allow-default** parameter).

Platforms

7705 SAR Gen 2

16.60 level

level

Syntax

level *priority-level* **rate** *pir-rate* [**cir** *cir-rate*]

level *priority-level* **percent-rate** *pir-percent* [**percent-cir** *cir-percent*]

no level *priority-level*

Context

[Tree] (config>port>ethernet>egr-scheduler-override level)

Full Context

configure port ethernet egress-scheduler-override level

Description

This command overrides the maximum and CIR rate parameters for a specific priority level on the port or channel's port scheduler instance. When the **level** command is executed for a priority level, the corresponding priority level command in the port-scheduler-policy associated with the port is ignored.

The override level command supports the keyword **max** for the **rate** and **cir** parameter. When executing the level override command, at least the **rate** or **cir** keywords and associated parameters must be specified for the command to succeed.

The **no** form of this command removes the local port priority level rate overrides. Once removed, the port priority level will use the port scheduler policies level command for that priority level.

Parameters

priority-level

Identifies which of the eight port priority levels are being overridden.

Values 1 to 8

pir-rate

Overrides the port scheduler policy's maximum level rate and requires either the **max** keyword or a rate defined in kilobits per second to follow.

Values For Ethernet: 1 to 6400000000, **max**

For SONET-SDH and TDM: 1 to 3200000000, **max**

cir-rate

Overrides the port scheduler policy's within-cir level rate and requires either the max keyword or a rate defined in kilobits per second to follow.

Values For Ethernet: 1 to 6400000000, **max**

For SONET-SDH and TDM: 1 to 3200000000, **max**

pir-percent

Specifies the PIR as a percentage.

Values 0.01 to 100.00

cir-percent

Specifies the CIR as a percentage.

Values 0.00 to 100.00

max

removes any existing rate limit imposed by the port scheduler policy for the priority level allowing it to use as much total bandwidth as possible.

Platforms

7705 SAR Gen 2

level

Syntax

level *level-number*

Context

[Tree] (config>service>vprn>isis>if level)

[Tree] (config>service>vprn>isis>link-group level)

[Tree] (config>service>vprn>isis level)

Full Context

configure service vprn isis interface level

configure service vprn isis link-group level

configure service vprn isis level

Description

This command creates the context to configure IS-IS Level 1 or Level 2 area attributes.

A router can be configured as a Level 1, Level 2, or Level 1/2 system. A Level 1 adjacency can be established if there is at least one area address shared by this router and a neighbor. A Level 2 adjacency cannot be established over this interface.

Level 1/2 adjacency is created if the neighbor is also configured as Level 1/2 router and has at least one area address in common. A Level 2 adjacency is established if there are no common area IDs.

A Level 2 adjacency is established if another router is configured as Level 2 or a Level 1/2 router with interfaces configured as Level 1/2 or Level 2. Level 1 adjacencies are not established over this interface.

To reset global and/or interface level parameters to the default, the following commands must be entered independently:

- **level>no hello-authentication-key**
- **level>no hello-authentication-type**
- **level>no hello-interval**
- **level>no hello-multiplier**
- **level>no metric**
- **level>no passive**
- **level>no priority**

Default

level 1 or level 2

Parameters

level-number

The IS-IS level number.

Values 1, 2

Platforms

7705 SAR Gen 2

level

Syntax

level *syslog-level*

Context

[Tree] (config>service>vprn>log>syslog level)

Full Context

configure service vprn log syslog level

Description

This command configures the syslog message severity level threshold. All messages with severity level equal to or higher than the threshold are sent to the syslog target host.

Only a single threshold level can be specified. If multiple levels are entered, the last **level** entered will overwrite the previously entered commands.

Default

level info

Parameters

syslog-level

The threshold severity level name.

Values emergency, alert, critical, error, warning, notice, info, debug

Router severity level	Numerical Severity (highest to lowest)	Configured Severity	Definition
	0	emergency	system is unusable
3	1	alert	action must be taken immediately
4	2	critical	critical condition
5	3	error	error condition
6	4	warning	warning condition
	5	notice	normal but significant condition
1 cleared 2 indeterminate	6	info	informational messages
	7	debug	debug-level messages

Platforms

7705 SAR Gen 2

level

Syntax

level *priority-level* **rate** *pir-rate* [**cir** *cir-rate*] **group** *name* [**weight** *weight*] [**monitor-threshold** *percent*]

level *priority-level* **percent-rate** *pir-percent* [**percent-cir** *cir-percent*] **group** *name* [**weight** *weight*] [**monitor-threshold** *percent*]

level *priority-level* **rate** *pir-rate* [**cir** *cir-rate*] [**monitor-threshold** *percent*]

level *priority-level* **percent-rate** *pir-percent* [**percent-cir** *cir-percent*] [**monitor-threshold** *percent*]

no level *priority-level*

Context

[Tree] (config>qos>port-scheduler-policy level)

Full Context

configure qos port-scheduler-policy level

Description

This command configures an explicit within-CIR bandwidth limit and a total bandwidth limit for each port scheduler's priority level. To understand how to set the level rate and CIR parameters, a basic understanding of the port-level scheduler bandwidth allocation mechanism is required. The port scheduler takes all available bandwidth for the port or channel (after the max-rate and any port egress-rate limits have been accounted for) and offers it to each of the eight priority levels twice.

The first pass is called the within-CIR pass and consists of providing the available port bandwidth to each of the 8 priority levels, starting with level 8 and moving down to level 1. Each level takes the offered load and distributes it to all child members that have a port-parent cir-level equal to the current priority level. (Any child with a cir-weight equal to 0 is skipped in this pass.) Each child may consume bandwidth up to the child's frame-based within-CIR offered load. The remaining available port bandwidth is then offered to the next lower priority level until level 1 is reached.

The second pass is called the above-CIR pass and consists of providing the remaining available port bandwidth to each of the eight priority levels a second time. Again, each level takes the offered load and distributes it to all child members that have a port-parent level equal to the current priority level. Each child may consume bandwidth up to the remainder of the child's frame-based offered load (some of the offered load may have been serviced during the within-CIR pass). The remaining available port bandwidth is then offered to the next priority level until level 1 is again reached.

If the port scheduling policy is using the default orphan behavior (orphan-override has not been configured on the policy), the system then takes any remaining port bandwidth and allocates it to the orphan queues and scheduler on priority level 1. In a non-override orphan state, all orphans are attached to priority level 1 using a weight of 0. The zero weight value causes the system to allocate bandwidth equally to all orphans based on each orphan queue or scheduler's ability to use the bandwidth. If the policy has an orphan-override configured, the orphans are handled based on the override commands parameters in a similar fashion to properly parented queues and schedulers.

The port scheduler priority level command rate keyword is used to optionally limit the total amount of bandwidth that is allocated to a priority level (total for the within-CIR and above-CIR passes). The cir keyword optionally limits the first pass bandwidth allocated to the priority level during the within-CIR pass.

When executing the level command, at least one of the optional keywords, **rate** or **cir,** must be specified. If neither keyword is included, the command will fail.

If a previous explicit value for rate or cir exists when the level command is executed, and either rate or cir is omitted, the previous value for the parameter is overwritten by the default value and the previous value is lost.

The configured priority level rate limits may be overridden at the egress port or channel using the egress-scheduler-override level priority-level command. When a scheduler instance has an override defined for a priority level, both the rate and cir values are overridden even when one of them is not explicitly expressed in the override command. For instance, if the cir kilobits per second portion of the override is not expressed, the scheduler instance defaults to not having a CIR rate limit for the priority level even when the port scheduler policy has an explicit CIR limit defined.

The **no** form of this command returns the level to its default value.

Default

no level priority-level

Parameters

priority-level

Specifies to which priority level the level command pertains. Each of the eight levels is represented by an integer value of 1 to 8, with 8 being the highest priority level.

Values 1 to 8 (8 is the highest priority)

pir-rate

Specifies the total bandwidth limits allocated to priority-level, in kilobits per second.

Values 1 to 6400000000, **max**

pir-percent

Specifies the percent bandwidth limits allocated to priority-level.

Values 0.01 to 100.00

cir-rate

The cir specified limits the total bandwidth allocated in the within-CIR distribution pass to priority-level. When cir is not specified, all the available port or channel bandwidth may be allocated to the specified priority level during the within-CIR pass.

Values 0 to 6400000000, **max**

The value given for kilobits per second is expressed in kilobits per second on a base 10 scale as is usual for line rate calculations. If a value of 1 is given, the result is 1000 bits per second (as opposed to a base 2 interpretation that would be 1024 bits per second).

cir-percent

Specifies the percent bandwidth limits allocated to priority-level.

Values 0.00 to 100.00

group ***name***

specifies the existing group that the weighted scheduler group this level maps to, up to 32 characters.

weight

Specifies the weight of the level within this weighted scheduler group.

Values 1 to 100

Default 1

monitor-threshold ***percent***

Specifies the percent of the configured rate. If the offered rate exceeds the configured threshold, a counter monitoring the threshold will be increased.

Values 0 to 100

Platforms

7705 SAR Gen 2

level

Syntax

level *syslog-level*

no level

Context

[Tree] (config>log>syslog level)

Full Context

configure log syslog level

Description

This command configures the syslog message severity level threshold. All messages with severity level equal to or higher than the threshold are sent to the syslog target host.

Only a single threshold level can be specified. If multiple levels are entered, the last **level** entered will overwrite the previously entered commands.

The **no** form of this command reverts to the default value.

Default

level info

Parameters

value

Specifies the threshold severity level name.

Values emergency, alert, critical, error, warning, notice, info, debug

Table 47: Level Parameter Value Descriptions

Router severity level	Numerical Severity (highest to lowest)	Configured Severity	Definition
	0	emergency	system is unusable
3	1	alert	action must be taken immediately
4	2	critical	critical condition
5	3	error	error condition
6	4	warning	warning condition
	5	notice	normal but significant condition
1 cleared 2 indeterminate	6	info	informational messages
	7	debug	debug-level messages

Platforms

7705 SAR Gen 2

level

Syntax

level {**1** | **2**}

Context

[Tree] (config>router>isis level)

[Tree] (config>router>isis>interface level)

Full Context

configure router isis level

configure router isis interface level

Description

This command creates the context to configure IS-IS Level 1 or Level 2 area attributes.

A router can be configured as a Level 1, Level 2, or Level 1/2 system. A Level 1 adjacency can be established if there is at least one area address shared by this router and a neighbor. A Level 2 adjacency cannot be established over this interface.

Level 1/2 adjacency is created if the neighbor is also configured as Level 1/2 router and has at least one area address in common. A Level 2 adjacency is established if there are no common area IDs.

A Level 2 adjacency is established if another router is configured as Level 2 or a Level 1/2 router with interfaces configured as Level 1/2 or Level 2. Level 1 adjacencies are not established over this interface.

To reset global and/or interface level parameters to the default, the following commands must be entered independently:

```
– level>no hello-authentication-key
    – level>no hello-authentication-type
    – level>no hello-interval
    – level>no hello-multiplier
    – level>no metric
    – level>no passive
    – level>no priority
```

Default

level 1 or level 2

Parameters

1

Specifies the IS-IS operational characteristics of the interface at level 1.

2

Specifies the IS-IS operational characteristics of the interface at level 2.

Platforms

7705 SAR Gen 2

level

Syntax

level {**1** | **2**}

no level

Context

[Tree] (config>router>policy-options>policy-statement>entry>from level)

[Tree] (config>router>policy-options>policy-statement>entry>to level)

Full Context

configure router policy-options policy-statement entry from level

configure router policy-options policy-statement entry to level

Description

This command specifies the ISIS route level as a match criterion for the entry.

Default

no level

Parameters

1 | 2

Matches the IS-IS route learned from level 1 or level 2.

Platforms

7705 SAR Gen 2

16.61 level-capability

level-capability

Syntax

level-capability {**level-1** | **level-2** | **level-1/2**}

no level-capability

Context

[Tree] (config>service>vprn>isis level-capability)

[Tree] (config>service>vprn>isis>if level-capability)

Full Context

configure service vprn isis level-capability

configure service vprn isis interface level-capability

Description

This command configures the routing level for an instance of the IS-IS routing process.

An IS-IS router and an IS-IS interface can operate at Level 1, Level 2 or both Level 1 *and* 2.

Table 48: Potential Adjacency Capabilities displays configuration combinations and the potential adjacencies that can be formed.

Table 48: Potential Adjacency Capabilities

Global Level	Interface Level	Potential Adjacency
L 1/2	L 1/2	Level 1 and/or Level 2
L 1/2	L 1	Level 1 only
L 1/2	L 2	Level 2 only

Global Level	Interface Level	Potential Adjacency
L 2	L 1/2	Level 2 only
L 2	L 2	Level 2 only
L 2	L 1	none
L 1	L 1/2	Level 1 only
L 1	L 2	none
L 1	L 1	Level 1 only

The **no** form of this command removes the level capability from the configuration.

Default

level-capability level-1/2

Parameters

level-1

Specifies the router/interface can operate at Level 1 only.

level-2

Specifies the router/interface can operate at Level 2 only.

level-1/2

Specifies the router/interface can operate at both Level 1 and Level 2.

Platforms

7705 SAR Gen 2

level-capability

Syntax

level-capability {**level-1** | **level-2** | **level-1/2**}

no level-capability

Context

[Tree] (config>router>isis level-capability)

[Tree] (config>router>isis>interface level-capability)

Full Context

configure router isis level-capability

configure router isis interface level-capability

Description

This command configures the routing level for an instance of the IS-IS routing process.

An IS-IS router and an IS-IS interface can operate at Level 1, Level 2 or both Level 1 and 2.

Table 49: Potential Adjacency displays configuration combinations and the potential adjacencies that can be formed.

Table 49: Potential Adjacency

Global Level	Interface Level	Potential Adjacency
L 1/2	L 1/2	Level 1 and/or Level 2
L 1/2	L 1	Level 1 only
L 1/2	L 2	Level 2 only
L 2	L 1/2	Level 2 only
L 2	L 2	Level 2 only
L 2	L 1	—
L 1	L 1/2	Level 1 only
L 1	L 2	—
L 1	L 1	Level 1 only

The **no** form of this command removes the level capability from the configuration.

Default

level-capability level-1/2

Parameters

level-1

Specifies the router/interface can operate at Level 1only.

level-2

Specifies the router/interface can operate at Level 2 only.

level-1/2

Specifies the router/interface can operate at both Level 1 and Level 2.

Platforms

7705 SAR Gen 2

16.62 lfa-policy-map

lfa-policy-map

Syntax

lfa-policy-map route-nh-template *template-name*

no lfa-policy-map

Context

[Tree] (config>service>vprn>isis>if lfa-policy-map)

Full Context

configure service vprn isis interface lfa-policy-map

Description

This command applies a route next-hop policy template to the IS-IS interface for the VPRN instance.

When a route next-hop policy template is applied to an interface in IS-IS, it is applied in both level 1 and level 2. When a route next-hop policy template is applied to an interface in OSPF, it is applied in all areas. However, the command in an OSPF interface context can only be executed under the area in which the specified interface is primary and then applied in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command fails.

If the user excluded the interface from LFA using the command **loopfree-alternate-exclude**, the LFA policy, if applied to the interface, has no effect.

Finally, if the user applied a route next-hop policy template to a loopback interface or to the system interface, the command will not be rejected, but it will result in no action being taken.

The **no** form deletes the mapping of a route next-hop policy template to an OSPF or IS-IS interface.

Parameters

template-name

Specifies the name of the template, up to 32 characters.

Platforms

7705 SAR Gen 2

lfa-policy-map

Syntax

lfa-policy-map route-nh-template *template-name*

no lfa-policy-map

Context

[Tree] (config>service>vprn>ospf>area>if lfa-policy-map)

[Tree] (config>service>vprn>ospf3>area>if lfa-policy-map)

[Tree] (config>router>ospf3>area>if lfa-policy-map)

[Tree] (config>router>isis>if lfa-policy-map)

[Tree] (config>router>ospf>area>if lfa-policy-map)

Full Context

configure service vprn ospf area interface lfa-policy-map

configure service vprn ospf3 area interface lfa-policy-map

configure router ospf3 area interface lfa-policy-map

configure router isis interface lfa-policy-map

configure router ospf area interface lfa-policy-map

Description

This command applies a route next-hop policy template to an OSPF or IS-IS interface.

When a route next-hop policy template is applied to an interface in IS-IS, it is applied in both level 1 and level 2. When a route next-hop policy template is applied to an interface in OSPF, it is applied in all areas. However, the command in an OSPF interface context can only be executed under the area in which the specified interface is primary and then applied in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command fails.

If the user excluded the interface from LFA using the command **loopfree-alternate-exclude**, the LFA policy, if applied to the interface, has no effect.

Finally, if the user applied a route next-hop policy template to a loopback interface or to the system interface, the command will not be rejected, but it results in no action being taken.

The **no** form deletes the mapping of a route next-hop policy template to an OSPF or IS-IS interface.

Default

no lfa-policy-map

Parameters

template-name

Specifies the name of the template, up to 32 characters.

Platforms

7705 SAR Gen 2

16.63 license

license

Syntax

license

Context

[Tree] (admin>system license)

Full Context

admin system license

Description

Enters a context for administrative commands related to licensing.

Platforms

7705 SAR Gen 2

16.64 license-file

license-file

Syntax

license-file *file-url*

no license-file

Context

[Tree] (bof license-file)

Full Context

bof license-file

Description

This command configures the license location and file name.

The **no** form of this command removes the file URL from the configuration.

Parameters

file-url

Specifies the *file-url*.

Values

file-url	{*local-url* | *remote-url*} (up to 180 characters)
local-url	[*cflash-id*/][*file-path*]
remote-url	[{ftp://| tftp://} *login:pswd@remote-locn*/][*file-path*]
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

16.65 lifetime

lifetime

Syntax

lifetime {*seconds* | **forever**}

Context

[Tree] (config>system>script-control>script-policy lifetime)

Full Context

configure system script-control script-policy lifetime

Description

This command is used to configure the maximum amount of time that a script may run.

Default

lifetime 3600

Parameters

seconds

Specifies the maximum amount of time that a script may run, in seconds.

Values 0 to 21474836

Default 3600 (1 hour)

forever

Specifies to allow a script to run indefinitely.

Platforms

7705 SAR Gen 2

16.66 limit-init-exchange

limit-init-exchange

Syntax

limit-init-exchange [**reduced-max-exchange-timeout** *seconds*]

no limit-init-exchange

Context

[Tree] (config>ipsec>ike-policy limit-init-exchange)

Full Context

configure ipsec ike-policy limit-init-exchange

Description

This command limits the number of ongoing IKEv2 initial exchanges per tunnel to 1. When the system receives a new IKEv2 IKE_SA_INIT request when there is an ongoing IKEv2 initial exchange from same peer, then system reduces the timeout value of the existing exchange to the specified **reduced-max-exchange-timeout**. If the **reduced-max-exchange-timeout** is **disabled**, then the system does not reduce the timeout value.

The **no** form of this command reverts to the default value.

Default

limit-init-exchange reduced-max-exchange-timeout 2

Parameters

seconds

Specifies the maximum timeout for the in-progress initial IKE exchange.

Values 2 to 60, disabled

Platforms

7705 SAR Gen 2

16.67 limit-mac-move

limit-mac-move

Syntax

limit-mac-move [**blockable** | **non-blockable**]

no limit-mac-move

Context

[Tree] (config>service>vpls>spoke-sdp limit-mac-move)

[Tree] (config>service>vpls>sap limit-mac-move)

Full Context

configure service vpls spoke-sdp limit-mac-move

configure service vpls sap limit-mac-move

Description

This command indicates whether or not the mac-move agent, when enabled using **config>service>vpls>mac-move** or **config>service>epipe>mac-move**, limits the MAC re-learn (move) rate on this SAP.

Default

limit-mac-move blockable

Parameters

blockable

Specifies that the agent monitors the MAC re-learn rate on the SAP, and it blocks it when the re-learn rate is exceeded.

non-blockable

Specifies that this SAP is not blocked, and another blockable SAP is blocked instead.

Platforms

7705 SAR Gen 2

limit-mac-move

Syntax

limit-mac-move [**blockable** | **non-blockable**]

no limit-mac-move

Context

[Tree] (config>service>pw-template limit-mac-move)

Full Context

configure service pw-template limit-mac-move

Description

This command indicates whether or not the mac-move agent will limit the MAC re-learn (move) rate.

Default

limit-mac-move blockable

Parameters

blockable

The agent will monitor the MAC re-learn rate, and it will block it when the re-learn rate is exceeded.

non-blockable

When specified, a SAP will not be blocked, and another blockable SAP will be blocked instead.

Platforms

7705 SAR Gen 2

16.68 link-address

link-address

Syntax

link-address *ipv6-address*

no link-address

Context

[Tree] (config>service>ies>if>sap>ipsec-gw>dhcp6 link-address)

[Tree] (config>service>vprn>if>sap>ipsec-gw>dhcp6 link-address)

Full Context

configure service ies interface sap ipsec-gw dhcp6 link-address

configure service vprn interface sap ipsec-gw dhcp6 link-address

Description

This command specifies the link address of the relayed DHCPv6 packets sent by the system.

Default

no link-address

Parameters

ipv6-address

Specifies a global unicast IPv6 address.

Platforms

7705 SAR Gen 2

16.69 link-bandwidth

link-bandwidth

Syntax

link-bandwidth

Context

[Tree] (config>service>vprn>bgp>group>neighbor link-bandwidth)

[Tree] (config>service>vprn>bgp>group link-bandwidth)

Full Context

configure service vprn bgp group neighbor link-bandwidth

configure service vprn bgp group link-bandwidth

Description

This command enables the configuration context for handling the link-bandwidth extended community attached to specific BGP routes.

When all used multipaths of an IP prefix correspond to BGP routes with a link-bandwidth extended community, the datapath is programmed to do weighted ECMP across the BGP next-hops in proportion to the bandwidth values.

Platforms

7705 SAR Gen 2

link-bandwidth

Syntax

link-bandwidth

Context

[Tree] (config>router>bgp>group>neighbor link-bandwidth)

[Tree] (config>router>bgp>group link-bandwidth)

Full Context

configure router bgp group neighbor link-bandwidth

configure router bgp group link-bandwidth

Description

This command enables the configuration context for handling the link-bandwidth extended community attached to specific BGP routes.

When all used multipaths of an IP prefix correspond to BGP routes with a link-bandwidth extended community, the datapath is programmed to do weighted ECMP across the BGP next-hops in proportion to the bandwidth values.

Platforms

7705 SAR Gen 2

16.70 link-group

link-group

Syntax

[**no**] **link-group** *link-group-name*

Context

[Tree] (config>service>vprn>isis link-group)

Full Context

configure service vprn isis link-group

Description

This command configures a link-group for the router or VPRN instance.

The **no** form of this command removes the specified link-group.

Parameters

link-group-name

Name of the link-group to be added or removed from the router or VPRN service.

Platforms

7705 SAR Gen 2

link-group

Syntax

link-group *link-group-name*

no link-group

Context

[Tree] (config>router>isis link-group)

Full Context

configure router isis link-group

Description

This command specifies the IS-IS link group associated with this particular level of the interface.

Default

no link-group

Parameters

link-group-name

Specifies an IS-IS link group name, up to 32 characters in length, on the system.

Platforms

7705 SAR Gen 2

16.71 link-local-address

link-local-address

Syntax

link-local-address *ipv6-address* [**dad-disable**]

no link-local-address

Context

[Tree] (config>service>vprn>if>ipv6 link-local-address)

[Tree] (config>router>if>ipv6 link-local-address)

[Tree] (config>service>ies>if>ipv6 link-local-address)

Full Context

configure service vprn interface ipv6 link-local-address

configure router interface ipv6 link-local-address

configure service ies interface ipv6 link-local-address

Description

This command configures the IPv6 Link Local address that is used as a virtual SRRP IPv6 address by the Master SRRP node. This address is sent in the Router Advertisements initiated by the Master SRRP node. Clients use this address as IPv6 default-gateway. Both SRRP nodes, Master and Backup, must be configured with the same Link Local address.

Only one link-local-address is allowed per interface.

Caution:
Removing a manually configured link local address may impact routing protocols or static routes that have a dependency on that address. It is not recommended to remove a link local address when there are active IPv6 subscriber hosts on an IES or VPRN interface.

The **no** form of this command reverts to the default.

Parameters

ipv6-address

Specifies the IPv6 address in the form:

Values

ipv6-address:	x:x:x:x:x:x:x:x
	x:x:x:x:x:x:d.d.d.d
	x - [0..FFFF]H
	d - [0..255]D

dad-disable

Disables Duplicate Address Detection (DAD) and sets the address to preferred, even if there is a duplicated address.

Platforms

7705 SAR Gen 2

16.72 link-local-modifier

link-local-modifier

Syntax

link-local-modifier *modifier*

no link-local-modifier

Context

[Tree] (config>service>ies>if>ipv6>secure-nd link-local-modifier)

Full Context

configure service ies interface ipv6 secure-nd link-local-modifier

Description

This command configures the Cryptographically Generated Address (CGA) modifier for link-local addresses.

Parameters

modifier

Specifies the modifier in 32 hexadecimal nibbles.

Values 0x0 to 0xFFFFFFFF

Platforms

7705 SAR Gen 2

link-local-modifier

Syntax

link-local-modifier *modifier*

no link-local-modifier

Context

[Tree] (config>service>vprn>if>send link-local-modifier)

Full Context

configure service vprn interface ipv6 secure-nd link-local-modifier

Description

This command configures the Cryptographically Generated Address (CGA) modifier for link-local addresses.

Parameters

modifier

Specifies the modifier in 32 hexadecimal nibbles.

Values 0x0–0xFFFFFFFF

Platforms

7705 SAR Gen 2

link-local-modifier

Syntax

link-local-modifier *modifier*

no link-local-modifier

Context

[Tree] (config>router>if>ipv6>secure-nd link-local-modifier)

Full Context

configure router interface ipv6 secure-nd link-local-modifier

Description

This command configures the Cryptographically Generated Address (CGA) modifier for link-local addresses.

Parameters

modifier

Specifies the modifier in 32 hexadecimal nibbles.

Values 0x0 to 0xFFFFFFFF

Platforms

7705 SAR Gen 2

16.73 link-state-export-enable

link-state-export-enable

Syntax

[**no**] **link-state-export-enable**

Context

[Tree] (config>router>bgp link-state-export-enable)

Full Context

configure router bgp link-state-export-enable

Description

This command enables the export of link-state information from the BGP-LS address family into the local Traffic Engineering Database (TED).

The **no** form of this command disables the export of link state information into the TED.

Default

no link-state-export-enable

Platforms

7705 SAR Gen 2

16.74 link-state-import-enable

link-state-import-enable

Syntax

[**no**] **link-state-import-enable**

Context

[Tree] (config>router>bgp link-state-import-enable)

Full Context

configure router bgp link-state-import-enable

Description

This command enables the import of link-state information into the BGP-LS address family for advertisement to other BGP neighbors.

The **no** form of this command disables the import of link state information into the BGP-LS address family.

Default

no link-state-import-enable

Platforms

7705 SAR Gen 2

16.75 link-type

link-type

Syntax

link-type {**pt-pt** | **shared**}

no link-type [**pt-pt** | **shared**]

Context

[Tree] (config>service>vpls>sap>stp link-type)

[Tree] (config>service>template>vpls-sap-template>stp link-type)

[Tree] (config>service>vpls>spoke-sdp>stp link-type)

Full Context

configure service vpls sap stp link-type

configure service template vpls-sap-template stp link-type

configure service vpls spoke-sdp stp link-type

Description

This command instructs STP on the maximum number of bridges behind this SAP or spoke-SDP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected via a shared media, their SAP or spoke-SDPs should all be configured as shared, and timer-based transitions are used.

The **no** form of this command returns the link type to the default value.

Default

link-type pt-pt

Platforms

7705 SAR Gen 2

link-type

Syntax

link-type {**pt-pt** | **shared**}

no link-type

Context

[Tree] (config>service>pw-template>stp link-type)

Full Context

configure service pw-template stp link-type

Description

This command instructs STP on the maximum number of bridges behind this SAP or spoke SDP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected via a shared media, their SAP or spoke SDPs should all be configured as shared, and timer-based transitions are used.

The **no** form of this command returns the link type to the default value.

Default

link-type pt-pt

Platforms

7705 SAR Gen 2

16.76 listen

listen

Syntax

listen

Context

[Tree] (config>system>netconf listen)

Full Context

configure system netconf listen

Description

Commands in this context configure NETCONF listening parameters.

Platforms

7705 SAR Gen 2

16.77 listening-port

listening-port

Syntax

listening-port *port*

no listening-port

Context

[Tree] (config>system>grpc listening-port)

Full Context

configure system grpc listening-port

Description

This command configures the listening port for the gRPC server.

The **no** form of this command reverts to the default.

Default

listening-port 57400

Parameters

port

Specifies the port number.

Values 1024 to 49151, 57400

Default 57400

Platforms

7705 SAR Gen 2

listening-port

Syntax

listening-port *port*

no listening-port

Context

[Tree] (config>system>security>ssh listening-port)

Full Context

configure system security ssh listening-port

Description

This command configures the default SSH port for SSH connections arriving in VPRN or base routing.

The **no** form of this command configures the default SSH port to 22.

Default

no listening-port

Parameters

port

Specifies the port number.

Values 1024 to 49151

Platforms

7705 SAR Gen 2

listening-port

Syntax

listening-port *port*

no listening-port

Context

[Tree] (config>system>security>telnet listening-port)

Full Context

configure system security telnet listening-port

Description

This command configures the default Telnet port for Telnet connections arriving in VPRN or base routing.

The **no** form of this command configures the default Telnet port to 23.

Default

no listening-port

Parameters

port

Specifies the port number.

Values 1024 to 49151

Platforms
7705 SAR Gen 2

16.78 lldp

lldp

Syntax
lldp

Context
[Tree] (config>port>ethernet lldp)

Full Context
configure port ethernet lldp

Description
Commands in this context configure Link Layer Discovery Protocol (LLDP) parameters on the specified port.

Platforms
7705 SAR Gen 2

lldp

Syntax
lldp

Context
[Tree] (config>port>ethernet lldp)

Full Context
configure port ethernet lldp

Description
Commands in this context configure Link Layer Discovery Protocol (LLDP) parameters on the specified port.

Platforms
7705 SAR Gen 2

lldp

Syntax
lldp

Context
[Tree] (config>system lldp)

Full Context
configure system lldp

Description
Commands in this context configure system-wide Link Layer Discovery Protocol parameters.

Platforms
7705 SAR Gen 2

16.79 lldp-member-template

lldp-member-template

Syntax
lldp-member-template

Context
[Tree] (config>lag lldp-member-template)

Full Context
configure lag lldp-member-template

Description
Commands in this context configure the LLDP parameters for member ports.

Platforms
7705 SAR Gen 2

16.80 load

load

Syntax

load *file-url* [**overwrite** | **insert** | **append**]

Context

[Tree] (candidate load)

Full Context

candidate load

Description

This command loads a previously saved candidate configuration into the current candidate. The edit point will be set to the end of the loaded configuration lines. The candidate configuration cannot be modified while a load is in progress.

Default

If the candidate is empty then a load without any of the optional parameters (such as overwrite, and so on) will load the file-url into the candidate. If the candidate is not empty then one of the options, such as overwrite, insert, and so on, must be specified.

Parameters

file-url

Specifies the directory and filename to load.

overwrite

Discards the contents of the current candidate and replace it with the contents of the file.

insert

Inserts the contents of the file at the current edit point.

append

Inserts the contents of the file at the end of the current candidate.

Platforms

7705 SAR Gen 2

16.81 load-balancing

load-balancing

Syntax

load-balancing

Context

[Tree] (config>service>epipe load-balancing)

Full Context

configure service epipe load-balancing

Description

This command enables the load-balancing context to configure interface per-flow load balancing options that will apply to traffic entering this interface and egressing over a LAG/ECMP on system-egress. This is a per interface setting. For load-balancing options that can also be enabled on the system level, the options enabled on the interface level overwrite system level configurations.

Default

not applicable

Platforms

7705 SAR Gen 2

load-balancing

Syntax

load-balancing

Context

[Tree] (config>service>template>vpls-template load-balancing)

Full Context

configure service template vpls-template load-balancing

Description

This command enables the load-balancing context to configure interface per-flow load balancing options that will apply to traffic entering this interface and egressing over a LAG/ECMP on system-egress. This is a per interface setting. For load-balancing options that can also be enabled on the system level, the options enabled on the interface level overwrite system level configurations.

Platforms

7705 SAR Gen 2

load-balancing

Syntax

load-balancing

Context

[Tree] (config>service>ies>if load-balancing)

Full Context

configure service ies interface load-balancing

Description

This command enables the load-balancing context to configure interface per-flow load balancing options that will apply to traffic entering this interface and egressing over a LAG/ECMP on system-egress. This is a per interface setting. For load balancing options that can also be enabled on the system level, the options enabled on the interface level overwrite system level configurations.

Platforms

7705 SAR Gen 2

load-balancing

Syntax

load-balancing

Context

[Tree] (config>service>vprn>nw-if load-balancing)

Full Context

configure service vprn network-interface load-balancing

Description

This command enables the load-balancing context to configure interface per-flow load balancing options that will apply to traffic entering this interface and egressing over a LAG/ECMP on system-egress. This is a per interface setting. For load-balancing options that can also be enabled on the system level, the options enabled on the interface level overwrite system level configurations.

Platforms

7705 SAR Gen 2

load-balancing

Syntax

load-balancing

Context

[Tree] (config>router>if load-balancing)

Full Context

configure router interface load-balancing

Description

This command enables the load-balancing context to configure interface per-flow load balancing options that will apply to traffic entering this interface and egressing over a LAG/ECMP on system-egress. This is a per interface setting. For load-balancing options that can also be enabled on the system level, the options enabled on the interface level overwrite system level configurations.

Platforms

7705 SAR Gen 2

load-balancing

Syntax

load-balancing

Context

[Tree] (config>system load-balancing)

Full Context

configure system load-balancing

Description

This command enables the load-balancing context to configure the interface per-flow load balancing options that will apply to traffic entering this interface and egressing over a LAG/ECMP on system-egress. This is a per interface setting. For load-balancing options that can also be enabled on the system level, the options enabled on the interface level overwrite system level configurations.

Platforms

7705 SAR Gen 2

16.82 load-balancing-algorithm

load-balancing-algorithm

Syntax

load-balancing-algorithm *option*

no load-balancing-algorithm

Context

[Tree] (config>port>ethernet load-balancing-algorithm)

Full Context

configure port ethernet load-balancing-algorithm

Description

This command specifies the load balancing algorithm to be used on this port.

In the default mode, **no load-balancing-algorithm**, the port inherits the global settings. The value is not applicable for ports that do not pass any traffic.

The configuration of load-balancing-algorithm at logical port level has three possible values:

- **include-l4** — Enables inherits system-wide settings including Layer 4 source and destination port value in hashing algorithm.
- **exclude-l4** — Layer 4 source and destination port value will not be included in hashing.
- **no load-balancing-algorithm** — Inherits system-wide settings.

The hashing algorithm addresses finer spraying granularity where many hosts are connected to the network. To address more efficient traffic distribution between network links (forming a LAG group), a hashing algorithm extension takes into account Layer 4 information (src/dst L4-protocol port). The hashing index can be calculated according to the following algorithm:

If [(TCP or UDP traffic) & enabled]

hash (<TCP/UDP ports>, <IP addresses>)

else if (IP traffic)

hash (<IP addresses>)

else

hash (<MAC addresses>)

endif

This algorithm will be used in all cases where IP information in per-packet hashing is included (refer to "Traffic Load Balancing Options” in the *7705 SAR Gen 2 Interface Configuration Guide*). However the Layer 4 information (TCP/UDP ports) will not be used in the following cases:

- fragmented packets

Default

no load-balancing-algorithm

Parameters

option

Specifies the load balancing algorithm to be used on this port.

Values **include-l4** — Specifies that the source and destination ports are used in the hashing algorithm.**exclude-l4** — Specifies that the source and destination ports are not used in the hashing algorithm.

Platforms

7705 SAR Gen 2

16.83 load-balancing-weight

load-balancing-weight

Syntax

load-balancing-weight *value*

no load-balancing-weight [*value*]

Context

[Tree] (config>service>vprn>static-route-entry>next-hop load-balancing-weight)

Full Context

configure service vprn static-route-entry next-hop load-balancing-weight

Description

This command configures a weighted ECMP load-balancing weight for a static route next-hop.

If all of the ECMP next-hops of a static route have a configured load-balancing-weight then packets matching the route are sprayed according to the relative weights. In other words, the next-hop interface with the largest load-balancing weight should receive the most forwarded traffic if weighted ECMP is applicable.

The **no** form of this command disables weighted ECMP for the interface and effectively disables weighted ECMP for the entire static route.

Parameters

value

Specifies the cost metric value.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

load-balancing-weight

Syntax

load-balancing-weight [*weight*]

no load-balancing-weight

Context

[Tree] (config>service>vprn>ospf3>area>if load-balancing-weight)

[Tree] (config>service>vprn>ospf>area>if load-balancing-weight)

Full Context

configure service vprn ospf3 area interface load-balancing-weight

configure service vprn ospf area interface load-balancing-weight

Description

This command configures the weighted ECMP load-balancing weight for an IS-IS, OSPF, and OSPF3 interface. If the interface becomes an ECMP next hop for an IPv4 or IPv6 route, and all the other ECMP next hops are interfaces with configured (non-zero) load-balancing weights, then the traffic distribution over the ECMP interfaces is proportional to the weights. This means that the interface with the largest load-balancing weight receives the most forwarded traffic if weighted ECMP is applicable.

The **no** form of this command disables weighted ECMP for the interface which effectively disables weighted ECMP for any IP prefix that has this interface as a next hop.

Default

no load-balancing-weight

Parameters

weight

Specifies the load balancing weight.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

load-balancing-weight

Syntax

load-balancing-weight *weight*

no load-balancing-weight

Context

[Tree] (config>service>vprn>isis>if load-balancing-weight)

Full Context

configure service vprn isis interface load-balancing-weight

Description

This command configures the weighted ECMP load-balancing weight for an IS-IS interface of the VPRN. If the interface becomes an ECMP next-hop for IPv4 or IPv6 route and all the other ECMP next-hops are interfaces with configured (non-zero) load-balancing weights, then the traffic distribution over the ECMP interfaces is proportional to the weights. In other words, the interface with the largest **load-balancing-weight** should receive the most forwarded traffic if weighted ECMP is applicable.

The **no** form of this command disables weighted ECMP for the interface and, therefore, effectively disables weighted ECMP for any IP prefix that has this interface as a next-hop.

Default

no load-balancing-weight

Parameters

weight

Specifies the load balancing weight.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

load-balancing-weight

Syntax

load-balancing-weight *weight*

no load-balancing-weight

Context

[Tree] (config>router>ldp>if-params>if load-balancing-weight)

Full Context

configure router ldp interface-parameters interface load-balancing-weight

Description

This command configures the load balancing weight for the LDP interface. The load balancing weight, normalized to 64, is used for weighted ECMP of LDP labeled packets over direct network IP interfaces.

If the interface becomes an ECMP next hop for an LDP FEC, and all the other ECMP next hops are interfaces with configured (non-zero) load-balancing weights, then the traffic distribution over the ECMP interfaces is proportional to the normalized weight with a granularity of 64.

If one or more of the LDP interfaces in the ECMP set does not have a configured load-balancing weight, then the system falls back to ECMP.

The **no** form of this command removes the load balancing weight for the LDP interface.

Parameters

weight

Specifies the load balancing weight value.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

load-balancing-weight

Syntax

load-balancing-weight *weight*

no load-balancing-weight

Context

[Tree] (config>router>mpls>lsp load-balancing-weight)

Full Context

configure router mpls lsp load-balancing-weight

Description

This command assigns a weight to an MPLS LSP for use in the weighted load-balancing, or weighted ECMP, over MPLS feature.

Parameters

weight

Specifies a 32-bit integer representing the weight of the LSP.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

load-balancing-weight

Syntax

load-balancing-weight *value*

no load-balancing-weight [*value*]

Context

[Tree] (config>router>static-route-entry>next-hop load-balancing-weight)

Full Context

configure router static-route-entry next-hop load-balancing-weight

Description

This command configures a weighted ECMP load-balancing weight for a static route next-hop.

If all of the ECMP next-hops of a static route have a configured load-balancing-weight then packets matching the route are sprayed according to the relative weights. In other words, the next-hop interface with the largest load-balancing weight should receive the most forwarded traffic if weighted ECMP is applicable.

The **no** form of this command disables weighted ECMP for the interface and effectively disables weighted ECMP for the entire static route.

Parameters

value

Specifies the load balancing weight value.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

load-balancing-weight

Syntax

load-balancing-weight [*value*]

no load-balancing-weight

Context

[Tree] (config>router>isis>interface load-balancing-weight)

Full Context

configure router isis interface load-balancing-weight

Description

This command configures the weighted ECMP load-balancing weight for an IS-IS interface. If the interface becomes an ECMP next hop for an IPv4 or IPv6 route, and all the other ECMP next hops are interfaces with configured (non-zero) load-balancing weights, then the traffic distribution over the ECMP interfaces is proportional to the weights. In other words, the interface with the largest load-balancing weight should receive the most forwarded traffic if weighted ECMP is applicable.

The **no** form of this command disables weighted ECMP for the interface and therefore effectively disables weighted ECMP for any IP prefix that has this interface as a next hop.

Default

no load-balancing-weight

Parameters

value

0 to 4294967295

Platforms

7705 SAR Gen 2

load-balancing-weight

Syntax

load-balancing-weight [*weight*]

no load-balancing-weight

Context

[Tree] (config>router>ospf3>area>if load-balancing-weight)

[Tree] (config>router>ospf>area>if load-balancing-weight)

Full Context

configure router ospf3 area interface load-balancing-weight

configure router ospf area interface load-balancing-weight

Description

This command configures the weighted ECMP load-balancing weight for an OSPF or OSPF3 interface. If the interface becomes an ECMP next hop for an IPv4 or IPv6 route, and all the other ECMP next hops are interfaces with configured (non-zero) load-balancing weights, then the traffic distribution over the ECMP interfaces is proportional to the weights. This means that the interface with the largest load-balancing weight receives the most forwarded traffic if weighted ECMP is applicable.

The **no** form of this command disables weighted ECMP for the interface which effectively disables weighted ECMP for any IP prefix that has this interface as a next hop.

Default

no load-balancing-weight

Parameters

weight

Specifies the load balancing weight.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

16.84 local

local

Syntax

local [**inherit** | **all** | **vc-only** | **none**]

Context

[Tree] (config>service>vprn>ttl-propagate local)

Full Context

configure service vprn ttl-propagate local

Description

This command overrides the global configuration of the TTL propagation for locally generated packets which are forwarded over a MPLS LSPs in a given VPRN service context.

The global configuration is performed under config>router>ttl-propagate>vprn-local.

The default behavior for a given VPRN instance is to inherit the global configuration for the same command. The user can explicitly set the default behavior by configuring the inherit value

Default

local inherit

Parameters

inherit

Specifies the TTL propagation behavior is inherited from the global configuration under **config**>**router**>**ttl-propagate**>**vprn-local**.

none

Specifies the TTL of the IP packet is not propagated into the VC label or labels in the transport label stack.

vc-only

Specifies the TTL of the IP packet is propagated into the VC label and not into the labels in the transport label stack.

all

Specifies the TTL of the IP packet is propagated into the VC label and all labels in the transport label stack.

Platforms

7705 SAR Gen 2

local

Syntax

local

Context

[Tree] (config>ipsec>ts-list local)

Full Context

configure ipsec ts-list local

Description

Commands in this context configure local TS-list parameters. The TS-list is the traffic selector of the local system, such as TSr, when the system acts as an IKEv2 responder.

Platforms

7705 SAR Gen 2

16.85 local-address

local-address

Syntax

local-address *ip-address*

no local-address

Context

[Tree] (config>service>vprn>bgp>group local-address)

[Tree] (config>service>vprn>bgp>group>neighbor local-address)

Full Context

configure service vprn bgp group local-address

configure service vprn bgp group neighbor local-address

Description

Configures the local IP address used by the group or neighbor when communicating with BGP peers.

Outgoing connections use the **local-address** as the source of the TCP connection when initiating connections with a peer.

When a local address is not specified, the OS uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGP peers. This command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of this command removes the configured local-address for BGP.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Parameters

no local-address

The router ID is used when communicating with IBGP peers and the interface address is used for directly connected EBGP peers.

ip-address

The local address expressed in dotted decimal notation. Allowed values are a valid routable IP address on the router, either an interface or system IP address.

Platforms

7705 SAR Gen 2

local-address

Syntax

local-address *ip-address*

no local-address

Context

[Tree] (config>router>pcep>pcc local-address)

Full Context

configure router pcep pcc local-address

Description

This command configures the local IPv4 address of the PCEP speaker.

The PCEP protocol operates over TCP using destination TCP port 4189. The PCE client (PCC) always initiates the connection. After the user configures the PCEP local IPv4 address and the peer IPv4 address on the PCC, the latter initiates a TCP connection to the PCE. If both a local IPv4 and a local IPv6 address are configured, the connection uses the local address that is the same family as the peer address. When the connection is established, the PCC and PCE exchange OPEN messages, which initializes the PCEP session and exchanges the session parameters to be negotiated.

By default, the PCC attempts to reach the remote PCE address out of band using the management port. If it cannot, it attempts to reach the remote PCE address in band. The user can change the configuration of the peer to attempt connecting in band only or out of band only. When the session comes up out of band, the management IP address is used as the local address. The local IPv4 address configured by the user is only used for in-band sessions and is otherwise ignored.

The **no** form of the command removes the configured local address of the PCEP speaker.

Parameters

ip-address

Specifies the IP address of the PCEP speaker to be used for in-band sessions.

Platforms

7705 SAR Gen 2

local-address

Syntax

local-address *ip-address*

no local-address

Context

[Tree] (config>router>origin-validation>rpki-session local-address)

Full Context

configure router origin-validation rpki-session local-address

Description

This command configures the local address to use for setting up the TCP connection used by an RPKI-Router session. The default local-address is the outgoing interface IPv4 or IPv6 address. The local-address cannot be changed without first shutting down the session.

Default

no local-address

Parameters

ip-address

Specifies an IPv4 address or an IPv6 address.

Platforms

7705 SAR Gen 2

local-address

Syntax

local-address [*ip-int-name* | *ip-address* | *ipv6-address*]

no local-address

Context

[Tree] (config>router>bgp>group local-address)

[Tree] (config>router>bgp>group>neighbor local-address)

Full Context

configure router bgp group local-address

configure router bgp group neighbor local-address

Description

This command configures the local IP address used by the group or neighbor when communicating with BGP peers.

Outgoing connections use the **local-address** as the source of the TCP connection when initiating connections with a peer.

When a local address is not specified, the router uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGP peers. This command is used at the neighbor level to revert to the value defined under the group level.

When set to a router interface, the **local-address** inherits the primary IPv4 or IPv6 address of the router interface depending on whether BGP is configured for IPv4 or IPv6. If the corresponding IPv4 or IPv6 address is not configured on the router interface, the BGP sessions that have this interface set as the **local-address** are kept down until an interface address is configured on the router interface.

The **no** form of this command removes the configured local-address for BGP.

The **no** form of this command used at the group level returns the configuration to the value defined at the global level.

The **no** form of this command used at the neighbor level returns the configuration to the value defined at the group level.

Default

no local-address

Parameters

ip-address

Specifies the local address expressed in dotted decimal notation. Allowed value is a valid routable IP address on the router, either an interface or system IP address.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address

Specifies the local address expressed in dotted decimal notation. Allowed value is a valid routable IPv6 address on the router, either an interface or system IPv6 address.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

ip-int-name

Specifies the IP interface name whose address the local address will inherit. The interface can be any network interface configured on the system.

Platforms

7705 SAR Gen 2

16.86 local-address-assignment

local-address-assignment

Syntax

[**no**] **local-address-assignment**

Context

[Tree] (config>service>ies>if>sap>ipsec-gw local-address-assignment)

[Tree] (config>service>vprn>if>sap>ipsec-gw local-address-assignment)

Full Context

configure service ies interface sap ipsec-gw local-address-assignment

configure service vprn interface sap ipsec-gw local-address-assignment

Description

Commands in this context configure local address assignments for the IPsec gateway.

Platforms

7705 SAR Gen 2

16.87 local-address-ipv6

local-address-ipv6

Syntax

local-address-ipv6 *ipv6-address*

no local-address-ipv6

Context

[Tree] (config>router>pcep>pcc local-address-ipv6)

Full Context

configure router pcep pcc local-address-ipv6

Description

This command configures the local IPv6 address of the PCEP speaker.

The PCEP protocol operates over TCP using destination TCP port 4189. The PCE client (PCC) always initiates the connection. After the user configures the PCEP local IPv6 address and the peer IPv6 address on the PCC, the latter initiates a TCP connection to the PCE. If both a local IPv4 and a local IPv6 address are configured, the connection uses the local address that is the same family as the peer address. When the connection is established, the PCC and PCE exchange OPEN messages, which initializes the PCEP session and exchanges the session parameters to be negotiated.

By default, the PCC attempts to reach the remote PCE address out of band using the management port. If it cannot, it attempts to reach the remote PCE address in-band. The user can change the configuration of the peer to attempt connecting in band only or out of band only. When the session comes up out of band, the management IP address is used as the local address. The local IPv6 address configured by the user is only used for in-band sessions and is otherwise ignored.

The **no** form of the command removes the configured local address of the PCEP speaker.

Parameters

ipv6-address

Specifies the IP address of the PCEP speaker to be used for in-band sessions.

Platforms

7705 SAR Gen 2

16.88 local-age

local-age

Syntax

local-age *aging-timer*

no local-age [*aging-timer*]

Context

[Tree] (config>service>vpls local-age)

[Tree] (config>service>template>vpls-template local-age)

Full Context

configure service vpls local-age

configure service template vpls-template local-age

Description

Specifies the aging time for locally learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance. In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a service destination point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the FDB. The **local-age** timer specifies the aging time for local learned MAC addresses.

The **no** form of this command returns the local aging timer to the default value.

Default

local age 300 — Local MACs aged after 300 seconds.

Parameters

aging-timer

Specifies the aging time for local MACs expressed in seconds

Values 60 to 86400

Platforms

7705 SAR Gen 2

16.89 local-as

local-as

Syntax

local-as *as-number* [**private**] [**no-prepend-global-as**]

no local-as

Context

[Tree] (config>service>vprn>bgp>group>neighbor local-as)

[Tree] (config>service>vprn>bgp local-as)

[Tree] (config>service>vprn>bgp>group local-as)

Full Context

configure service vprn bgp group neighbor local-as

configure service vprn bgp local-as

configure service vprn bgp group local-as

Description

This command configures a BGP virtual autonomous system (AS) number.

In addition to the global AS number configured for BGP in the config>router>autonomous-system context, a virtual (local) AS number can be configured to support various AS number migration scenarios. The local AS number is added to the to the beginning the as-path attribute ahead of the router's AS number.

This configuration parameter can be set at three levels: global level (applies to all EBGP peers), group level (applies to all EBGP peers in peer-group) or neighbor level (only applies to EBGP specified peer). Thus, by specifying this at each neighbor level, it is possible to have a separate local-as per EBGP session. The local-as command is not supported for IBGP sessions. When the optional **private** keyword is specified in the command the local-as number is not added to inbound routes from the EBGP peer that has **local-as** in effect.

When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The **private** attribute can be added or removed dynamically by reissuing the command.

Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number. Changing the local AS at the global level in an active BGP instance causes BGP to re-establish the peer relationships with all peers in the group with the new local AS number. Changing the local AS at the neighbor level in an active BGP instance causes BGP to re-establish the peer relationship with the new local AS number.

This is an optional command and can be used in the following circumstance:

Provider router P is moved from AS1 to AS2. The customer router that is connected to P, however, is configured to belong to AS1. To avoid reconfiguring the customer router, the **local-as** value on router P can be set to AS1. Thus, router P adds AS1 to the as-path message for routes it advertises to the customer router.

The **no** form of this command used at the global level removes any virtual AS number configured.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no local-as

Parameters

as-number

The virtual autonomous system number, expressed as a decimal integer.

Values 1 to 65535

private

Specifies the local-as is hidden in paths learned from the peering.

no-prepend-global-as

Specifies that the global-as is hidden in paths announced to the EBGP peer.

Platforms

7705 SAR Gen 2

local-as

Syntax

local-as *as-number* [**private**] [no-prepend-global-as]

no local-as

Context

[Tree] (config>router>bgp local-as)

[Tree] (config>router>bgp>group>neighbor local-as)

[Tree] (config>router>bgp>group local-as)

Full Context

configure router bgp local-as

configure router bgp group neighbor local-as

configure router bgp group local-as

Description

This command configures a BGP local autonomous system (AS) number. In addition to the global AS number configured for BGP using the autonomous-system command, a local AS number can be configured to support various AS number migration scenarios.

When the **local-as** command is applied to a BGP neighbor and the local-as is different from the peer-as, the session comes up as EBGP and by default the global-AS number and then (in that order) the local-as number are prepended to the AS_PATH attribute in outbound routes sent to the peer. In received routes from the EBGP peer, the local AS is prepended to the AS path by default, but this can be disabled with the **private** option.

When the **local-as** command is applied to a BGP neighbor and the local-as is the same as the peer-as, the session comes up as IBGP, and by default, the global-AS number is prepended to the AS_PATH attribute in outbound routes sent to the peer.

This configuration parameter can be set at three levels: global level (applies to all BGP peers), group level (applies to all BGP peers in group) or neighbor level (only applies to one specific BGP neighbor). By specifying this at the neighbor level, it is possible to have a separate **local-as** for each BGP session.

When the optional **no-prepend-global-as** command is configured, the global-as number is not added in outbound routes sent to an IBGP or EBGP peer.

When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The private option can be added or removed dynamically by reissuing the command. Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number. Changing the local AS at the global level in an active BGP instance causes BGP to re-establish the peer relationships with all peers in the group with the new local AS number. Changing the local AS at the neighbor level in an active BGP instance causes BGP to re-establish the peer relationship with the new local AS number.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no local-as

Parameters

as-number

Specifies the virtual autonomous system number expressed as a decimal integer.

Values 1 to 4294967295

private

Specifies the local-as is hidden in paths learned from the peering.

no-prepend-global-as

Specifies that the global-as is hidden in paths announced to the BGP peer.

Platforms

7705 SAR Gen 2

16.90 local-attachment-circuit

local-attachment-circuit

Syntax

local-attachment-circuit *ac-name* [**endpoint** *endpoint-name*] [**bgp** *bgp-instance*] [**create**]

no local-attachment-circuit *ac-name*

Context

[Tree] (config>service>epipe>bgp-evpn local-attachment-circuit)

Full Context

configure service epipe bgp-evpn local-attachment-circuit

Description

This command configures a local attachment circuit (AC) in which the local Ethernet tag can be configured.

The **no** form of this command disables the context.

Default

no local-attachment-circuit

Parameters

ac-name

Specifies the name of the local attachment circuit, up to 32 characters.

endpoint-name

Specifies the name of the endpoint, up to 32 characters.

bgp-instance

Specifies the BGP instance ID.

Values 1 to 2

Default 1

create

Keyword used to create the local AC.

Platforms

7705 SAR Gen 2

16.91 local-dhcp-server

local-dhcp-server

Syntax

local-dhcp-server *server-name* [**create**]

no local-dhcp-server *server-name*

Context

[Tree] (config>router>dhcp local-dhcp-server)

[Tree] (config>service>vprn>dhcp local-dhcp-server)

Full Context

configure router dhcp local-dhcp-server

configure service vprn dhcp local-dhcp-server

Description

This command instantiates a local DHCP server. A local DHCP server can serve multiple interfaces but is limited to the routing context it was which it was created.

The **no** form of this command reverts to the default.

Parameters

server-name

Specifies the name of local DHCP server, up to 32 characters.

create

Keyword used to create the local DHCP server. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

local-dhcp-server

Syntax

local-dhcp-server *server-name* [**create**] [**auto-provisioned**]

no local-dhcp-server *server-name*

Context

[Tree] (config>router>dhcp6 local-dhcp-server)

Full Context

configure router dhcp6 local-dhcp-server

Description

This command instantiates a DHCP6 server. A local DHCP6 server can serve multiple interfaces but is limited to the routing context it was which it was created.

The **no** form of this command reverts to the default.

Parameters

server-name

Specifies the name of local DHCP6 server, up to 32 characters.

create

Keyword used to create the local DHCP or DHCP6 server. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

auto-provisioned

Specifies the auto provisioning mode. This parameter only applies to DHCP6 creation to configure DHCP6 default values.

Platforms

7705 SAR Gen 2

local-dhcp-server

Syntax

[**no**] **local-dhcp-server** *server-name* [**lease-address** *ip-prefix*[*prefix-length*]]

[**no**] **local-dhcp-server** *server-name* [**mac** *ieee-address*]

[**no**] **local-dhcp-server** *server-name* [**link-local-address** *ipv6z-address*]

Context

[Tree] (debug>router local-dhcp-server)

Full Context

debug router local-dhcp-server

Description

This command enables, disables or configures debugging for a local DHCP server.

Parameters

server-name

Specifies an existing local DHCP server name.

ip-prefix[/prefix-length]

Specifies the IP prefix and prefix length of the subnet.

Values ip-prefix — a.b.c.d (host bits must be 0)
length — 0 to 32

ieee-address
Specifies that the provisioned MAC address for the local DHCP server.

ipv6z-address
Specifies the IPv6z address.

ipv6-address:	x:x:x:x:x:x:x:x [-interface]
	x:x:x:x:x:x:d.d.d.d [-interface]
	x: [0 to FFFF]H
	d: [0 to 255]D
interface	up to 32 characters, mandatory for link local addresses

Platforms
7705 SAR Gen 2

local-dhcp-server

Syntax
[**no**] **local-dhcp-server**

Context
[Tree] (config>redundancy>multi-chassis>peer>sync local-dhcp-server)

Full Context
configure redundancy multi-chassis peer sync local-dhcp-server

Description
This command synchronizes DHCP server information.

Default
no local-dhcp-server

Platforms
7705 SAR Gen 2

local-dhcp-server

Syntax

local-dhcp-server *local-server-name*

no local-dhcp-server

Context

[Tree] (config>service>ies>if local-dhcp-server)

[Tree] (config>service>vprn>if>ipv6 local-dhcp-server)

Full Context

configure service ies interface local-dhcp-server

configure service vprn interface ipv6 local-dhcp-server

Description

This command assigns a DHCP server to the interface.

Parameters

local-server-name

Specifies an existing local server name.

Platforms

7705 SAR Gen 2

local-dhcp-server

Syntax

local-dhcp-server *local-server-name*

no local-dhcp-server

Context

[Tree] (config>router>if local-dhcp-server)

[Tree] (config>router>if>ipv6 local-dhcp-server)

Full Context

configure router interface local-dhcp-server

configure router interface ipv6 local-dhcp-server

Description

This command instantiates a local DHCP server. A local DHCP server can serve multiple interfaces but is limited to the routing context in which it was created.

The **no** form of this command reverts to the default value.

Default

no local-dhcp-server

Parameters

local-server-name

Specifies the name of local DHCP server, up to 32 characters.

Platforms

7705 SAR Gen 2

16.92 local-end

local-end

Syntax

local-end {*ip-address* | *ipv6-address*}

no local-end

Context

[Tree] (config>service>sdp local-end)

Full Context

configure service sdp local-end

Description

This command configures the local-end address of the following SDP encapsulation types:

- IPv6 address of the termination point of a SDP of encapsulation **l2tpv3** (L2TP v3 tunnel).
- IPv4/IPv6 source address of a SDP of encapsulation **eth-gre-bridged** (L2oGRE SDP).
- IPv4 source address of a SDP of encapsulation **gre** (GRE SDP).

A change to the value of the local-end parameter requires that the SDP be shut down.

When used as the source address of a SDP of encapsulation **gre** (GRE SDP), the primary IPv4 address of any local network IP interface, loopback or otherwise, may be used.

The address of the following interfaces are not supported:

- unnumbered network IP interface

- IES interface
- VPRN interface
- CSC VPRN interface

The **local-end** parameter value adheres to the following rules:

- A maximum of 15 distinct address values can be configured for all GRE SDPs under the **config**>**service**>**sdp**>**local-end** context, and all L2oGRE SDPs under the **config**>**service**>**system**>**gre-eth-bridged**>**tunnel-termination** context.
- The same source address cannot be used in both contexts since an address configured for a L2oGRE SDP matches an internally created interface that is not available to other applications.
- The **local-end** address of a GRE SDP, when different from system, need not match the primary address of an interface that has the MPLS-over-GRE termination subnet configured, unless a GRE SDP or tunnel from the far-end router terminates on this address.

The **no** form of the command removes the address from the local-end configuration.

Parameters

ip-address* | *ipv6-address

Specifies a IPv4 or IPv6 address for local-end of an SDP in dotted decimal notation.

Values

ip-address	a.b.c.d
ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0..FFFF]H
	d - [0..255]D

Platforms

7705 SAR Gen 2

16.93 local-gateway-address

local-gateway-address

Syntax

local-gateway-address [*ip-address* | *ipv6-address*]

no local-gateway-address

Context

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel local-gateway-address)

[Tree] (config>router>if>ipsec>ipsec-tunnel local-gateway-address)

Full Context

configure service ies interface ipsec ipsec-tunnel local-gateway-address

configure router interface ipsec ipsec-tunnel local-gateway-address

Description

This command configures local gateway address of the IPsec gateway.

Parameters

ip-address

Specifies a unicast IPv4 address, up to 64 characters.

ipv6-address

Specifies a unicast global unicast IPv6 address, up to 64 characters.

Platforms

7705 SAR Gen 2

local-gateway-address

Syntax

local-gateway-address *ip-address*

no local-gateway-address

Context

[Tree] (config>service>ies>if>sap>ipsec-gw local-gateway-address)

[Tree] (config>service>vprn>if>sap>ipsec-gw local-gateway-address)

Full Context

configure service ies interface sap ipsec-gw local-gateway-address

configure service vprn interface sap ipsec-gw local-gateway-address

Description

This command configures local gateway address of the IPsec gateway.

Parameters

ip-address

Specifies a unicast IPv4 address or a global unicast IPv6 address. This address must be within the subnet of the public interface.

Platforms

7705 SAR Gen 2

local-gateway-address

Syntax

local-gateway-address *ip-address* **peer** *ip-address* **delivery-service** *service-id*

no local-gateway-address

Context

[Tree] (config>service>vprn>if>sap>ipsec-tunnel local-gateway-address)

Full Context

configure service vprn interface sap ipsec-tunnel local-gateway-address

Description

This command specifies the local gateway address used for the tunnel and the address of the remote security gateway at the other end of the tunnel remote peer IP address to use.

Default

no local-gateway-address

Parameters

ip-address

IP address of the local end of the tunnel.

delivery-service *service-id*

The ID of the IES or VPRN (front-door) delivery service of this tunnel. Use this service-id to find the VPRN used for delivery.

Values *service-id*: 1 to 2147483648

svc-name: Specifies an existing service name up to 64 characters in length.

Platforms

7705 SAR Gen 2

16.94 local-id

local-id

Syntax

local-id *type* [**value** *value*]

no local-id

Context

[Tree] (config>router>if>ipsec>ipsec-tunnel>dyn local-id)

[Tree] (config>service>vprn>if>sap>ipsec-tun>dyn local-id)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>dyn local-id)

[Tree] (config>service>vprn>if>sap>ipsec-gw local-id)

[Tree] (config>service>ies>if>sap>ipsec-gw local-id)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>dyn local-id)

[Tree] (config>ipsec>trans-mode-prof>dyn local-id)

Full Context

configure router interface ipsec ipsec-tunnel dynamic-keying local-id

configure service vprn interface sap ipsec-tunnel dynamic-keying local-id

configure service vprn interface ipsec ipsec-tunnel dynamic-keying local-id

configure service vprn interface sap ipsec-gw local-id

configure service ies interface sap ipsec-gw local-id

configure service ies interface ipsec ipsec-tunnel dynamic-keying local-id

configure ipsec ipsec-transport-mode-profile dynamic-keying local-id

Description

This command specifies the local ID used for IDi or IDr for IKEv2 negotiation.

The default behavior depends on the local-auth-method as follows:

- Psk: local tunnel IP address
- Cert-auth: subject of the local certificate

The **no** form of this command removes the parameters from the configuration.

Default

no local-id

Parameters

type

Specifies the type of local ID payload, which could be IPv4 or IPv6 address or FQDN domain name or distinguish the name of the subject in the X.509 certificate.

Values ipv4 — Specifies to use IPv4 as the local ID type; the default value is the local tunnel end-point address.

ipv6 — Specifies to use IPv6 as the local ID type; the default value is the local tunnel end-point address.

fq1dn — Specifies to use FQDN as the local ID type. The value must be configured.

value

Specifies the data type as an enumerated integer that describes the local identifier type used for IDi or IDr for IKEv2, up to 255 characters.

Platforms

7705 SAR Gen 2

16.95 local-ip

local-ip

Syntax

local-ip {*ip-prefix/prefix-length* | *ip-prefix netmask* | **any**}

Context

[Tree] (config>service>vprn>ipsec>sec-plcy>entry local-ip)

[Tree] (config>router>ipsec>sec-plcy>entry local-ip)

Full Context

configure service vprn ipsec security-policy entry local-ip

configure router ipsec security-policy entry local-ip

Description

This command configures the local (from the VPN) IP prefix/mask for the policy parameter entry.

Only one entry is necessary to describe a potential flow. The **local-ip** and **remote-ip** commands can be defined only once. The system evaluates:

- the local IP as the source IP when traffic is examined in the direction of the flows from private to public and as the destination IP when traffic flows from public to private
- the remote IP as the source IP when traffic flows public to private and as the destination IP when traffic flows from private to public

Parameters

ip-prefix

The destination address of the aggregate route in dotted decimal notation

Values a.b.c.d (host bits must be 0)

prefix-length 1 to 32

netmask

The subnet mask in dotted decimal notation

any

keyword to specify that it can be any address

Platforms

7705 SAR Gen 2

16.96 local-lsr-id

local-lsr-id

Syntax

local-lsr-id {**system** | **interface**} [**32bit-format**]

local-lsr-id *interface-name* [**32bit-format**]

no local-lsr-id

Context

[Tree] (config>router>ldp>if-params>if>ipv4 local-lsr-id)

[Tree] (config>router>ldp>if-params>if>ipv6 local-lsr-id)

Full Context

configure router ldp interface-parameters interface ipv4 local-lsr-id

configure router ldp interface-parameters interface ipv6 local-lsr-id

Description

This command enables the use of the address of the local LDP interface, or any other network interface configured on the system, as the LSR-ID to establish link LDP Hello adjacency and LDP session with directly connected LDP peers. The network interface can be a loopback or not.

Link LDP sessions to all peers discovered over a given LDP interface share the same local LSR-ID. However, LDP sessions on different LDP interfaces can use different network interface addresses as their local LSR-ID.

By default, the LDP session to a peer uses the system interface address as the LSR-ID unless explicitly configured using this command. The system interface must always be configured on the router, or the LDP protocol will not come up on the node. There is no requirement to include the system interface in any routing protocol.

At initial configuration, the LDP session to a peer will remain down while the network interface used as LSR-ID is down. LDP will not try to bring it up using the system interface.

If the network IP interface used as LSR-ID goes down, the LDP sessions to all discovered peers using this LSR-ID go down.

When an interface other than the system is used as the LSR-ID, the transport connection (TCP) for the link LDP session will also use the address of that interface as the transport address. If the system or interface value is configured in the **config**>**router**>**ldp**>**if-params**>**if**>**ipv4** or **config**>**router**>**ldp**>**if-**

params>if>ipv6> transport-address context, it will be overridden with the address of the LSR-ID interface.

When the **local-lsr-id** command is enabled with the **32bit-format** option, an SR OS LSR will be able to establish an LDP IPv6 Hello adjacency and an LDP IPv6 session with an RFC 7552 compliant peer LSR. The LSR uses a 32-bit LSR-ID set to the value of the IPv4 address of the specified local LSR-ID interface and a 128-bit transport address set to the value of the IPv6 address of the specified local LSR-ID interface.

Note:
The system interface cannot be used as a local LSR-ID with the **32bit-format** option enabled because the system interface is the default LSR-ID and transport address for all LDP sessions to peers on this LSR. This configuration is blocked in the CLI.

If the user enables the **32bit-format** option in the IPv6 context of a running LDP interface, the already established LDP IPv6 Hello adjacency and LDP IPv6 session will be brought down and re-established with the new 32-bit LSR-ID value.

If the user changes the LSR-ID value between **system**, **interface**, and *interface-name*, or enables the **32bit-format** option while the LDP session is up, LDP will immediately tear down all sessions using this LSR-ID and will attempt to re-establish them using the new LSR-ID.

The **no** form of this command returns to the default behavior, in which case the system interface address is used as the LSR-ID.

Default

no local-lsr-id

Parameters

system

Specifies the use of the address of the system interface as the value of the LSR-ID of this LDP LSR.

interface

Specifies the use of the address of the local LDP interface as the value of the LSR-ID of this LDP LSR.

interface-name interface-name

Specifies the name, up to 32 character, of the network IP interface (which address is used as the LSR-ID of this LDP LSP). An interface name cannot be in the form of an IP address. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

32bit-format

Specifies the use of the IPv4 address of the local LSR-ID interface as the LSR-ID of this LDP LSR.

Platforms

7705 SAR Gen 2

local-lsr-id

Syntax

local-lsr-id *interface-name* [**32bit-format**]

no local-lsr-id

Context

[Tree] (config>router>ldp>targ-session>peer-template local-lsr-id)

[Tree] (config>router>ldp>targ-session>peer local-lsr-id)

Full Context

configure router ldp targeted-session peer-template local-lsr-id

configure router ldp targeted-session peer local-lsr-id

Description

This command enables the use of the address of any network interface configured on the system, as the LSR-ID to establish a targeted LDP Hello adjacency and a targeted LDP session with an LDP peer. The network interface can be a loopback or not.

By default, the targeted LDP session to a peer uses the system interface address as the LSR-ID and as the transport address, unless explicitly configured using this command. The system interface must always be configured on the router, or the LDP protocol will not come up on the node. There is no requirement to include the system interface in any routing protocol.

When the **local-lsr-id** command is enabled with the **32bit-format** option, an SR OS LSR will be able to establish a targeted LDP IPv6 Hello adjacency and a targeted LDP IPv6 session with an RFC 7552 compliant peer LSR. The LSR uses a 32-bit LSR-ID set to the value of the IPv4 address of the specified local LSR-ID interface and a 128-bit transport address set to the value of the IPv6 address of the specified local LSR-ID interface.

Note:
The system interface cannot be used as a local LSR-ID with the 3**2bit-format** option enabled because the system interface is the default LSR-ID and transport address for all targeted LDP sessions to peers on this LSR. This configuration is blocked in the CLI.

If the user enables the **32bit-format** option in the IPv6 context of a running targeted LDP peer, the already established targeted LDP IPv6 Hello adjacency and targeted LDP IPv6 session will be brought down and re-established with the new 32-bit LSR-ID value.

If the user changes the local LSR-ID value or enables/disables the **32bit-format** option, while the targeted LDP session is up, LDP will immediately tear down the targeted session using this LSR-ID and will attempt to re-establish it using the new LSR-ID.

The **no** form of this command returns to the default behavior, in which case the system interface address is used as the LSR-ID.

Default

no local-lsr-id

Parameters

interface-name

Specifies the name, up to 32 characters, of the network IP interface (which address is used as the LSR-ID of this LDP LSP). An interface name cannot be in the form of an IP address. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

32bit-format

Specifies the use of the IPv4 address of the local LSR-ID interface as the LSR-ID of this LDP LSR.

Platforms

7705 SAR Gen 2

16.97 local-max-checkpoints

local-max-checkpoints

Syntax

local-max-checkpoints [*number-of-files*]

no local-max-checkpoints

Context

[Tree] (config>system>rollback local-max-checkpoints)

Full Context

configure system rollback local-max-checkpoints

Description

This command configures the maximum number of rollback checkpoint files when the rollback-location is on local compact flash.

Default

no local-max-checkpoints

Parameters

number of files

Specifies the maximum rollback files on a compact flash.

Values 1 to 50

Platforms

7705 SAR Gen 2

16.98 local-monitoring-policer

local-monitoring-policer

Syntax

[**no**] **local-monitoring-policer** *policer-name* [**create**]

Context

[Tree] (config>sys>security>dist-cpu-protection>policy local-monitoring-policer)

Full Context

configure system security dist-cpu-protection policy local-monitoring-policer

Description

This command configures a monitoring policer that is used to monitor the aggregate rate of several protocols arriving on an object (for example, SAP). When the **local-monitoring-policer** is determined to be in a nonconforming state (at the end of a minimum monitoring time of 60 seconds) then the system will attempt to allocate dynamic policers for the particular object for any protocols associated with the local monitor (for example, using the **protocol** *name* **enforcement dynamic** *policer-name* CLI command).

If the system cannot allocate all the dynamic policers within 150 seconds, it will stop attempting to allocate dynamic policers, raise a LocMonExcdAllDynAlloc log event, and go back to using the local monitor. The local monitor may then detect exceeded packets again and make another attempt at allocating dynamic policers.

Once this *policer-name* is referenced by a protocol then this policer will be instantiated for each "object" that is created and references this DDoS policy. If there is no policer free then the object will be blocked from being created.

Parameters

policy-name

Specifies name of the policy, up to 32 characters.

Platforms

7705 SAR Gen 2

16.99 local-preference

local-preference

Syntax

local-preference *local-preference*

no local-preference

Context

[Tree] (config>service>vprn>bgp>group local-preference)

[Tree] (config>service>vprn>bgp>group>neighbor local-preference)

[Tree] (config>service>vprn>bgp local-preference)

Full Context

configure service vprn bgp group local-preference

configure service vprn bgp group neighbor local-preference

configure service vprn bgp local-preference

Description

This command enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute. This value is used if the BGP route arrives from a BGP peer without the **local-preference** integer set.

The specified value can be overridden by any value set via a route policy. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no local-preference - Does not override the local-preference value set in arriving routes and analyze routes without local preference with value of 100.

Parameters

local-preference

The local preference value to be used as the override value, expressed as a decimal integer.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

local-preference

Syntax

local-preference *local-preference*

no local-preference

Context

[Tree] (config>router>bgp>group>neighbor local-preference)

[Tree] (config>router>bgp local-preference)

[Tree] (config>router>bgp>group local-preference)

Full Context

configure router bgp group neighbor local-preference

configure router bgp local-preference

configure router bgp group local-preference

Description

This command enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute.

This value is used if the BGP route arrives from a BGP peer without the **local-preference** integer set.

The specified value can be overridden by any value set via a route policy. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to the specified peer). The most specific value is used.

The **no** form of this command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no local-preference

Parameters

local-preference

Specifies the local preference value to be used as the override value expressed as a decimal integer.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

local-preference

Syntax

local-preference *preference* [**equal** | **or-higher** | **or-lower**]

no local-preference

Context

[Tree] (config>router>policy-options>policy-statement>entry>from local-preference)

Full Context

configure router policy-options policy-statement entry from local-preference

Description

This command matches BGP routes based on local preference (the value in the LOCAL_PREF attribute).

If no comparison qualifiers are present (**equal**, **or-higher**, **or-lower**), then **equal** is the implied default.

A non-BGP route does not match a policy entry if it contains the **local-preference** command.

Default

no local-preference

Parameters

preference

Specifies the local preference value.

Values 0 to 4294967295, or a parameter name delimited by starting and ending at-sign (@) characters

equal

Specifies that matched routes should have the same local preference as the value specified.

or-higher

Specifies that matched routes should have the same or a greater local preference as the value specified.

or-lower

Specifies that matched routes should have the same or a lower local preference as the value specified.

Platforms

7705 SAR Gen 2

local-preference

Syntax

local-preference *preference*

no local-preference

Context

[Tree] (config>router>policy-options>policy-statement>default-action local-preference)

Full Context

configure router policy-options policy-statement default-action local-preference

Description

This command assigns a BGP local preference to routes matching a route policy statement entry.

If no local preference is specified, the BGP configured local preference is used.

The **no** form of this command disables assigning a local preference in the route policy entry.

Default

no local-preference

Parameters

preference

Specifies the local preference expressed as a decimal integer.

Values 0 to 4294967295 name — Specifies the local preference parameter variable name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@”.

Platforms

7705 SAR Gen 2

16.100 local-prefix

local-prefix

Syntax

local-prefix *local-prefix* [**create**]

no local-prefix *local-prefix*

Context

[Tree] (config>service>pw-routing local-prefix)

Full Context

configure service pw-routing local-prefix

Description

This command configures one or more node prefix values to be used for MS-PW routing. At least one prefix must be configured on each node that is an S-PE or a T-PE.

The **no** form of this command removes a previously configured prefix, and will cause the corresponding route to be withdrawn if it has been advertised in BGP.

Default

no local-prefix

Parameters

local-prefix

Specifies a 32 bit prefix for the AII. One or more prefix values, up to a maximum of 16, may be assigned to the 7705 SAR Gen 2 node. The global ID can contain the 2-octet or 4-octet value of the provider's Autonomous System Number (ASN). The presence of a global ID based on the provider's ASN ensures that the AII for spoke-SDPs configured on the node will be globally unique.

Values <global-id>:<ip-addr>| <raw-prefix>

ip-addr	a.b.c.d
raw-prefix	1 to 4294967295
global-id	1 to 4294967295

Platforms

7705 SAR Gen 2

16.101 local-proxy-arp

local-proxy-arp

Syntax

[**no**] **local-proxy-arp**

Context

[Tree] (config>service>ies>if local-proxy-arp)

[Tree] (config>service>vprn>if local-proxy-arp)

Full Context

configure service ies interface local-proxy-arp

configure service vprn interface local-proxy-arp

Description

This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus becomes the forwarding point for all traffic between hosts in that subnet.

When **local-proxy-arp** is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

The **no** form of this command reverts to the default.

Platforms

7705 SAR Gen 2

local-proxy-arp

Syntax

[**no**] **local-proxy-arp**

Context

[Tree] (config>router>if local-proxy-arp)

Full Context

configure router interface local-proxy-arp

Description

This command enables local proxy ARP on the interface.

Default

no local-proxy-arp

Platforms

7705 SAR Gen 2

16.102 local-proxy-nd

local-proxy-nd

Syntax

[**no**] **local-proxy-nd**

Context

[Tree] (config>service>vprn>if>ipv6 local-proxy-nd)

[Tree] (config>service>ies>if>ipv6 local-proxy-nd)

Full Context

configure service vprn interface ipv6 local-proxy-nd

configure service ies interface ipv6 local-proxy-nd

Description

This command enables local proxy neighbor discovery on the interface.

When this command is enabled, the interface replies to neighbor solicitation requests when both the hosts are on the same subnet. In this case, ICMP redirects are disabled. When this command is disabled, the interface does not reply to neighbor solicitation requests if both the hosts are on the same subnet.

The **no** form of this command reverts to the default.

Platforms

7705 SAR Gen 2

local-proxy-nd

Syntax

[**no**] **local-proxy-nd**

Context

[Tree] (config>router>if>ipv6 local-proxy-nd)

Full Context

configure router interface ipv6 local-proxy-nd

Description

This command enables local proxy neighbor discovery on the interface.

The **no** form of this command disables local proxy neighbor discovery.

Platforms

7705 SAR Gen 2

16.103 local-routes-domain-id

local-routes-domain-id

Syntax

local-routes-domain-id [*global-field*:*local-field*]

no local-routes-domain-id

Context

[Tree] (config>service>vprn local-routes-domain-id)

Full Context

configure service vprn local-routes-domain-id

Description

This command specifies the domain ID that is used in the D-PATH attribute for local routes before those routes are exported to a BGP neighbor using BGP-IPVPN, EVPN-IFF, EVPN-IFL or PE-CE BGP. A local route is a non-BGP route installed in the VPRN route table and learned using static route or an IGP.

The domain IDs are used in the D-PATH attribute, in accordance with *draft-ietf-bess-evpn-ipvpn-interworking*. The D-PATH attribute is modified by gateway routers, where a gateway is defined as a PE where a VPRN is instantiated, and that VPRN advertises or receives routes from multiple BGP owners (for example, EVPN-IFL and BGP-IPVPN).

The D-PATH attribute is used on gateways to detect loops (for received routes where the D-PATH contains a local domain ID) and to make BGP best path selection decisions based on the D-PATH length (shorter D-PATH is preferred).

The **no** form of this command removes the domain ID for local routes.

Default

no local-routes-domain-id

Parameters

global-field:local-field

Specifies the domain ID for local routes.

Values *4byte-GlobalAdminValue:2byte-LocalAdminValue*

4byte-GlobalAdminValue:	0 to 4294967295
2byte-LocalAdminValue	0 to 65535

Platforms

7705 SAR Gen 2

16.104 local-source-address

local-source-address

Syntax

local-source-address {*ip-int-name* | *ip-address*}

no local-source-address

Context

[Tree] (config>system>telemetry>persistent>subscription local-source-address)

Full Context

configure system telemetry persistent-subscriptions subscription local-source-address

Description

This command is used to assign a source IP address in the respective persistent subscription context for use when packets are sent out.

The **no** form of this command removes this address from the configuration.

Parameters

ip-int-name

Specifies the source IP address name, up to 64 characters.

ip-address

Specifies the source IP address.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x:-[0 to FFFF]H
	d: [0 to 255]D

Platforms

7705 SAR Gen 2

local-source-address

Syntax

local-source-address {*ip-int-name* | *ip-address*}

no local-source-address

Context

[Tree] (config>system>grpc-tunnel>destination-group>destination local-source-address)

Full Context

configure system grpc-tunnel destination-group destination local-source-address

Description

This command configures a local source IP address in the destination group context for use when packets are sent out.

The **no** form of this command removes this address from the configuration.

Default

no local-source-address

Parameters

ip-int-name

Specifies the source IP address name, up to 64 characters.

ip-address

Specifies the source IPv4 address (in dotted decimal notation) or IPv6 address.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x:-[0 to FFFF]H
	d: [0 to 255]D

Platforms

7705 SAR Gen 2

16.105 local-sr-protection

local-sr-protection

Syntax

local-sr-protection *local-sr-protection*

no local-sr-protection

Context

[Tree] (config>router>mpls>lsp local-sr-protection)

[Tree] (config>router>mpls>lsp-template local-sr-protection)

Full Context

configure router mpls lsp local-sr-protection

configure router mpls lsp-template local-sr-protection

Description

This command configures the SR LFA protection needed for the adjacencies used in the path computation of an SR-TE LSP by the local CSPF.

The default value of the command is **preferred**. The local CSPF will prefer a protected adjacency over an unprotected adjacency whenever both exist for a TE link. However, the entire computed path can combine both types of adjacencies.

When the user enables the **mandatory** value, CSPF uses it as an additional path constraint and selects protected adjacencies exclusively in computing the path of the SR-TE LSP. CSPF will return no path if all candidate paths that otherwise satisfy all other LSP path constraints do not have an unprotected SID for each of their TE links.

Similarly, if the user enables the value **none**, CSPF uses it as an additional path constraint and selects unprotected adjacencies exclusively in computing the path of the SR-TE LSP. CSPF will return no path if all candidate paths that otherwise satisfy all other LSP path constraints do not have a protected SID for each of their TE links.

The **no** form of this command returns the command to its default value.

Default

no local-sr-protection

Parameters

local-sr-protection

Specifies the local-sr-protection for LSPs.

Values none — Selects unprotected adjacencies only in the SR-TE LSP path computation.

preferred — Prefers protected adjacencies in the SR-TE LSP path computation.

mandatory — Selects protected adjacencies only in the SR-TE LSP path computation.

Platforms

7705 SAR Gen 2

16.106 local-user-db

local-user-db

Syntax

local-user-db *local-user-db-name* [**create**]

no local-user-db *local-user-db-name*

Context

[Tree] (config>subscr-mgmt local-user-db)

Full Context

configure subscriber-mgmt local-user-db

Description

Commands in this context configure a local user database.

The **no** form of this command reverts to the default.

Parameters

local-user-db-name

Specifies the name of a local user database, up to 32 characters.

Platforms

7705 SAR Gen 2

16.107 local-v6-ip

local-v6-ip

Syntax

local-v6-ip *ipv6-prefix/prefix-length*

local-v6-ip any

no local-v6-ip

Context

[Tree] (config>service>vprn>ipsec>sec-plcy>entry local-v6-ip)

[Tree] (config>router>ipsec>sec-plcy>entry local-v6-ip)

Full Context

configure service vprn ipsec security-policy entry local-v6-ip

configure router ipsec security-policy entry local-v6-ip

Description

This command specifies the local v6 prefix for the security-policy entry.

Parameters

ipv6-prefix/prefix-length

Specifies the local v6 prefix and length

Values	ipv6-address/prefix: ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x [0 to FFFF]H
		d [0 to 255]D
		host bits must be 0
		:: not allowed
		prefix-length [1 to 128]

any

keyword to specify that it can be any address.

Platforms

7705 SAR Gen 2

16.108 location

location

Syntax

location *cflash-id*

no location

Context

[Tree] (config>system>persistence>nat-fwd location)

[Tree] (config>system>persistence>dhcp-server location)

[Tree] (config>system>persistence>ancp location)

Full Context

configure system persistence nat-port-forwarding location

configure system persistence dhcp-server location

configure system persistence ancp location

Description

This command instructs the system where to write the persistency files for the corresponding application. Each application creates two files on the flash card, one with suffix .i<version>, referencing an index file, and the other with suffix .0<version>, where <version> is a 2-digit number reflecting the file version. These versions are not related to the SR OS release running on the node. The <version> can remain the same over two major releases, for example, when no format change is made to the persistency file. On boot, the system scans the file systems looking for the corresponding persistency files, and the load begins.

For example, in the subscriber management context, the location specifies the flash device on a CPM card where the data for handling subscriber management persistency is stored.

The **no** form of this command returns the system to the default. If there is a change in file location while persistence is running, a new file will be written on the new flash, and then the old file will be removed.

Default

no location

Parameters

cflash-id

Specifies the compact flash device name.

Values cf1:, cf2:, cf3:

Platforms

7705 SAR Gen 2

location

Syntax

location *location*

no location

Context

[Tree] (config>system location)

Full Context

configure system location

Description

This command creates a text string that identifies the system location for the device.

Only one location can be configured. If multiple locations are configured, the last one entered overwrites the previous entry.

The **no** form of the command reverts to the default value.

Parameters

location

Specifies the location as a character string. The string may be up to 80 characters. Any printable, seven-bit ASCII characters can be used within the string. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

location

Syntax

location *file-url*

no location

Context

[Tree] (config>system>script-control>script location)

Full Context

configure system script-control script location

Description

This command is used to identify the location of a script to be scheduled.

The **no** form of the command removes the location.

Default

no location

Parameters

file-url

Specifies the location to search for scripts.

Values *local-url* | *remote-url*

local-url — [*cflash-id/*] [*file-path*] 200 chars max, including *cflash-id* directory length 99 characters max each

remote url — [{ftp:// | tftp://}*login*:*password@remote-location/*][*file-path*] 255 characters max directory length 99 characters max each

remote-location — [*hostname* | *ipv4-address* | *ipv6-address*]

ipv4-address — *a.b.c.d*

ipv6-address — x:x:x:x:x:x:x:x[-*interface*]

x:x:x:x:x:x:d.d.d.d[-*interface*]

x — [0 to FFFF]H

d — [0 to 255]D

interface — 32 characters max, for link local addresses

cflash-id — cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

location

Syntax

location *cflash-id* [*backup-cflash-id*]

no location

Context

[Tree] (config>log>file-id location)

Full Context

configure log file-id location

Description

This command specifies the primary and optional backup location where the log or billing file will be created.

The **location** command is optional. If the location command not explicitly configured, log files will be created on cf1: and accounting files will be created on cf2: without overflow onto other devices. Generally, cf3: is reserved for system files (configurations, images, and so on).

When multiple location commands are entered in a single file ID context, the last command overwrites the previous command.

When the location of a file ID that is associated with an active log ID is changed, the log events are not immediately written to the new location. The new location does not take effect until the log is rolled over either because the rollover period has expired or a **clear log log-id** command is entered to manually rollover the log file.

When creating files, the primary location is used as long as there is available space. If no space is available, an attempt is made to delete unnecessary files that are past their retention date.

If sufficient space is not available an attempt is made to remove the oldest to newest closed log or accounting files. After each file is deleted, the system attempts to create the new file.

A medium severity trap is issued to indicate that a compact flash is either not available or that no space is available on the specified flash and that the backup location is being used.

A high priority alarm condition is raised if none of the configured compact flash devices for this file ID are present or if there is insufficient space available. If space does become available, then the alarm condition will be cleared.

Log files are created on cf1: and accounting files are created on cf2.

Use the **no** form of this command to revert to default settings.

Default

no location

Parameters

cflash-id

Specify the primary location.

Values cflash-id: cf1:, cf2:, cf3:

backup-cflash-id

Specify the secondary location.

Values cflash-id: cf1:, cf2:, cf3:

Platforms

7705 SAR Gen 2

location

Syntax

location *location-id* [**primary-ip-address** *ipv4-address*] [**secondary-ip-address** *ipv4-address*] [**tertiary-ip-address** *ipv4-address*]

Context

[Tree] (config>router>bgp>optimal-route-reflection location)

Full Context

configure router bgp optimal-route-reflection location

Description

This command configures the location ID for the for the route reflector. A BGP neighbor can be associated with a location if it is a route-reflector client.

Parameters

location-id

Specifies an optimal-route-reflection location.

Values 1 to 255

ipv4-address

Specifies the primary, secondary, or tertiary IP address.

Values primary ipv4-address, secondary ipv4-address, tertiary ipv4-address

Platforms

7705 SAR Gen 2

16.109 lock

lock

Syntax

[**no**] **lock**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization lock)

Full Context

configure system security profile netconf base-op-authorization lock

Description

This command authorizes a user associated with the profile to send a NETCONF <lock> RPC. This lock RPC allows a NETCONF client to lock a configuration datastore.

The **no** form of the command denies the user from requesting a lock.

Default

no lock

Platforms

7705 SAR Gen 2

16.110 lock-override

lock-override

Syntax

[**no**] **lock-override**

Context

[Tree] (config>system>script-control>script-policy lock-override)

Full Context

configure system script-control script-policy lock-override

Description

This command allows a triggered EHS/CRON script to execute while there is a datastore lock, started by an MD interface, in place.

A triggered EHS/CRON script queues until an **ongoing commit** (or **confirmed-commit**) is done. When an EHS/CRON script is triggered while the **lock-override** CLI knob is on, SR OS behaves as follows.

When an exclusive session is in place:

- Keep if it is an MD-CLI session. Disconnect if it is a NETCONF session
- Lose the exclusive lock
- Lose any uncommitted configuration changes

When a global session is in place:

- Keep the MD-CLI or NETCONF session
- Keep the uncommitted configuration changes
- An update may be required after committing the EHS/CRON script configuration changes

The **no** form of this command does not allow the script to execute while there is a datastore lock in place.

Default

lock-override

Platforms

7705 SAR Gen 2

16.111 lockout

lockout

Syntax

lockout failed-attempts *count* **duration** *duration-minutes* **block** *block-minutes* [**max-port-per-ip** *number-of-ports*]

no lockout

Context

[Tree] (config>ipsec>ike-policy lockout)

Full Context

configure ipsec ike-policy lockout

Description

This command enables the lockout mechanism for the IPsec tunnel. The system will lock out an IPsec client for the configured time interval if the number of failed authentications exceeds the configured value within the specified duration. This command only applies when the system acts as a tunnel responder.

A client is defined as the tunnel IP address plus the port.

Optionally, the **max-port-per-ip** parameter can be configured as the maximum number of ports allowed behind the same IP address. If this threshold is exceeded, then all ports behind the IP address are blocked.

The **no** form of this command disables the lockout mechanism.

Default

no lockout

Parameters

count

Specifies the maximum number of failed authentications allowed during the *duration-minutes* interval.

Values 1 to 64

Default 3

duration-minutes

Specifies the interval of time, in minutes, during which the configured failed authentication count must be exceeded in order to trigger a lockout.

Values 1 to 60

Default 5

block-minutes

Specifies the number of minutes that the client is blocked if the configured failed authentication count is exceeded.

Values 1 to 1440, infinite

Default 10

number-of-ports

Specifies the maximum number of ports allowed behind the same IP address.

Values 1 to 32000

Default 16

Platforms

7705 SAR Gen 2

lockout

Syntax

clear lockout {**user** *user-name* | **all**}

Context

[Tree] (admin>clear lockout)

Full Context

admin clear lockout

Description

This command is used to clear any lockouts for a specific user, or for all users.

Parameters

user-name

Clears the locked username.

all

Clears all locked usernames.

Platforms

7705 SAR Gen 2

16.112 log

log

Syntax

log

Context

[Tree] (config log)

Full Context

configure log

Description

Commands in this context are used to configure both event logs and accounting logs. Event logs control the generation, dissemination and recording of system events for monitoring status and troubleshooting faults within the system. Event logging configuration includes syslog, snmp notifications (traps), NETCONF notifications and other types of event log outputs. Accounting logs collect comprehensive accounting statistics and write them to XML files on the compact flash in order to support a variety of billing models.

Platforms

7705 SAR Gen 2

log

Syntax

[**no**] **log**

Context

[Tree] (config>service>vprn>static-route-entry>indirect>cpe-check log)

[Tree] (config>service>vprn>static-route-entry>next-hop>cpe-check log)

Full Context

configure service vprn static-route-entry indirect cpe-check log

configure service vprn static-route-entry next-hop cpe-check log

Description

This optional parameter enables the ability to log transitions between active and in-active based on the CPE connectivity check. Events will be sent to the system log, syslog and SNMP traps.

Default

no log

Platforms

7705 SAR Gen 2

log

Syntax

log

Context

[Tree] (config>service>vprn log)

Full Context

configure service vprn log

Description

Commands in this context configure event logging within a specific VPRN.

By default, the log events in a VPRN log are a subset of the complete set of possible log events in SR OS. See the **config**>**log**>**services-all-events** command for more details.

Platforms

7705 SAR Gen 2

log

Syntax

log *log-id*

no log

Context

[Tree] (config>filter>ipv6-filter>entry log)

[Tree] (config>filter>ip-filter>entry log)

Full Context

configure filter ipv6-filter entry log

configure filter ip-filter entry log

Description

This command associates a filter log to the current filter policy entry and therefore enables logging for that filter entry.

The filter log must exist before a filter entry can be enabled to use the filter log.

The **no** form of the command disables logging for the filter entry.

Default

no log

Parameters

log-id

Specifies the filter log ID expressed as a decimal integer.

Values 101 to 199

Platforms

7705 SAR Gen 2

log

Syntax

log *log-id* [**create**]

no log *log-id*

Context

[Tree] (config>filter log)

Full Context

configure filter log

Description

This command, creates a configuration context for the specified filter log if it does not exist, and enables the context to configure the specified filter log.

The **no** form of the command deletes the filter log. The log cannot be deleted if there are filter entries configured to write to the log. All filter entry logging associations need to be removed before the log can be deleted.

Default

log 101

Parameters

log-id

Specifies the filter log ID expressed as a decimal integer.

Values 101 to 199

create

This keyword is required to create the configuration context. After it is created, the context can be enabled with or without the **create** keyword.

Platforms

7705 SAR Gen 2

log

Syntax

[**no**] **log**

Context

[Tree] (config>router>static-route-entry>indirect>cpe-check log)

[Tree] (config>router>static-route-entry>next-hop>cpe-check log)

Full Context

configure router static-route-entry indirect cpe-check log

configure router static-route-entry next-hop cpe-check log

Description

This optional parameter enables the ability to log transitions between active and in-active based on the CPE connectivity check. Events will be sent to the system log, syslog and SNMP traps.

Default

no log

Platforms

7705 SAR Gen 2

log

Syntax

[**no**] log

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry log)

[Tree] (config>system>security>mgmt-access-filter>ip-filter>entry log)

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter>entry log)

Full Context

configure system security management-access-filter mac-filter entry log

configure system security management-access-filter ip-filter entry log

configure system security management-access-filter ipv6-filter entry log

Description

This command enables match logging. When enabled, matches on this entry will cause the Security event mafEntryMatch to be raised.

Default

no log

Platforms

7705 SAR Gen 2

16.113 log-events

log-events

Syntax

log-events [**verbose**]

no log-events

Context

[Tree] (config>sys>security>dist-cpu-protection>policy>local-monitoring-policer log-events)

Full Context

configure system security dist-cpu-protection policy local-monitoring-policer log-events

Description

This command controls the creation of log events related to **local-monitoring-policer** status and activity.

Default

log-events

Parameters

verbose

Sends the same events as just "log-events” plus DcpLocMonExcd, DcpLocMonExcdAllDynAlloc, and DcpLocMonExcdAllDynFreed. The optional "verbose” includes some events that are more likely used during debug/tuning/investigations

Platforms

7705 SAR Gen 2

log-events

Syntax

[**no**] **log-events** [**verbose**]

no log-events

Context

[Tree] (config>sys>security>dist-cpu-protection>policy>protocol>dyn-para log-events)

Full Context

configure system security dist-cpu-protection policy protocol dynamic-parameters log-events

Description

This command controls the creation of log events related to dynamic enforcement policer status and activity.

Default

log-events

Parameters

verbose

This parameter sends the same events as just "log-events" plus Hold Down Start, Hold Down End, DcpDynamicEnforceAlloc and DcpDynamicEnforceFreed events. This includes the allocation/de-allocation events (typically used for debug/tuning only – could be very noisy even when there is nothing much of concern).

Platforms

7705 SAR Gen 2

log-events

Syntax

log-events [**verbose**]

no log-events

Context

[Tree] (config>sys>security>dist-cpu-protection>policy>static-policer log-events)

Full Context

configure system security dist-cpu-protection policy static-policer log-events

Description

This command controls the creation of log events related to static-policer status and activity.

Default

log-events

Parameters

verbose

Sends the same events as just "log-events" plus Hold Down Start and Down End events. The optional "verbose" includes some events that are more likely used during debug/tuning/investigations.

Platforms

7705 SAR Gen 2

16.114 log-files-total-size

log-files-total-size

Syntax

log-files-total-size *megabytes*

no log-files-total-size

Context

[Tree] (config>log>storage log-files-total-size)

Full Context

configure log file-storage-control log-files-total-size

Description

This command configures the limit for the total space that all log files can occupy on each storage device on the active CPM.

When this threshold is reached, log events are no longer written to the files in the `\log` directory until SR OS removes older log files and the occupancy is below the limit.

When unconfigured, there is no specific limit for the total size of all log files.

Only log files in the `\log` directory with system generated names (including no file extension) are applicable toward the total size limit.

If a user manually adds or deletes log files from the `\log` directory, the size of the files is not taken into account for up to 1 hour.

The configured total size limit is not validated against the actual size of the installed storage devices. If the configured limit is larger than the installed CF device, the limit is never reached.

Default

no log-files-total-size

Parameters

megabytes

Specifies the total size limit for log files, in MB.

Values 50 to 4,194,304 MB (4 TBytes, 2^{22} MB)

Default 0

Platforms

7705 SAR Gen 2

16.115 log-filter

log-filter

Syntax

log-filter *filter-id*

no log-filter

Context

[Tree] (config>log>event-trigger>event>trigger-entry log-filter)

Full Context

configure log event-trigger event trigger-entry log-filter

Description

This command configures the log filter to be used for this trigger entry. The log filter defines the matching criteria that must be met in order for the log event to trigger the handler execution. The log filter is applied to the log event and, if the filtering decision results in a forward action, then the handler is triggered.

It is typically unnecessary to configure match criteria for the application or number in the log filter used for EHS since the particular filter is only applied for a specific log event application and number, as configured under the **config**>**log**>**event-trigger** context.

The **no** form of this command removes the log filter configuration.

Parameters

filter-id

Specifies the identifier of the filter.

Values 1 to 1500

Platforms

7705 SAR Gen 2

16.116 log-id

log-id

Syntax

log-id *log-id* [**name** *log-name*]

no log-id *log-id*

Context

[Tree] (config>service>vprn>log log-id)

Full Context

configure service vprn log log-id

Description

This command creates a context to configure destinations for event streams.

The **log-id** context is used to direct events, alarms or traps, and debug information to respective destinations.

A maximum of 30 logs can be configured.

Before an event can be associated with this log-id, the **from** command identifying the source of the event must be configured.

Only one destination can be specified for a *log-id*. The destination of an event stream can be an in-memory buffer, console, session, snmp-trap-group, Syslog, or file.

Use the **event-control** command to suppress the generation of events, alarms, and traps for all log destinations.

An event filter policy can be applied in the log-id context to limit which events, alarms, and traps are sent to the specified log-id.

By default, the log events in a VPRN log are a subset of the complete set of possible log events in SR OS. See the **config**>**log**>**services-all-events** command for more details.

The **no** form of this command deletes the log destination ID from the configuration.

Default

No log destinations are defined.

Parameters

log-id

Specifies the log ID number, expressed as a decimal integer.

Values 1 to 100

name log-name

Configures an optional log name, up to 64 characters, that can be used to refer to the log destination after it is created.

Platforms

7705 SAR Gen 2

log-id

Syntax

log-id *log-id* [**name** *log-name*]

no log-id *log-id*

Context

[Tree] (config>log log-id)

Full Context

configure log log-id

Description

This command creates a context to configure destinations for event streams.

The **log-id** context is used to direct events, alarms or traps, and debug information for specific destinations.

A maximum of 30 logs can be configured.

Before an event can be associated with this log ID, the **from** command identifying the source of the event must be configured.

Only one destination can be specified for a *log-id*. The destination of an event stream can be an in-memory buffer, console, session, snmp-trap-group, syslog, or file.

Use the **event-control** command to suppress the generation of events, alarms, and traps for all log destinations.

An event filter policy can be applied in the log-id context to limit which events, alarms, and traps are sent to the specified log-id.

Log-IDs 99 and 100 are created by the agent. Log-ID 99 captures all log messages. Log-ID 100 captures log messages with a severity level of major and above.

Note:
Log-ID 99 provides valuable information for the admin-tech file. Removing or changing the log configuration may hinder debugging capabilities. It is strongly recommended not to alter the configuration for Log-ID 99.

The **no** form of this command deletes the log destination ID from the configuration.

Parameters

log-id

Specifies the log ID, expressed as a decimal integer.

Values 1 to 101

name log-name

Configures an optional log name, up to 64 characters, that can be used to refer to the log destination after it is created.

Platforms

7705 SAR Gen 2

16.117 log-prefix

log-prefix

Syntax

log-prefix *log-prefix-string*

no log-prefix

Context

[Tree] (config>service>vprn>log>syslog log-prefix)

Full Context

configure service vprn log syslog log-prefix

Description

This command adds the string prepended to every syslog message sent to the syslog host.

RFC 3164, *The BSD syslog Protocol,* allows an alphanumeric string (tag) to be prepended to the content of every log message sent to the syslog host. This alphanumeric string can, for example, be used to identify the node that generates the log entry. The software appends a colon (:) and a space to the string and it is inserted in the syslog message after the date stamp and before the syslog message content.

Only one string can be entered. If multiple strings are entered, the last string overwrites the previous string. The alphanumeric string can contain lowercase (a-z), uppercase (A-Z) and numeric (0-9) characters.

The **no** form of this command removes the log prefix string.

Default

log-prefix "TMNX".

Parameters

log-prefix-string

Specifies the alphanumeric string of up to 32 characters. Spaces and colons (:) cannot be used in the string.

Platforms

7705 SAR Gen 2

log-prefix

Syntax

log-prefix *log-prefix-string*

no log-prefix

Context

[Tree] (config>log>syslog log-prefix)

Full Context

configure log syslog log-prefix

Description

This command adds the string prepended to every syslog message sent to the syslog host.

RFC 3164, allows an alphanumeric string (tag) to be prepended to the content of every log message sent to the syslog host. This alphanumeric string can, for example, be used to identify the node that generates the log entry. The software appends a colon (:) and a space to the string and it is inserted in the syslog message after the date stamp and before the syslog message content.

Only one string can be entered. If multiple strings are entered, the last string overwrites the previous string. The alphanumeric string can contain lowercase (a-z), uppercase (A-Z) and numeric (0 to 9) characters.

The **no** form of this command removes the log prefix string.

Default

no log-prefix

Parameters

log-prefix-string

Specifies an alphanumeric string up to 32 characters in length. Spaces and colons (:) cannot be used in the string.

Platforms

7705 SAR Gen 2

16.118 logger-event-bundling

logger-event-bundling

Syntax

[**no**] **logger-event-bundling**

Context

[Tree] (config>router>mpls logger-event-bundling)

Full Context

configure router mpls logger-event-bundling

Description

This feature merges two of the most commonly generated MPLS traps, vRtrMplsXCCreate and vRtrMplsXCDelete, which can be generated at both LER and LSR into a new specific trap vRtrMplsSessionsModified. In addition, this feature perform bundling of traps of multiple RSVP sessions, that is LSPs, into this new specific trap.

The intent is to provide a tool for the user to minimize trap generation in an MPLS network. Note that the MPLS trap throttling will not be applied to this new trap.

The **no** version of this command disables the merging and bundling of the above MPLS traps.

Platforms

7705 SAR Gen 2

16.119 login-banner

login-banner

Syntax

[**no**] **login-banner**

Context

[Tree] (config>system>login-control login-banner)

Full Context

configure system login-control login-banner

Description

This command enables or disables the display of a login banner. The login banner contains the SR OS copyright and build date information for a console login attempt.

The **no** form of this command causes only the configured pre-login-message and a generic login prompt to display.

Platforms

7705 SAR Gen 2

16.120 login-control

login-control

Syntax

login-control

Context

[Tree] (config>system login-control)

Full Context

configure system login-control

Description

This command creates the context to configure the session control for console, Telnet, SSH, and FTP sessions.

Platforms

7705 SAR Gen 2

16.121 login-exec

login-exec

Syntax

login-exec *url-prefix*: *source-url*

no login-exec

Context

[Tree] (config>system>security>user>console login-exec)

[Tree] (config>system>security>user-template>console login-exec)

Full Context

configure system security user console login-exec

configure system security user-template console login-exec

Description

This command configures a user's login exec file which executes whenever the user successfully logs in to a console session.

Only one exec file can be configured. If multiple **login-exec** commands are entered for the same user, each subsequent entry overwrites the previous entry.

The **no** form of this command disables the login exec file for the user.

Default

no login-exec

Parameters

url-prefix: source-url

Specifies either a local or remote URL, up to 200 characters, that identifies the exec file that is executed after the user successfully logs in.

Platforms

7705 SAR Gen 2

16.122 login-scripts

login-scripts

Syntax

login-scripts

Context

[Tree] (config>system>login-control login-scripts)

Full Context

configure system login-control login-scripts

Description

Commands in this context configure CLI scripts that execute when a user (authenticated via any method including local user database, TACACS+, or RADIUS) first logs into a CLI session.

Platforms

7705 SAR Gen 2

16.123 logout

logout

Syntax

logout

Context

[Tree] (logout)

Full Context

logout

Description

This command logs out of the router session.

When the **logout** command is issued from the console, the login prompt is displayed, and any log IDs directed to the console are discarded. When the console session resumes (regardless of the user), the log output to the console resumes.

When a Telnet session is terminated from a **logout** command, all log IDs directed to the session are removed. When a user logs back in, the log IDs must be re-created.

Platforms

7705 SAR Gen 2

16.124 long-lived

long-lived

Syntax

[**no**] **long-lived**

Context

[Tree] (config>service>vprn>bgp>graceful-restart long-lived)

[Tree] (config>service>vprn>bgp>group>graceful-restart long-lived)

[Tree] (config>service>vprn>bgp>group>neighbor>graceful-restart long-lived)

Full Context

configure service vprn bgp graceful-restart long-lived

configure service vprn bgp group graceful-restart long-lived

configure service vprn bgp group neighbor graceful-restart long-lived

Description

Commands in this context configure BGP Long-Lived Graceful-Restart (LLGR) procedures.

LLGR, known informally as BGP persistence, is an extension of BGP graceful restart that allows a session to stay down for a longer period of time. During this time, learned routes are marked and re-advertised as stale but they can continue to be used as routes of last resort.

The LLGR handling of a session failure can be invoked immediately or it can be delayed until the end of the traditional GR restart window.

Default

no long-lived

Platforms

7705 SAR Gen 2

long-lived

Syntax

[**no**] **long-lived**

Context

[Tree] (config>router>bgp>graceful-restart long-lived)

[Tree] (config>router>bgp>group>graceful-restart long-lived)

[Tree] (config>router>bgp>group>neighbor>graceful-restart long-lived)

Full Context

configure router bgp graceful-restart long-lived

configure router bgp group graceful-restart long-lived

configure router bgp group neighbor graceful-restart long-lived

Description

Commands in this context enter commands related to BGP Long-Lived Graceful-Restart (LLGR) procedures.

LLGR, known informally as BGP persistence, is an extension of BGP GR that allows a session to stay down for a longer period of time. During this time, learned routes are marked and re-advertised as stale but they can continue to be used as routes of last resort.

The LLGR handling of a session failure can be invoked immediately or it can be delayed until the end of the traditional GR restart window.

Default

no long-lived

Platforms

7705 SAR Gen 2

16.125 loop-detect

loop-detect

Syntax

loop-detect {**drop-peer** | **discard-route** | **ignore-loop** | **off**}

no loop-detect

Context

[Tree] (config>service>vprn>bgp loop-detect)

[Tree] (config>service>vprn>bgp>group>neighbor loop-detect)

[Tree] (config>service>vprn>bgp>group loop-detect)

Full Context

configure service vprn bgp loop-detect

configure service vprn bgp group neighbor loop-detect

configure service vprn bgp group loop-detect

Description

This command configures how the BGP peer session handles loop detection in the AS path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

Dynamic configuration changes of **loop-detect** are not recognized.

The **no** form of this command used at the global level reverts to default, which is **loop-detect ignore-loop**.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

loop-detect ignore-loop

Parameters

drop-peer

Sends a notification to the remote peer and drops the session.

discard-route

Discards routes received with loops in the AS path.

ignore-loop

ignores routes with loops in the AS path but maintains peering.

off

Disables loop detection.

Platforms

7705 SAR Gen 2

loop-detect

Syntax

loop-detect {**drop-peer** | **discard-route** | **ignore-loop** | **off**}

no loop-detect

Context

[Tree] (config>router>bgp>group loop-detect)

[Tree] (config>router>bgp>group>neighbor loop-detect)

[Tree] (config>router>bgp loop-detect)

Full Context

configure router bgp group loop-detect

configure router bgp group neighbor loop-detect

configure router bgp loop-detect

Description

This command configures how the BGP peer session handles loop detection in the AS path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

Note:
Dynamic configuration changes of **loop-detect** are not recognized.

The **no** form of this command used at the global level reverts to default, which is **loop-detect ignore-loop**.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

loop-detect ignore-loop

Parameters

drop-peer

Sends a notification to the remote peer and drops the session.

discard-route

Discards routes received from a peer with the same AS number as the router itself. This option prevents routes looped back to the router from being added to the routing information base and consuming memory. When this option is changed, the change will not be active for an established peer until the connection is re-established for the peer.

ignore-loop

Ignores routes with loops in the AS path but maintains peering.

off

Disables loop detection.

Platforms

7705 SAR Gen 2

16.126 loop-detect-threshold

loop-detect-threshold

Syntax

loop-detect-threshold *loop-detect-threshold*

no loop-detect-threshold

Context

[Tree] (config>service>vprn>bgp>group>neighbor loop-detect-threshold)

[Tree] (config>service>vprn>bgp loop-detect-threshold)

[Tree] (config>service>vprn>bgp>group loop-detect-threshold)

Full Context

configure service vprn bgp group neighbor loop-detect-threshold

configure service vprn bgp loop-detect-threshold

configure service vprn bgp group loop-detect-threshold

Description

This command provides additional control over the behavior enabled by the **loop-detect** command. If this command specifies a threshold value of *n,* then a route received by the local BGP speaker with an AS path

that contains up to *n* occurrences of the local speaker's AS number is considered valid and not treated as an AS path loop. An AS loop is considered to occur only when the received AS path has more than *n* occurrences of the local speaker's AS number.

The **no** form of this command removes the configuration and sets the value to 0. One or more occurrence of the local speaker's AS number in the received AS path triggers the **loop-detect** behavior.

Default

no loop-detect-threshold

Parameters

loop-detect-threshold

The maximum number of occurrences of the local speaker's AS number in the received AS path before the AS path is considered to be a loop.

Values 0 to 15

Default 0

Platforms

7705 SAR Gen 2

loop-detect-threshold

Syntax

loop-detect-threshold *loop-detect-threshold*

no loop-detect-threshold

Context

[Tree] (config>router>bgp loop-detect-threshold)

[Tree] (config>router>bgp>group>neighbor loop-detect-threshold)

[Tree] (config>router>bgp>group loop-detect-threshold)

Full Context

configure router bgp loop-detect-threshold

configure router bgp group neighbor loop-detect-threshold

configure router bgp group loop-detect-threshold

Description

This command provides additional control over the behavior enabled by the **loop-detect** command. If this command specifies a threshold value of *n*, then a route received by the local BGP speaker with an AS path that contains up to *n* occurrences of the local speaker's AS number is considered valid and not treated as an AS path loop. An AS loop is considered to occur only when the received AS path has more than *n* occurrences of the local speaker's AS number.

The **no** form of this command removes the configuration and sets the value to 0. One or more occurrence of the local speaker's AS number in the received AS path triggers the **loop-detect** behavior.

Default

no loop-detect-threshold

Parameters

loop-detect-threshold

The maximum number of occurrences of the local speaker's AS number in the received AS path before the AS path is considered to be a loop.

Values 0 to 15

Default 0

Platforms

7705 SAR Gen 2

16.127 loopback

loopback

Syntax

[**no**] **loopback**

Context

[Tree] (config>service>vprn>if loopback)

[Tree] (config>service>ies>if loopback)

Full Context

configure service vprn interface loopback

configure service ies interface loopback

Description

This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated IES/VPRN interface cannot be bound to a SAP.

Note:
Configure an IES interface as a loopback interface by issuing the **loopback** command instead of the **sap** *sap-id* command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

The **no** form of this command reverts to the default.

Platforms

7705 SAR Gen 2

loopback

Syntax

[**no**] **loopback**

Context

[Tree] (config>service>vprn>nw-if loopback)

Full Context

configure service vprn network-interface loopback

Description

This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated interface cannot be bound to a SAP.

When using mtrace/mstat in a Layer 3 VPN context then the configuration for the VPRN should have a loopback address configured which has the same address as the core instance's system address (BGP next-hop).

Default

no loopback

Platforms

7705 SAR Gen 2

loopback

Syntax

[**no**] **loopback**

Context

[Tree] (config>router>if loopback)

Full Context

configure router interface loopback

Description

This command configures the interface as a loopback interface. The **vas-if-type** and **loopback** commands are mutually exclusive.

Default

Not enabled

Platforms

7705 SAR Gen 2

16.128 loopfree-alternate-exclude

loopfree-alternate-exclude

Syntax

[**no**] **loopfree-alternate-exclude**

Context

[Tree] (config>service>vprn>isis>level loopfree-alternate-exclude)

[Tree] (config>service>vprn>isis>interface loopfree-alternate-exclude)

[Tree] (config>router>isis>interface loopfree-alternate-exclude)

[Tree] (config>router>isis>level loopfree-alternate-exclude)

Full Context

configure service vprn isis level loopfree-alternate-exclude

configure service vprn isis interface loopfree-alternate-exclude

configure router isis interface loopfree-alternate-exclude

configure router isis level loopfree-alternate-exclude

Description

This command instructs IGP to not include a specific interface or all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.

When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the above OSPF command can only be executed under the area in which the specified interface is primary and once enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

The **no** form of this command re-instates the default value for this command.

Default

no loopfree-alternate-exclude

Platforms

7705 SAR Gen 2

loopfree-alternate-exclude

Syntax

[**no**] **loopfree-alternate-exclude**

Context

[Tree] (config>service>vprn>ospf3>area loopfree-alternate-exclude)

[Tree] (config>service>vprn>ospf>area>if loopfree-alternate-exclude)

[Tree] (config>service>vprn>ospf3>area>if loopfree-alternate-exclude)

[Tree] (config>service>vprn>ospf>area loopfree-alternate-exclude)

Full Context

configure service vprn ospf3 area loopfree-alternate-exclude

configure service vprn ospf area interface loopfree-alternate-exclude

configure service vprn ospf3 area interface loopfree-alternate-exclude

configure service vprn ospf area loopfree-alternate-exclude

Description

This command instructs IGP to not include a specific interface or all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.

When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the above OSPF command can only be executed under the area in which the specified interface is primary and once enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command fails.

The **no** form of this command re-instates the default value for this command.

Default

no loopfree-alternate-exclude

Platforms

7705 SAR Gen 2

loopfree-alternate-exclude

Syntax

[**no**] **loopfree-alternate-exclude**

Context

[Tree] (config>router>ospf>area loopfree-alternate-exclude)

[Tree] (config>router>ospf>area>interface loopfree-alternate-exclude)

[Tree] (config>router>ospf3>area loopfree-alternate-exclude)

[Tree] (config>router>ospf3>area>interface loopfree-alternate-exclude)

Full Context

configure router ospf area loopfree-alternate-exclude

configure router ospf area interface loopfree-alternate-exclude

configure router ospf3 area loopfree-alternate-exclude

configure router ospf3 area interface loopfree-alternate-exclude

Description

This command instructs IGP to not include a specific interface or all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.

When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the above OSPF command can only be executed under the area in which the specified interface is primary and once enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

The **no** form of this command re-instates the default value for this command.

Default

no loopfree-alternate-exclude

Platforms

7705 SAR Gen 2

16.129 loopfree-alternates

loopfree-alternates

Syntax

[**no**] **loopfree-alternates**

Context

[Tree] (config>service>vprn>isis loopfree-alternates)

Full Context

configure service vprn isis loopfree-alternates

Description

This command enables Loop-Free Alternate (LFA) computation by SPF under the IS-IS routing protocol level or under the OSPF routing protocol instance level.

When this command is enabled, it instructs the IGP SPF to attempt to pre-compute both a primary next-hop and an LFA next-hop for every learned prefix. When found, the LFA next-hop is populated into the routing table along with the primary next-hop for the prefix.

The **no** form of this command disables the LFA computation by IGP SPF.

Default

no loopfree-alternates

Platforms

7705 SAR Gen 2

loopfree-alternates

Syntax

[**no**] **loopfree-alternates**

Context

[Tree] (config>service>vprn>ospf3 loopfree-alternates)

[Tree] (config>service>vprn>ospf loopfree-alternates)

Full Context

configure service vprn ospf3 loopfree-alternates

configure service vprn ospf loopfree-alternates

Description

This command enables Loop-Free Alternate (LFA) computation by SPF under the IS-IS routing protocol level, or under the OSPF routing protocol instance level.

When this command is enabled, it instructs the IGP SPF to attempt to pre-compute both a primary next-hop and an LFA next-hop for every learned prefix. IS-IS computes the primary SPF first and then computes the LFA SPF. The LFA backup next-hop is only available after the LFA SPF is completed. When found, the LFA next-hop is populated into the routing table along with the primary next-hop for the prefix.

The **no** form of this command disables the LFA computation by IGP SPF.

Default

no loopfree-alternates

Platforms

7705 SAR Gen 2

loopfree-alternates

Syntax

[**no**] **loopfree-alternates**

Context

[Tree] (config>router>isis>flex-algos>flex-algo loopfree-alternates)

[Tree] (config>router>ospf>flex-algos>flex-algo loopfree-alternates)

Full Context

configure router isis flexible-algorithms flex-algo loopfree-alternates

configure router ospf flexible-algorithms flex-algo loopfree-alternates

Description

This command enables the advertisement of flexible-algorithm aware loop free alternates (LFAs).

The flexible algorithm LFA configuration (for example, LFA, remote-LFA or TI-LFA) inherits the LFA configuration for base SPF algorithm 0.

LFAs are administratively disabled for flexible algorithms in which IS-IS is participating. LFAs must be explicitly enabled using the **loopfree-alternates** command.

The **no** form of this command disables LFAs for the specific flexible algorithm in which the router is participating.

Default

no loopfree-alternates

Platforms

7705 SAR Gen 2

loopfree-alternates

Syntax

[**no**] **loopfree-alternates**

Context

[Tree] (config>router>isis loopfree-alternates)

Full Context

configure router isis loopfree-alternates

Description

This command enables Loop-Free Alternate (LFA) computation by SPF for the IS-IS routing protocol.

When this command is enabled, it instructs the IGP SPF to attempt to pre-compute both a primary nexthop and an LFA next-hop for every learned prefix. When found, the LFA next-hop is populated into the routing table along with the primary next-hop for the prefix.

The user enables the remote LFA next-hop calculation by the IGP LFA SPF by appending the remote-lfa option. When this option is enabled in an IGP instance, SPF performs the remote LFA additional computation following the regular LFA next-hop calculation when the latter resulted in no protection for one or more prefixes which are resolved to a given interface.

Remote LFA extends the protection coverage of LFA-FRR to any topology by automatically computing and establishing/tearing-down shortcut tunnels, also referred to as repair tunnels, to a remote LFA node which puts the packets back into the shortest without looping them back to the node which forwarded them over the repair tunnel. The remote LFA node is referred to as PQ node. A repair tunnel can in theory be an RSVP LSP, a LDP-in-LDP tunnel, or a SR tunnel. In this feature, it is restricted to use SR repair tunnel to the remote LFA node.

The remote LFA algorithm is a per-link LFA SPF calculation and not a per-prefix like the regular LFA one. So, it provides protection for all destination prefixes which share the protected link by using the neighbor on the other side of the protected link as a proxy for all these destinations.

The Topology-Independent LFA (TI-LFA) further improves the protection coverage of a network topology by computing and automatically instantiating a repair tunnel to a Q node which is not in shortest path from the computing node. The repair tunnel uses shortest path to the P node and a source routed path from the P node to the Q node.

In addition, the TI-LFA algorithm selects the backup path which matches the post-convergence path. This helps the capacity planning in the network since traffic will always flow on the same path when transitioning to the FRR next-hop and then onto the new primary next-hop.

At a high level, the TI-LFA protection algorithm is searching for a candidate P-Q set separated with a number of hops such that the label stack size does not exceed the value of **ti-lfa max-sr-frr-labels**, on each of the post-convergence paths to each destination node or prefix D.

When the **ti-lfa** option is enabled in IS-IS, it provides TI-LFA node-protect or link-protect backup path in IS-IS MT=0 for an SR-ISIS IPV4/IPv6 tunnel (node SID and adjacency SID), for an IPv4 SR-TE LSP, and for LDP IPv4 FEC when the LDP **fast-reroute backup-sr-tunnel** option is enabled.

The **max-sr-frr-labels** parameter is used to limit the search for the TI-LFA backup next-hop:

1. 0 — The IGP LFA SPF restricts the search to TI-LFA backup next-hop which does not require a repair tunnel, meaning that P node and Q node are the same and match a neighbor. This is also the case when both P and Q node match the advertising router for a prefix.
2. 1 to 3 — The IGP LFA SPF will widen the search to include a repair tunnel to a P node which itself is connected to the Q nodes with a 0-to-2 hops for a total of maximum of three labels: one node SID to P node and two adjacency SIDs from P node to the Q node. If the P node is a neighbor of the computing node, its node SID is compressed and meaning that up to three adjacency SIDs can separate the P and Q nodes.
3. 2 (default) — Corresponds to a repair tunnel to a non-adjacent P which is adjacent to the Q node. If the P node is a neighbor of the computing node, then the node SID of the P node is compressed and the default value of two labels corresponds to two adjacency SIDs between the P and Q nodes.

When the **node-protect** command is enabled, the router will prefer a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node. This

node-protect backup protects against the failure of a downstream node in the path of the prefix of a node SID except for the node owner of the node SID.

The parameter **max-pq-nodes** in Remote LFA controls the maximum number of PQ nodes found in the LFA SPFs for which the node protection check is performed. The node-protect condition means the router must run the original Remote LFA algorithm plus one extra forward SPF on behalf of each PQ node found, potentially after applying the **max-pq-cost** parameter, to check if the path from the PQ node to the destination does not traverse the protected node. Setting this parameter to a lower value means the LFA SPFs will use less computation time and resources but may result in not finding a node-protect repair tunnel.

The **no** form of this command disables the LFA computation by IGP SPF.

Default

no loopfree-alternates

Platforms

7705 SAR Gen 2

loopfree-alternates

Syntax

[**no**] **loopfree-alternates**

Context

[Tree] (config>router>ospf loopfree-alternates)

[Tree] (config>router>ospf3 loopfree-alternates)

Full Context

configure router ospf loopfree-alternates

configure router ospf3 loopfree-alternates

Description

This command enables Loop-Free Alternate (LFA) computation by SPF under the OSPF or OSPFv3 routing protocol instance.

When this command is enabled, it instructs the IGP SPF to attempt to precalculate both a primary next hop and an LFA next hop for every learned prefix. When found, the LFA next hop is populated into the routing table along with the primary next hop for the prefix.

The user enables the remote LFA next hop calculation by the IGP LFA SPF by appending the **remote-lfa** option. When this option is enabled in an IGP instance, SPF performs the remote LFA additional computation following the regular LFA next hop calculation when the latter resulted in no protection for one or more prefixes which are resolved to a particular interface.

Remote LFA extends the protection coverage of LFA-FRR to any topology by automatically computing and establishing or tearing down shortcut tunnels, also referred to as repair tunnels, to a remote LFA node that puts the packets back into the shortest path without looping them back to the node that forwarded them over the repair tunnel. The remote LFA node is referred to as a PQ node. A repair tunnel can, in theory, be

an RSVP-TE LSP, an LDP-in-LDP tunnel, or a segment routing (SR) tunnel. In this command, **remote-lfa** is restricted to using an SR repair tunnel to the remote LFA node.

The remote LFA algorithm is a per-link LFA SPF calculation and not a per-prefix calculation like the regular LFA algorithm. The remote LFA algorithm provides protection for all destination prefixes that share the protected link by using the neighbor on the other side of the protected link as a proxy for all the destinations.

The Topology-Independent LFA (TI-LFA) further improves the protection coverage of a network topology by computing and automatically instantiating a repair tunnel to a Q node which is not in shortest path from the computing node. The repair tunnel uses shortest path to the P node and a source routed path from the P node to the Q node.

In addition, the TI-LFA algorithm selects the backup path which matches the post-convergence path. This helps the capacity planning in the network since traffic will always flow on the same path when transitioning to the FRR next hop and then onto the new primary next hop.

At a high level, the TI-LFA protection algorithm is searching for a candidate P-Q set separated with a number of hops such that the label stack size does not exceed the value of **ti-lfa max-sr-frr-labels**, on each of the post-convergence paths to each destination node or prefix D.

When the **ti-lfa** option is enabled in OSPF, it provides TI-LFA node-protect or link-protect backup path for a SR-OSPF IPV4 tunnel (node SID and adjacency SID), and for a IPv4 SR-TE LSP.

The **max-sr-frr-labels** parameter is used to limit the search for the TI-LFA backup next hop:

1. 0 — The IGP LFA SPF restricts the search to TI-LFA backup next hop which does not require a repair tunnel, meaning that P node and Q node are the same and match a neighbor. This is also the case when both P and Q node match the advertising router for a prefix.
2. 1 to 3 — The IGP LFA SPF will widen the search to include a repair tunnel to a P node which itself is connected to the Q nodes with a 0-to-2 hops for a total of maximum of three labels: one node SID to P node and two adjacency SIDs from P node to the Q node. If the P node is a neighbor of the computing node, its node SID is compressed and meaning that up to three adjacency SIDs can separate the P and Q nodes.
3. 2 (default) — Corresponds to a repair tunnel to a non-adjacent P which is adjacent to the Q node. If the P node is a neighbor of the computing node, then the node SID of the P node is compressed and the default value of two labels corresponds to two adjacency SIDs between the P and Q nodes.

The TI-LFA repair tunnel can have a maximum of three labels pushed in addition to the label of the destination node or prefix. The user can set a lower maximum value for the additional FRR labels by configuring the CLI option **max-sr-frr-labels** *labels*. The default value is 2.

When the **node-protect** command is enabled, the router will prefer a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node. This node-protect backup protects against the failure of a downstream node in the path of the prefix of a node SID except for the node owner of the node SID.

The parameter **max-pq-nodes** in Remote LFA controls the maximum number of PQ nodes found in the LFA SPFs for which the node protection check is performed. The node-protect condition means the router must run the original Remote LFA algorithm plus one extra forward SPF on behalf of each PQ node found, potentially after applying the **max-pq-cost** parameter, to check if the path from the PQ node to the destination does not traverse the protected node. Setting this parameter to a lower value means the LFA SPFs will use less computation time and resources but may result in not finding a node-protect repair tunnel.

The **no** form of this command disables the LFA computation by the IGP SPF.

Default
no loopfree-alternates

Platforms
7705 SAR Gen 2

16.130 loss

loss

Syntax
loss

Context
[Tree] (config>oam-pm>session>ip>twamp-light loss)

Full Context
configure oam-pm session ip twamp-light loss

Description
Commands in this context configure loss parameters for the TWAMP-Light test.

Platforms
7705 SAR Gen 2

16.131 loss-event

loss-event

Syntax
loss-event rising-threshold *threshold* [**falling-threshold** *threshold*] [*direction*]
no loss-event

Context
[Tree] (config>saa>test loss-event)

Full Context
configure saa test loss-event

Description

Specifies that at the termination of an SAA test run, the calculated loss event value is evaluated against the configured rising and falling loss event thresholds. SAA threshold events are generated as required.

The configuration of loss event thresholds is optional.

The **no** form of this command disables the loss-event test run.

Parameters

rising-threshold *threshold*

Specifies a rising threshold loss event value, in packets. When the test run is completed, the calculated loss event value is compared to the configured loss event rising threshold. If the test run loss event value is greater than the configured rising threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Values 0 to 2147483647

Default 0

falling-threshold *threshold*

Specifies a falling threshold loss event value, in packets. When the test run is completed, the calculated loss event value is compared to the configured loss event falling threshold. If the test run loss event value is greater than the configured falling threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Values 0 to 2147483647

Default 0

direction

Specifies the direction for OAM ping responses received for an OAM ping test run.

Values **inbound** — Monitor the value of jitter calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run.

outbound — Monitor the value of jitter calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run.

roundtrip — Monitor the value of jitter calculated for the round trip, two-way, OAM ping requests and replies for an OAM ping test run.

Default roundtrip

Platforms

7705 SAR Gen 2

16.132 loss-events

loss-events

Syntax

[**no**] **loss-events**

Context

[Tree] (config>oam-pm>session>meas-intvl>event-mon loss-events)

Full Context

configure oam-pm session meas-interval event-mon loss-events

Description

This enables the monitoring of all configured loss events. Adding this functionality starts the monitoring of the configured loss events at the start of the next measurement interval. If the function is removed using the **no** command, all monitoring of configured loss events, logging, and recording of new events for that session are suspended. Any existing events at the time of the shutdown are maintained until the active measurement window in which the removal was performed has completed. The state of this monitoring function can be changed without having to shut down all the tests in the session.

The **no** form of this command disables the monitoring of all configured loss events.

Platforms

7705 SAR Gen 2

loss-events

Syntax

loss-events

Context

[Tree] (config>oam-pm>session>ip>twamp-light loss-events)

Full Context

configure oam-pm session ip twamp-light loss-events

Description

This context allows the operator to define the loss events and thresholds that are to be tracked.

Platforms

7705 SAR Gen 2

16.133 low

low

Syntax

low

Context

[Tree] (config>service>vpls>sap>ingress>queue-override>queue>drop-tail low)

[Tree] (config>service>ies>if>sap>ingress>queue-override>queue>drop-tail low)

[Tree] (config>service>vpls>sap>egress>queue-override>queue>drop-tail low)

[Tree] (config>service>ies>if>sap>egress>queue-override>queue>drop-tail low)

Full Context

configure service vpls sap ingress queue-override queue drop-tail low

configure service ies interface sap ingress queue-override queue drop-tail low

configure service vpls sap egress queue-override queue drop-tail low

configure service ies interface sap egress queue-override queue drop-tail low

Description

Commands in this context configure the queue **low drop-tail** parameters. The low drop tail defines the queue depth beyond which out-of-profile packets are not accepted into the queue and are discarded.

Platforms

7705 SAR Gen 2

low

Syntax

low

Context

[Tree] (config>port>eth>access>egr>qgrp>qover>q>drop-tail low)

[Tree] (config>port>eth>access>ing>qgrp>qover>q>drop-tail low)

[Tree] (config>port>ethernet>network>egr>qgrp>qover>q>drop-tail low)

Full Context

configure port ethernet access egress queue-group queue-overrides queue drop-tail low

configure port ethernet access ingress queue-group queue-overrides queue drop-tail low

configure port ethernet network egress queue-group queue-overrides queue drop-tail low

Description

Commands in this context configure the queue low drop tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets will not be accepted into the queue and will be discarded.

Platforms

7705 SAR Gen 2

low

Syntax

low

Context

[Tree] (config>service>epipe>sap>ingress>queue-override>queue>drop-tail low)

[Tree] (config>service>epipe>sap>egress>queue-override>queue>drop-tail low)

Full Context

configure service epipe sap ingress queue-override queue drop-tail low

configure service epipe sap egress queue-override queue drop-tail low

Description

Commands in this context configure the queue low drop-tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets are not accepted into the queue and will be discarded.

Platforms

7705 SAR Gen 2

low

Syntax

low

Context

[Tree] (config>service>vprn>if>sap>egress>queue-override>queue>drop-tail low)

[Tree] (config>service>vprn>if>sap>ingress>queue-override>queue>drop-tail low)

Full Context

configure service vprn interface sap egress queue-override queue drop-tail low

configure service vprn interface sap ingress queue-override queue drop-tail low

Description

Commands in this context configure the queue low drop tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets are not accepted into the queue and will be discarded.

Platforms

7705 SAR Gen 2

low

Syntax

low

Context

[Tree] (config>qos>sap-ingress>queue>drop-tail low)

[Tree] (config>qos>sap-egress>queue>drop-tail low)

Full Context

configure qos sap-ingress queue drop-tail low

configure qos sap-egress queue drop-tail low

Description

Commands in this context configure the queue low drop-tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets will not be accepted into the queue and will be discarded.

Platforms

7705 SAR Gen 2

low

Syntax

low

Context

[Tree] (config>qos>network-queue>queue>drop-tail low)

Full Context

configure qos network-queue queue drop-tail low

Description

Commands in this context configure the queue low drop tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets will not be accepted into the queue and will be discarded.

Platforms

7705 SAR Gen 2

low

Syntax

low

Context

[Tree] (config>qos>qgrps>ing>qgrp>queue>drop-tail low)

[Tree] (config>qos>qgrps>egr>qgrp>queue>drop-tail low)

Full Context

configure qos queue-group-templates ingress queue-group queue drop-tail low

configure qos queue-group-templates egress queue-group queue drop-tail low

Description

Commands in this context configure the queue low drop-tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets will not be accepted into the queue and will be discarded.

Platforms

7705 SAR Gen 2

16.134 low-octets-discarded-count

low-octets-discarded-count

Syntax

[**no**] **low-octets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>queue>i-counters low-octets-discarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters low-octets-discarded-count)

Full Context

configure log accounting-policy custom-record queue i-counters low-octets-discarded-count

configure log accounting-policy custom-record ref-queue i-counters low-octets-discarded-count

Description

This command includes the low octets discarded count.

The **no** form of this command excludes the low octets discarded count.

Default

no low-octets-discarded-count

Platforms

7705 SAR Gen 2

16.135 low-octets-offered-count

low-octets-offered-count

Syntax

[**no**] **low-octets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>queue>i-counters low-octets-offered-count)

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters low-octets-offered-count)

Full Context

configure log accounting-policy custom-record queue i-counters low-octets-offered-count

configure log accounting-policy custom-record ref-queue i-counters low-octets-offered-count

Description

This command includes the low octets discarded count.

The **no** form of this command excludes the low octets discarded count.

Default

no low-octets-offered-count

Platforms

7705 SAR Gen 2

16.136 low-packets-discarded-count

low-packets-discarded-count

Syntax

[**no**] **low-packets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>queue>i-counters low-packets-discarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters low-packets-discarded-count)

Full Context

configure log accounting-policy custom-record queue i-counters low-packets-discarded-count

configure log accounting-policy custom-record ref-queue i-counters low-packets-discarded-count

Description

This command includes the low packets discarded count.

The **no** form of this command excludes the low packets discarded count.

Default

no low-packets-discarded-count

Platforms

7705 SAR Gen 2

16.137 low-packets-offered-count

low-packets-offered-count

Syntax

[**no**] **low-packets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters low-packets-offered-count)

[Tree] (config>log>acct-policy>cr>queue>i-counters low-packets-offered-count)

Full Context

configure log accounting-policy custom-record ref-queue i-counters low-packets-offered-count

configure log accounting-policy custom-record queue i-counters low-packets-offered-count

Description

This command includes the low packets discarded count.

The **no** form of this command excludes the low packets discarded count.

Default

no low-packets-offered-count

Platforms

7705 SAR Gen 2

16.138 lower-bound

lower-bound

Syntax

lower-bound *microseconds*

no lower-bound

Context

[Tree] (config>oam-pm>bin-group>bin-type>bin lower-bound)

Full Context

configure oam-pm bin-group bin-type bin lower-bound

Description

This command allows the operator to specify the individual floors thresholds for the bins. The operator does not have to specific a lower threshold for every bin that was previously defined by the bin-count for the specific type. By default, each bin is the bin-number times 5000 microseconds. Lower thresholds in the previous adjacent bin must be lower than the threshold of the next higher bin threshold. A separate line per bin is required to configure an operator-specific threshold. An error prevents the bin from entering the active state if this is not maintained, at the time the **no shutdown** is issued. Bin 0 is the result of the difference between 0 and the configured lower-threshold of bin 1. The highest bin in the bin-count captures every result above the threshold. Any negative delay metric result is treated as zero and placed in bin 0.

The **no** form of this command removes the user configured threshold value and applies the default for the bin.

Parameters

microseconds

Specifies the threshold that defines the floor of the bin. The bin range is the difference between its configured threshold and the threshold of the next higher bin in microsecond threshold value.

Values 1 to 4294967295

Default bin-number * 5000

Platforms

7705 SAR Gen 2

16.139 lsa-accumulate

lsa-accumulate

Syntax

lsa-accumulate *lsa-accumulate*

no lsa-accumulate

Context

[Tree] (config>router>ospf3>timers lsa-accumulate)

[Tree] (config>router>ospf>timers lsa-accumulate)

Full Context

configure router ospf3 timers lsa-accumulate

configure router ospf timers lsa-accumulate

Description

This command sets the internal OSPF delay to allow for the accumulation of multiple LSA so OSPF messages can be sent as efficiently as possible. The **lsa-accumulate** timer applies to all LSAs except Type 1 and Type 2 LSAs, which are sent immediately. LSAs are accumulated and then sent when:

- its size reaches the MTU size of the interface
- a new LSA is received on the interface
- the **lsa-accumulate** timer expires

Shorting this delay can speed up the advertisement of LSAs to OSPF neighbors but may increase the number of OSPF messages sent.

The **no** form of this command reverts to the default value.

Note:
The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Default

lsa-accumulate 1000

Parameters

lsa-accumulate

Specifies the LSA accumulation delay in milliseconds.

Values 0 to 1000

Platforms

7705 SAR Gen 2

16.140 lsa-arrival

lsa-arrival

Syntax

lsa-arrival *lsa-arrival-time*

no lsa-arrival

Context

[Tree] (config>service>vprn>ospf3>timers lsa-arrival)

[Tree] (config>service>vprn>ospf>timers lsa-arrival)

Full Context

configure service vprn ospf3 timers lsa-arrival

configure service vprn ospf timers lsa-arrival

Description

This parameter defines the minimum delay that must pass between receipt of the same Link State Advertisements (LSAs) arriving from neighbors.

It is recommended that the neighbor's configured **lsa-generate** *lsa-second-wait* interval is equal to or greater than the **lsa-arrival** timer configured here.

Use the **no** form of this command to return to the default.

Note:
The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is ≥ 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Default

lsa-arrival 1000

Parameters

lsa-arrival-time

Specifies the timer in milliseconds.

Values 0 to 600000

Platforms

7705 SAR Gen 2

lsa-arrival

Syntax

lsa-arrival *lsa-arrival-time*

no lsa-arrival

Context

[Tree] (config>router>ospf>timers lsa-arrival)

[Tree] (config>router>ospf3>timers lsa-arrival)

Full Context

configure router ospf timers lsa-arrival

configure router ospf3 timers lsa-arrival

Description

This parameter defines the minimum delay that must pass between receipt of the same Link State Advertisements (LSAs) arriving from neighbors.

It is recommended that the neighbors configured **lsa-generate** *lsa-second-wait* interval is equal or greater than the **lsa-arrival** timer configured here.

The **no** form of this command reverts to the default.

Note:
The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Default

lsa-arrival 1000

Parameters

lsa-arrival-time

Specifies the timer, in milliseconds.

Values 0 to 600000

Platforms

7705 SAR Gen 2

16.141 lsa-filter-out

lsa-filter-out

Syntax

lsa-filter-out [**all** | **except-own-rtrlsa** | **except-own-rtrlsa-and-defaults**]

no lsa-filter-out

Context

[Tree] (config>service>vprn>ospf>area>if lsa-filter-out)

[Tree] (config>router>ospf>area>if lsa-filter-out)

[Tree] (config>service>vprn>ospf3>area>if lsa-filter-out)

[Tree] (config>router>ospf3>area>if lsa-filter-out)

Full Context

configure service vprn ospf area interface lsa-filter-out

configure router ospf area interface lsa-filter-out

configure service vprn ospf3 area interface lsa-filter-out

configure router ospf3 area interface lsa-filter-out

Description

This command enables filtering of outgoing OSPF LSAs on the selected OSPFv2 or OSPFv3 interface. Three filtering options are provided:

- Do not flood any LSAs out the interface. This option is suitable if the neighbor is simply-connected and has a statically configured default route with the address of this interface as next-hop.
- Flood the router's own router-LSA out the interface and suppress all other flooded LSAs. This option is suitable if the neighbor is simply-connected and has a statically configured default route with a loopback or system interface address (contained in the router-LSA) as next-hop.

- Flood the router's own router-LSA and all self-generated type-3, type-5 and type-7 LSAs advertising a default route (0/0) out the interface; suppress all other flooded LSAs. This option is suitable if the neighbor is simply-connected and does not have a statically configured default route.

The **no** form of this command disables OSPF LSA filtering (normal operation).

Default

no lsa-filter-out

Platforms

7705 SAR Gen 2

16.142 lsa-generate

lsa-generate

Syntax

lsa-generate *max-lsa-wait* [lsa-initial-wait *lsa-initial-wait* [*lsa-second-wait lsa-second-wait*]]

no lsa-generate-interval

Context

[Tree] (config>service>vprn>ospf3>timers lsa-generate)

[Tree] (config>service>vprn>ospf>timers lsa-generate)

Full Context

configure service vprn ospf3 timers lsa-generate

configure service vprn ospf timers lsa-generate

Description

This parameter customizes the throttling of OSPF LSA-generation. Timers that determine when to generate the first, second, and subsequent LSAs can be controlled with this command. Subsequent LSAs are generated at increasing intervals of the *lsa-second-wait* timer until a maximum value is reached.

Configuring the **lsa-arrival** interval to equal or less than the *lsa-second-wait* interval configured in the **lsa-generate** command is recommended.

The **no** form of this command reverts to the default.

Note:
The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is ≥ 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Parameters

max-lsa-wait

Specifies the maximum interval, in milliseconds, between two consecutive occurrences of an LSA being generated.

Values 10 to 600000

Default 5000

lsa-initial-wait

Specifies the first waiting period between link-state advertisements (LSA) originate(s), in milliseconds. When the LSA exceeds the **lsa-initial-wait** timer value and the topology changes, there is no wait period and the LSA is immediately generated.

When an LSA is generated, the initial wait period commences. If, within the specified **lsa-initial-wait** period and another topology change occurs, then the **lsa-initial-wait** timer applies.

Values 10 to 600000

Default 5000

lsa-second-wait

Specifies the hold time in milliseconds between the first and second LSA generation. The next topology change is subject to this second wait period. With each subsequent topology change, the wait time doubles (this is 2x the previous wait time). This assumes that each failure occurs within the relevant wait period.

Values 10 to 600000

Default 5000

Platforms

7705 SAR Gen 2

lsa-generate

Syntax

lsa-generate *max-lsa-wait* [**lsa-initial-wait** *lsa-initial-wait* [**lsa-second-wait** *lsa-second-wait*]]

no lsa-generate

Context

[Tree] (config>router>ospf3>timers lsa-generate)

[Tree] (config>router>ospf>timers lsa-generate)

Full Context

configure router ospf3 timers lsa-generate

configure router ospf timers lsa-generate

Description

This parameter customizes the throttling of OSPF LSA-generation. Timers that determine when to generate the first, second, and subsequent LSAs can be controlled with this command. Subsequent LSAs are generated at increasing intervals of the *lsa-second-wait* timer until a maximum value is reached.

Configuring the **lsa-arrival** interval to equal or less than the *lsa-second-wait* interval configured in the **lsa-generate** command is recommended.

The **no** form of this command reverts to the default.

Note:
The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Default

lsa-generate 5000

Parameters

max-lsa-wait

Specifies the maximum interval in milliseconds between two consecutive occurrences of an LSA being generated.

Values 10 to 600000

Default 5000

lsa-initial-wait

Specifies the first waiting period between link-state advertisements (LSA) originate(s), in milliseconds. When the LSA exceeds the **lsa-initial-wait** timer value and the topology changes, there is no wait period and the LSA is immediately generated.

When an LSA is generated, the initial wait period commences. If, within the specified **lsa-initial-wait** period and another topology change occurs, then the **lsa-initial-wait** timer applies.

Values 10 to 600000

Default 5000

lsa-second-wait

Specifies the hold time in milliseconds between the first and second LSA generation. The next topology change is subject to this second wait period. With each subsequent topology change, the wait time doubles (this is 2x the previous wait time). This assumes that each failure occurs within the relevant wait period.

Values 10 to 600000

Default 5000

Platforms

7705 SAR Gen 2

16.143 lsdb

lsdb

Syntax

[**no**] **lsdb** [*level-number*] [*system-id* | *lsp-id*]

Context

[Tree] (debug>router>isis lsdb)

Full Context

debug router isis lsdb

Description

This command enables debugging for Link State DataBase (LSDB).

The **no** form of the command disables debugging.

Parameters

system-id

When specified, only the specified system-id is debugged. Host name up to 38 characters.

lsp-id

When specified, only the specified lsp-id is debugged. Hostname up to 38 characters.

level-number

Specifies the interface level (1, 2, or 1 and 2).

Platforms

7705 SAR Gen 2

lsdb

Syntax

lsdb [*type*] [*ls-id*] [*adv-rtr-id*] [**area** *area-id*]

no lsdb

Context

[Tree] (debug>router>ospf3 lsdb)

[Tree] (debug>router>ospf lsdb)

Full Context

debug router ospf3 lsdb

debug router ospf lsdb

Description

This command enables debugging for an OSPF link-state database (LSDB).

Parameters

type

Specifies the OSPF link-state database (LSDB) type.

Values in the **ospf** context — router, network, summary, asbr, extern, nssa, area-opaque, as-opaque, link-opaque

in the **ospf3** context — router, network, inter-area-pfx, inter-area-rtr, external, nssa, intra-area-pfx, rtr-info-link, rtr-info-area, rtr-info-as

ls-id

Specifies an LSA type specific field containing either a router ID or an IP address. It identifies the piece of the routing domain being described by the advertisement.

adv-rtr-id

Specifies the router identifier of the router advertising the LSA.

area area-id

Specifies a 32-bit integer uniquely identifying an area.

Values ip-address — a.b.c.d

area — 0 to 4294967295

Platforms

7705 SAR Gen 2

16.144 lsn

lsn

Syntax

lsn router *router-instance* [**b4** *ipv6-address*] [**aftr** *ipv6-address*] **ip** *ip-address* **protocol** {**tcp** | **udp**} [**port** *port*] [**outside-ip** *ipv4-address*] [**outside-port** *port*] [**nat-policy** *nat-policy-name*] [*force*]

no lsn router *router-instance* [**b4** *ipv6-address*] **ip** *ip-address* **protocol** {**tcp** | **udp**} **port** *port* [**nat-policy** *nat-policy-name*]

Context

[Tree] (config>service>nat>fwd lsn)

Full Context

configure service nat port-forwarding lsn

Description

This command creates NAT static port forwards for LSN44, Ds-Lite and NAT64. Static port forwards (SPF) are static mappings created so that certain applications on the inside (private side) can be reached from host that are on the outside of the NAT. SPF statically map the subscriber (inside IP address in LSN44, CPE IPv6 address/prefix in DS-Lite and IPv6 prefix in NAT64), inside port and protocol to an outside IPv4 address, port and the same protocol.

If only the inside router, the inside IPv4/v6 address/prefix and the protocol are configured as parameters in the SPF request, the remaining fields in the mapping (outside port and outside IPv4 address) will be selected automatically by the node and reported in CLI once the command execution is completed.

Specifying the outside IPv4 address in the SPF request, mandates that all other, otherwise optional, parameters be also specified in the request (inside port and outside port). This creates a fully specified SPF request. Fully specified SPF request can be used in multi-chassis NAT redundancy deployments where the SPF is manually replicated between the SR OS nodes. In single chassis NAT deployments, fully specified SPF request is guaranteed to work only in the system with a single MS-ISA in it. Otherwise (multiple MS-ISAs in the system) a conflict may arise where two distinct inside IP addresses that may reside on separate MS-ISAs are requested to be mapped to the same outside IPv4 address. This will not be possible since the outside IPv4 address cannot be split across the MS-ISAs (each IP address, inside or outside, is tied to a single MS-ISA).

In non-fully specified SPF requests (missing the inside port and/or outside port and the outside IPv4 address within the SPF request), the outside IPv4 address selection will depend on the configuration of the outside port in the SPF request:

- If the outside port is not specified or is specified from the configured **port-forwarding-range** [1024..port-forwarding-range], then the outside IPv4 address will be the same as the outside IPv4 address in an existing dynamic mapping for the same subscriber. If the subscriber does not exist (no dynamic mappings exist at the time of SPF creation request), then the subscriber will be automatically created and an outside IPv4 address will be assigned. In case that the outside ports are not available from the outside IPv4 address of the corresponding dynamic mapping, then the SPF request will fail. In other words, the dynamic and static mappings (created in this manner) for the same subscriber must use the same outside IPv4 address.
- If the outside port from the well-known port range [0 to 1023] is requested, then the outside IPv4 address does not have to match the outside IPv4 address of an existing dynamic mapping for the same subscriber, but can instead be any outside IPv4 address.

If multiple NAT policies per inside routing context are used, then the NAT policy must be specified in the SPF creation request. This is needed so the SPF be created in the correct pool.

SPFs are disabled by default and they must be explicitly enabled by the **port-limits forwarding** command within the NAT policy.

Configured SPFs, unlike SPFs created with the **tools** commands, are preserved across reboots without having to configure persistency (**config>system>persistence>nat-port-forwarding**) since they are part of the configuration. When the pool is shutdown the SPFs are deactivated. When the pool is enabled (no shutdown), the SPFs (as created by the **tools** command or by configuration) are activated.

To avoid possible persistency related conflicts, SPFs can only be created using one method on a given node: either as configuration (the CLI **configure** branch) or using the **tools** command. For example: if a first SPF entry is created via CLI **tools** commands, the node prevents SPF creation via configuration (the CLI **configure** branch) and vice versa.

The **no** form of the command deletes NAT static port forwards for LSN44, Ds-Lite and NAT64.

Parameters

router *router-instance*

This mandatory parameter specifies the inside routing instance; router name or service-id.

Values router-name, service-id

b4 *ipv6-address*

This optional parameter specifies the IPv6 address of the B4 element in DS-Lite.

Values <ipv6-address> : ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x - [0..FFFF]H
d - [0..255]D

aftr *ipv6-address*

This optional parameter specifies IPv6 address of the AFTR element in DS-Lite.

Values <ip-address> : ipv4-address - a.b.c.d
ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x - [0..FFFF]H
d - [0..255]D

protocol {tcp | udp}

This mandatory parameter specifies the protocol to use, either TCP or UDP.

port *port*

This optional parameter specifies a source port.

Values 1 to 65535

outside-ip *ipv4-address*

This mandatory parameter specifies the outside IPv4 address. If the outside IPv4 address is specified, then all other optional parameters become mandatory.

outside-port *port*

This optional parameter specifies the outside port.

nat-policy *policy-name*

If multiple NAT policies are used inside the routing context, then the NAT policy should be specified in the SPF request so the SPF is created in the correct NAT pool. Otherwise, the default NAT policy from the inside routing context will be used.

force

Force allocation of static port forwards.

This command is only applicable to a LSN44 pool with flexible port allocations. Such pools allow interleaving of static port forwards with other dynamically allocated ports triggered by traffic flow.

If the requested outside port in the static port forward command is already occupied by a dynamically allocated port, using the force keyword preempts the dynamically allocated port and reassigns it as a port forwarding port.

During preemption, all the flows associated with the dynamically allocated port are terminated.

When the requested port forward is occupied by another static port and not a dynamic port, the command fails even if the force keyword is configured.

When a static port forwarding command requests an arbitrary outside port that is not specified in the command, the force keyword has an effect if the entire port space is already occupied. In this case, one of the dynamic ports is preempted.

Values Keyword

Platforms

7705 SAR Gen 2

16.145 lsp

lsp

Syntax

[**no**] **lsp** *lsp-name*

Context

[Tree] (config>router>ldp>targ-session>peer>tunneling lsp)

Full Context

configure router ldp targeted-session peer tunneling lsp

Description

This command configures a specific LSP destined to this peer and to be used for tunneling of LDP FEC over RSVP. A maximum of 4 RSVP LSPs can be explicitly used for tunneling LDP FECs to the T-LDP peer.

It is not necessary to specify any RSVP LSP in this context unless there is a need to restrict the tunneling to selected LSPs. All RSVP LSPs with a to address matching that of the T-LDP peer are eligible by default. The user can also exclude specific LSP names by using the ldp-over-rsvp exclude command in the **config>router>mpls>lsp** context.

Platforms

7705 SAR Gen 2

lsp

Syntax

[**no**] **lsp** *lsp-name* [**bypass-only** | **p2mp-lsp** | **mpls-tp** *src-tunnel-num* | **sr-te**]

Context

[Tree] (config>router>mpls lsp)

Full Context

configure router mpls lsp

Description

This command creates an LSP that is either signaled dynamically by the router, or a statically provisioned MPLS-TP LSP.

When the LSP is created, the egress router must be specified using the **to** command and at least one **primary** or **secondary** path must be specified for signaled LSPs, or at least one working path for MPLS-TP LSPs. All other statements under the LSP hierarchy are optional.

LSPs are created in the administratively down (**shutdown**) state.

The **no** form of this command deletes the LSP. All configuration information associated with this LSP is lost. The LSP must be administratively shutdown before it can be deleted. The LSP must also be unbound from all SDPs before it can be deleted.

Parameters

lsp-name

Specifies the name that identifies the LSP. The LSP name can be up to 64 characters long and must be unique.

bypass-only

Defines an LSP as a manual bypass LSP exclusively. When a path message for a new LSP requests bypass protection, the PLR first checks if a manual bypass tunnel satisfying the path constraints exists. If one if found, the router selects it. If no manual bypass tunnel is found, the router dynamically signals a bypass LSP in the default behavior. The CLI for this feature includes a knob that provides the user with the option to disable dynamic bypass creation on a per node basis.

p2mp-lsp

Defines an LSP as a point-to-multipoint LSP. The following parameters can be used with a P2MP LSP: adaptive, adspec, cspf, exclude, fast-reroute, from, hop-limit, include, metric, retry-limit, retry-timer, resignal-timer. The following parameters cannot be used with a P2MP LSP: primary, secondary, to, dest-global-id, dest-tunnel-number, working-tp-path, protect-tp-path.

mpls-tp ***src-tunnel-num***

Defines an LSP as an MPLS-TP LSP. The *src-tunnel-num* is a mandatory create time parameter for mpls-tp LSPs, and has to be assigned by the user based on the configured range of tunnel IDs. The following parameters can only be used with an MPLS-TP LSP: to, dest-global-id, dest-tunnel-number, working-tp-path, protect-tp-path. Other parameters defined for the above LSP types cannot be used.

sr-te

Defines an LSP of type Segment Routing Traffic Engineering (SR-TE) LSP. The user can associate an empty path or a path with strict or loose explicit hops with the primary path of the SR-TE LSP. A hop which corresponds to an adjacency SID must be identified with its far-end host IP address (next-hop) on the subnet. If the local end host IP address is provided, this hop is ignored since this router can have multiple adjacencies (next-hops) on the same subnet. A hop which corresponds to a node SID is identified by the prefix address. The user is only allowed to configure a primary path for the SR-TE LSP.

Platforms

7705 SAR Gen 2

lsp

Syntax

[**no**] **lsp** *lsp-name*

Context

[Tree] (config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter>sr-te lsp)

[Tree] (config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter>rsvp-te lsp)

Full Context

configure router static-route-entry indirect tunnel-next-hop resolution-filter sr-te lsp

configure router static-route-entry indirect tunnel-next-hop resolution-filter rsvp-te lsp

Description

This command restricts the search for a resolving LSP to a specific set of named LSPs. Only those LSPs named in the associated name list will be searched for a match to resolve the associated static route.

Parameters

lsp-name

Specifies the name of the LSP to be searched for a valid resolving tunnel for the static route's next-hop.

Platforms

7705 SAR Gen 2

lsp

Syntax

[**no**] **lsp** *lsp-name*

Context

[Tree] (config>service>sdp lsp)

Full Context

configure service sdp lsp

Description

This command creates associations between one or more label switched paths (LSPs) and an Multi-Protocol Label Switching (MPLS) service destination point (SDP). This command is implemented *only* on MPLS-type encapsulated SDPs.

In MPLS SDP configurations *either* one or more LSP names can be specified *or* LDP can be enabled. The SDP **ldp** and **lsp** commands are mutually exclusive except if the mixed-lsp-mode option is also enabled. If an LSP is specified on an MPLS SDP, then LDP cannot be enabled on the SDP. To enable LDP on the SDP when an LSP is already specified, the LSP must be removed from the configuration using the **no lsp lsp-name** command.

Alternatively, if LDP is already enabled on an MPLS SDP, then an LSP cannot be specified on the SDP. To specify an LSP on the SDP, the LDP must be disabled or the mixed-lsp-mode option is also enabled. The LSP must have already been created in the **config>router>mpls** context. with a valid far-end IP address. RSVP must be enabled.

If no LSP is associated with an MPLS SDP, the SDP cannot enter the operationally up state. The SDP can be administratively enabled (**no shutdown)** with no LSP associations. The *lsp-name* may be shutdown, causing the association with the SDP to be operationally down (the LSP will not be used by the SDP).

Up to 16 LSP names can be entered on a single command line.

The **no** form of this command deletes one or more LSP associations from an SDP. If the *lsp-name* does not exist as an association or as a configured LSP, no error is returned. An *lsp-name* must be removed from all SDP associations before the *lsp-name* can be deleted from the system. The SDP must be administratively disabled (**shutdown)** before the last *lsp-name* association with the SDP is deleted.

Parameters

lsp-name

Specifies the name of the LSP to associate with the SDP. An LSP name is case sensitive and is limited to 32 ASCII 7-bit printable characters with no spaces. If an exact match of *lsp-name* does not already exist as a defined LSP, an error message is generated. If the *lsp-name* does exist and the LSP **to** IP address matches the SDP **far-end** IP address, the association is created.

Platforms

7705 SAR Gen 2

lsp

Syntax

lsp *lsp-name*

[**no**] **lsp**

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls>rsvp-te lsp)

Full Context

configure oam-pm session ip tunnel mpls rsvp-te lsp

Description

This command configures the name of the RSVP-TE LSP to transport the test packets.

The **no** form of this command removes the *lsp-name* from the configuration.

Parameters

lsp-name

Specifies the LSP name, up to 64 characters.

Platforms

7705 SAR Gen 2

lsp

Syntax

lsp *lsp-name*

no lsp

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls>sr-te lsp)

Full Context

configure oam-pm session ip tunnel mpls sr-te lsp

Description

This command configures specification of SR-TE specific tunnel information that is used to transport the test packets. Entering this context removes all other tunnel type options configured under the **configure oam-pm session ip tunnel mpls** context. Only a single **mpls** type can be configured for an OAM-PM session.

The **no** form of this command removes the SR-TE LSP name from the configuration.

Default

no lsp

Parameters

lsp-name

Specifies the SR-TE LSP name, up to 64 characters.

Platforms

7705 SAR Gen 2

16.146 lsp-bsid-block

lsp-bsid-block

Syntax

lsp-bsid-block *name*

no lsp-bsid-block

Context

[Tree] (config>router>mpls lsp-bsid-block)

Full Context

configure router mpls lsp-bsid-block

Description

This command configures a reference to a pre-existing reserved label block for statically configured binding SIDs.

The **no** form of this command removes the use of the label block as a pool of binding SIDs.

Parameters

name

Specifies an existing reserved label block name, up to 64 characters.

Platforms

7705 SAR Gen 2

16.147 lsp-exp

lsp-exp

Syntax

lsp-exp *lsp-exp-value* **fc** *fc-name* **profile** {**in** | **out**}

no lsp-exp

Context

[Tree] (config>qos>network>ingress lsp-exp)

Full Context

configure qos network ingress lsp-exp

Description

This command creates a mapping between the LSP EXP bits of the network ingress traffic and the forwarding class.

Ingress traffic that matches the specified LSP EXP bits will be assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all eight LSP EXP bit values to the forwarding class. For undefined values, packets are assigned to the forwarding class specified under the **default-action** command.

The **no** form of this command removes the association of the LSP EXP bit value to the forwarding class. The **default-action** then applies to that LSP EXP bit pattern.

Default

no lsp-exp

Parameters

lsp-exp-value

Specify the LSP EXP values to be associated with the forwarding class.

Values 0 to 8 (Decimal representation of three EXP bit field)

fc *fc-name*

Enter this required parameter to specify the fc-name that the EXP bit pattern will be associated with.

Values be, l2, af, l1, h2, ef, h1, nc

profile {in | out}

Enter this required parameter to indicate whether the LSP EXP value is the in-profile or out-of-profile value.

Values in, out

Platforms

7705 SAR Gen 2

16.148 lsp-exp-in-profile

lsp-exp-in-profile

Syntax

lsp-exp-in-profile *lsp-exp-value*

no lsp-exp-in-profile

Context

[Tree] (config>qos>network>egress>fc lsp-exp-in-profile)

Full Context

configure qos network egress fc lsp-exp-in-profile

Description

This command specifies the in-profile LSP EXP value for the forwarding class. The EXP value will be used for all LSP labeled packets requiring marking that require marking at egress on this forwarding class queue, and that are in-profile. The inplus-profile traffic is marked with the same value as in-profile traffic.

When multiple EXP values are associated with the forwarding class at network egress, the last name entered will overwrite the previous value.

The **no** form of this command resets the configuration to the factory default in-profile EXP setting.

Default

Policy-id 1:	Factory setting
Policy-id 2 to 65535:	Policy-id setting

Parameters

lsp-exp-value

Specifies the 3-bit LSP EXP bit value, expressed as a decimal integer.

Values 0 to 7

Platforms

7705 SAR Gen 2

16.149 lsp-exp-out-profile

lsp-exp-out-profile

Syntax

lsp-exp-out-profile *lsp-exp-value*

no lsp-exp-out-profile

Context

[Tree] (config>qos>network>egress>fc lsp-exp-out-profile)

Full Context

configure qos network egress fc lsp-exp-out-profile

Description

This command specifies the out-of-profile LSP EXP value for the forwarding class. The EXP value will be used for all LSP labeled packets that require marking at egress on this forwarding class queue, and that are out-of-profile. The exceed-profile traffic is marked with the same value as out-of-profile traffic.

When multiple EXP values are associated with the forwarding class at network egress, the last name entered will overwrite the previous value.

The **no** form of this command resets the configuration to the factory default out-of-profile EXP setting.

Default

Policy-id 1:	Factory setting
Policy-id 2 to 65535:	Policy-id setting

Parameters

mpls-exp-value

Specifies the 3-bit MPLS EXP bit value, expressed as a decimal integer.

Values 0 to 7

Platforms

7705 SAR Gen 2

16.150 lsp-history

lsp-history

Syntax

[no] lsp-history

Context

[Tree] (config>router>mpls lsp-history)

Full Context

configure router mpls lsp-history

Description

This command allocates memory, which may be used to store up to the last 100 significant events for each RSVP-TE and SR-TE LSP.

The **no** form of this command deallocates memory for storing LSP history events and deletes any event history .

Default

no lsp-history

Platforms

7705 SAR Gen 2

16.151 lsp-init-retry-timeout

lsp-init-retry-timeout

Syntax

lsp-init-retry-timeout *seconds*

no lsp-init-retry-timeout

Context

[Tree] (config>router>mpls lsp-init-retry-timeout)

Full Context

configure router mpls lsp-init-retry-timeout

Description

This command configures the initial LSP path retry-timer.

The new LSP path initial retry-timer is used instead of the retry-timer to abort the retry cycle when no RESV is received. The retry-timer exclusively governs the time between two retry cycles and to handle retrying of an LSP path in a failure case with PATH errors or RESVTear.

The intent is that the user can now control how many refreshes of the pending PATH state can be performed before starting a new retry-cycle with a new LSP ID. This is all done without affecting the ability to react faster to failures of the LSP path, which will continue to be governed by the retry-timer.

The **no** form of this command returns the timer to the default value.

Default

lsp-init-retry-timeout 30

Parameters

seconds

Specifies the value (in s), used as the fast retry timer for a secondary path.

Values 10 to 600

Default 30

Platforms

7705 SAR Gen 2

16.152 lsp-lifetime

lsp-lifetime

Syntax

lsp-lifetime *seconds*

no lsp-lifetime

Context

[Tree] (config>service>vprn>isis lsp-lifetime)

Full Context

configure service vprn isis lsp-lifetime

Description

This command sets the time, in seconds, the router wants the LSPs it originates to be considered valid by other routers in the domain.

Each LSP received is maintained in an LSP database until the **lsp-lifetime** expires unless the originating router refreshes the LSP. By default, each router refreshes its LSPs every 20 minutes (1200 seconds) so other routers will not age out the LSP.

The LSP refresh timer is derived from this formula: lsp-lifetime/2

LSPs originated by the router should be valid for 1200 seconds (20 minutes).

The **no** form of this command reverts to the default value.

Default

lsp-lifetime 1200

Parameters

seconds

Specifies the time, in seconds, that the router wants the LSPs it originates to be considered valid by other routers in the domain.

Values 350 to 65535

Platforms

7705 SAR Gen 2

lsp-lifetime

Syntax

lsp-lifetime *seconds*

no lsp-lifetime

Context

[Tree] (config>router>isis lsp-lifetime)

Full Context

configure router isis lsp-lifetime

Description

This command sets the time, in seconds, the router wants the LSPs it originates to be considered valid by other routers in the domain.

Each LSP received is maintained in an LSP database until the **lsp-lifetime** expires unless the originating router refreshes the LSP. By default, each router refreshes its LSPs every 20 minutes (1200 seconds) so other routers will not age out the LSP.

The LSP refresh timer is derived from this formula: lsp-lifetime/2

The **no** form of this command reverts to the default value.

Default

lsp-lifetime 1200

Parameters

seconds

Specifies the time, in seconds, that the router wants the LSPs it originates to be considered valid by other routers in the domain.

Values 350 to 65535

Platforms

7705 SAR Gen 2

16.153 lsp-minimum-remaining-lifetime

lsp-minimum-remaining-lifetime

Syntax

lsp-minimum-remaining-lifetime *seconds*

no lsp-minimum-remaining-lifetime

Context

[Tree] (config>service>vprn>isis lsp-minimum-remaining-lifetime)

Full Context

configure service vprn isis lsp-minimum-remaining-lifetime

Description

This command configures the minimum value to which the remaining lifetime of the LSP is set. The value is a counter that decrements, in seconds, starting from the value in the received LSP (if not self-originated) or from **lsp-lifetime** *seconds* (if self-originated). When the remaining lifetime becomes zero, the contents of the LSP is purged. The remaining lifetime of an LSP is not changed when there is no **lsp-minimum-remaining-lifetime** value configured.

The configured value must be greater than or equal to the **lsp-lifetime** value.

The **no** form of this command removes the *seconds* value from the configuration.

Default

no lsp-minimum-remaining-lifetime

Parameters

seconds

Specifies the decrementing counter, in seconds. The configured value must be greater than or equal to the locally configured value of lsp-lifetime (MaxAge).

Values 350 to 65535

Platforms

7705 SAR Gen 2

lsp-minimum-remaining-lifetime

Syntax

lsp-minimum-remaining-lifetime *seconds*

no lsp-minimum-remaining-lifetime

Context

[Tree] (config>router>isis lsp-minimum-remaining-lifetime)

Full Context

configure router isis lsp-minimum-remaining-lifetime

Description

This command configures the minimum value to which the remaining lifetime of the LSP is set. The value is a counter that decrements, in seconds, starting from the value in the received LSP (if not self-originated) or from **lsp-lifetime** *seconds* (if self-originated). When the remaining lifetime becomes zero, the contents of the LSP is purged. The remaining lifetime of an LSP is not changed when there is no **lsp-minimum-remaining-lifetime** value configured.

The configured value must be greater than or equal to the **lsp-lifetime** value.

The **no** form of this command removes the *seconds* value from the configuration.

Parameters

seconds

Specifies the decrementing counter, in seconds. The configured value must be greater than or equal to the locally configured value of lsp-lifetime (MaxAge).

Values 350 to 65535

Platforms

7705 SAR Gen 2

16.154 lsp-mtu-size

lsp-mtu-size

Syntax

lsp-mtu-size *size*

no lsp-mtu-size

Context

[Tree] (config>service>vprn>isis lsp-mtu-size)

[Tree] (config>service>vprn>isis>level lsp-mtu-size)

Full Context

configure service vprn isis lsp-mtu-size

configure service vprn isis level lsp-mtu-size

Description

This command configures the LSP MTU size. If the *size* value is changed from the default using CLI or SNMP, then ISIS must be restarted for the change to take effect. This can be done by performing a **shutdown** command and then a **no shutdown** command in the **config>router>isis** context.

Note:
Using the **exec** command to execute a configuration file to change the LSP MTU size from its default value will automatically restart IS-IS for the change to take effect.

The **no** form of this command reverts to the default value.

Default

lsp-mtu-size 1492

Parameters

size

Specifies the LSP MTU size.

Values 490 to 9778

Platforms

7705 SAR Gen 2

lsp-mtu-size

Syntax

lsp-mtu-size *size*

no lsp-mtu-size

Context

[Tree] (config>router>isis lsp-mtu-size)

[Tree] (config>router>isis>level lsp-mtu-size)

Full Context

configure router isis lsp-mtu-size

configure router isis level lsp-mtu-size

Description

This command configures the LSP MTU size. If the *size* value is changed from the default using CLI or SNMP, then IS-IS must be restarted in order for the change to take effect. This can be done by performing a **shutdown** command and then a **no shutdown** command in the **config>router>isis** context.

Note:
Using the **exec** command to execute a configuration file to change the LSP MTU-size from its default value automatically restarts IS-IS for the change to take effect.

The **no** form of this command reverts to the default value.

Default

lsp-mtu-size 1492

Parameters

size

Specifies the LSP MTU size.

Values 490 to 9778

Platforms

7705 SAR Gen 2

16.155 lsp-pacing-interval

lsp-pacing-interval

Syntax

lsp-pacing-interval *milliseconds*

no lsp-pacing-interval

Context

[Tree] (config>service>vprn>isis>if lsp-pacing-interval)

Full Context

configure service vprn isis interface lsp-pacing-interval

Description

This command configures the interval at which LSPs are sent from the interface.

To avoid overwhelming neighbors that have less CPU processing power with LSPs, the pacing interval can be configured to limit how many LSPs are sent at the interval. LSPs are sent in bursts at the interval up to the configured limit. If a value of 0 is configured, no LSPs are sent from the interface.

If configured to the default LSP pacing interval of 100, LSPs are sent in 100 millisecond intervals.

The **no** form of this command reverts to the default value.

Note:
The IS-IS LSP pacing interval is 100 milliseconds for values < 100 milliseconds, and 1 second for values ≥ 100 milliseconds. For example, a pacing interval of 2 milliseconds means that a maximum of 50 LSPs are sent in a burst at 100 millisecond intervals. The default pacing interval of 100 milliseconds means that a maximum of 10 LSPs are sent in a burst at 1 second intervals.

Default

lsp-pacing-interval 100

Parameters

milliseconds

Specifies the pacing interval in milliseconds at which IS-IS LSPs are sent from the interface at each interval expressed as a decimal integer.

Values 0 to 65535

Platforms

7705 SAR Gen 2

lsp-pacing-interval

Syntax

lsp-pacing-interval *milliseconds*

no lsp-pacing-interval

Context

[Tree] (config>router>isis>interface lsp-pacing-interval)

Full Context

configure router isis interface lsp-pacing-interval

Description

This command configures the interval at which LSPs are sent from the interface.

To avoid overwhelming neighbors that have less CPU processing power with LSPs, the pacing interval can be configured to limit how many LSPs are sent at the interval. LSPs are sent in bursts at the interval up to the configured limit. If a value of 0 is configured, no LSPs are sent from the interface. The interval applies to all LSPs: LSPs generated by the router, and LSPs received from other routers.

If configured to the default LSP pacing interval of 100, LSPs are sent in 100 millisecond intervals.

The **no** form of this command reverts to the default value.

Note:
The IS-IS LSP pacing interval is 100 milliseconds for values < 100 milliseconds, and 1 second for values ≥ 100 milliseconds. For example, a pacing interval of 2 milliseconds means that a maximum of 50 LSPs are sent in a burst at 100 millisecond intervals. The default pacing interval of 100 milliseconds means that a maximum of 10 LSPs are sent in a burst at 1 second intervals.

Default

lsp-pacing-interval 100

Parameters

milli-seconds

Specifies the interval in milliseconds during which IS-IS LSPs are sent from the interface expressed as a decimal integer.

Values 0 to 65535

Platforms

7705 SAR Gen 2

16.156 lsp-ping

lsp-ping

Syntax

lsp-ping *lsp-name* [**path** *path-name*]

lsp-ping bgp-label prefix *ip-prefix/prefix-length* [**path-destination** *ip-address* [{**interface** *if-name* | **next-hop** *ip-address*}]]

lsp-ping ldp prefix *ip-prefix/prefix-length* [**path-destination** *ip-address* [{**interface** *if-name* | **next-hop** *ip-address*}]]

lsp-ping prefix *ip-prefix/prefix-length* [**path-destination** *ip-address* [{**interface** *if-name* | **next-hop** *ip-address*}]]

lsp-ping rsvp-te *lsp-name* [**path** *path-name*]

lsp-trace sr-isis prefix *ip-prefix/prefix-length* [**igp-instance** *igp-instance*] [**detail**] [**downstream-map-tlv** *downstream-map-tlv*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**flex-algo** *flex-algo-num*] [**interval** *interval*] [**max-fail** *no-response-count*] [**min-ttl** *min-label-ttl*] [**max-ttl** *max-label-ttl*] [**path-destination** *ip-address* [**interface** *if-name* | **next-hop** *ip-address*]] [**probe-count** *probes-per-hop*] [**size** *octets*] [**src-ip-address** *ip-address*] [**timeout** *timeout*]

lsp-trace sr-ospf prefix *ip-prefix/prefix-length* [**igp-instance** *igp-instance*] [**detail**] [**downstream-map-tlv** *downstream-map-tlv*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**flex-algo** *flex-algo-num*] [**interval** *interval*] [**max-fail** *no-response-count*] [**min-ttl** *min-label-ttl*] [**max-ttl** *max-label-ttl*] [**path-destination** *ip-address* [**interface** *if-name* | **next-hop** *ip-address*]] [**probe-count** *probes-per-hop*] [**size** *octets*] [**src-ip-address** *ip-address*] [**timeout** *timeout*]

lsp-ping sr-ospf3 prefix *ip-prefix/prefix-length* [**igp-instance** *igp-instance*] [**path-destination** *ip-address* [{**interface** *if-name* | **next-hop** *ip-address*}]]

lsp-ping sr-policy color *color-id* **endpoint** *ip-address* [**segment-list** *segment-list-id*] [**detail**] [**path-destination** *ip-address* [{ **interface** *if-name* | **next-hop** *ip-address*}]]

lsp-ping sr-te *lsp-name* [**path** *path-name*] [**path-destination** *ip-address* [{**interface** *if-name* | **next-hop** *ip-address*}]]

lsp-ping static *lsp-name* [**assoc-channel** {**ipv4** | **non-ip** | **none**}] [**dest-global-id** *global-id* **dest-node-id** *node-id*] [**force**] [**path-type** {**active** | **working** | **protect**}]

NOTE: Options common to all **lsp-ping** cases: [**detail**] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**flex-algo** *flex-algo-num*] [**interval** *interval*] [**send-count** *send-count*] [**size** *octets*] [**src-ip-address** *ip-address*] [**timeout** *timeout*] [**ttl** *label-ttl*]

Context

[Tree] (oam lsp-ping)

[Tree] (config>saa>test>type lsp-ping)

Full Context

oam lsp-ping

configure saa test type lsp-ping

Description

This command performs in-band LSP connectivity tests.

This command performs an LSP ping using the protocol and data structures defined in the RFC 8029, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*.

The LSP ping operation is modeled after the IP ping utility which uses ICMP echo request and reply packets to determine IP connectivity.

In an LSP ping, the originating device creates an MPLS echo request packet for the LSP and path to be tested. The MPLS echo request packet is sent through the data plane and awaits an MPLS echo reply packet from the device terminating the LSP. The status of the LSP is displayed when the MPLS echo reply packet is received.

This command, when used with the **static** option, performs in-band on-demand LSP connectivity verification tests for static MPLS-TP LSPs. For other LSP types, the **static** option should be excluded and these are described elsewhere in this user guide.

The **lsp-ping static** command performs an LSP ping using the protocol and data structures defined in the RFC 8029, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*, as extended by RFC 6426, *MPLS On-Demand Connectivity Verification and Route Tracing*.

In MPLS-TP, the echo request and echo reply messages are always sent in-band over the LSP, either in a G-ACh channel or encapsulated as an IP packet below the LSP label.

The timestamp format to be sent, and to be expected when received in a PDU, is as configured by the **config>test-oam>mpls-time-stamp-format** command. If RFC 4379 (obsoleted by RFC 8029) is selected, then the timestamp is in seconds and microseconds since 1900, otherwise it is in seconds and microseconds since 1970.

Default

The active LSP path

Values: Any path name associated with the LSP

Parameters

lsp-name

Specifies the name of the target RSVP-TE LSP, up to 64 characters.

rsvp-te lsp-name

Specifies the name of the target RSVP-TE LSP, up to 64 characters.

Note:
The **rsvp-te** explicit target FEC type is not supported under the SAA context.

path-name

Specifies the LSP path name, up to 32 characters, to which to send the LSP ping request.

Values Any path name associated with the LSP.

Default The active LSP path.

bgp-label prefix *ip-prefix/prefix-length*

Specifies the address prefix and subnet mask of the target BGP IPv4 /32 label route or the target BGP IPv6 /128 label route.

Values <ipv4-prefix>/32 | <ipv6-prefix>/128

ipv4-prefix a.b.c.d

ipv6-prefix x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

path-destination *ip-address*

Specifies the IP address of the path destination from the range 127/8. When the LDP FEC prefix is IPv6, the user must enter a 127/8 IPv4 mapped IPv6 address, that is, in the range ::ffff:127/104.

ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

flex-algo *flex-algo-num*

Specifies the Segment Routing Flexible Algorithm for the test.

Values 128 to 255

interface if-name

Specifies the name of an IP interface, up to 32 characters, to send the MPLS echo request message to. The name must already exist in the **config>router>interface** context.

next-hop *ip-address*

Specifies the next-hop address to send the MPLS echo request message to.

Values

ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

prefix *ip-prefix/prefix-length*

Specifies the address prefix and subnet mask of the target LDP FEC.

Values <ipv4-prefix>/32 | <ipv6-prefix>/128

ipv4-prefix a.b.c.d

ipv6-prefix x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

ldp prefix *ip-prefix/prefix-length*

Specifies the address prefix and subnet mask of the target LDP FEC.

Values <ipv4-prefix>/32 | <ipv6-prefix>/128

ipv4-prefix a.b.c.d

ipv6-prefix x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

sr-isis prefix *ip-prefix/prefix-length*

Specifies the address prefix and subnet mask of the target node SID of the SR-IS-IS tunnel.

Values <ipv4-prefix>/32 | <ipv6-prefix>/128

ipv4-prefix a.b.c.d

ipv6-prefix x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

igp-instance

Specifies the IGP instance of the node SID prefix.

Values isis-inst: 0 to 127

ospf3-inst: 0 to 31, 64 to 95

ospf-inst: 0 to 31

sr-ospf prefix *ip-prefix/prefix-length*

Specifies the address prefix and subnet mask of the target node SID of the SR-OSPF tunnel.

Values <ipv4-prefix>/32 | <ipv6-prefix>/128

ipv4-prefix - a.b.c.d

ipv6-prefix - x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

sr-ospf3 prefix *ip-prefix/prefix-length*

Specifies the address prefix and subnet mask of the target node SID of the SR-OSPF3 tunnel. Note that only IPv6 prefixes in OSPFv3 instance ID 0-31 are supported.

Values

ipv6-prefix - x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

sr-policy color *color-id* endpoint *ip-address* segment-list *segment-list-id*

Specifies the name of the target IPv4 or IPv6 SR policy.

Note:
The **sr-policy** target FEC type is supported under the OAM context and under **type-multi-line node** in the SAA context.

color *color-id* — Specifies the color ID.

Values 0 to 4294967295

endpoint *ip-address* — Specifies the endpoint address.

Values

ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

segment-list *segment-list-id* — Specifies the segment list ID.

Values 1 to 32

detail

Displays detailed information.

sr-te *lsp-name*

Specifies the name of the target SR-TE LSP, up to 64 characters.

static

Specifies the target FEC stack sub-type "Static LSP”.

assoc-channel {ipv4 | non-ip | none}

Specifies the launched echo request’s usage of the Associated Channel (ACH) mechanism, when testing an MPLS-TP LSP.

Values **ipv4** — Use an Associated Channel with IP encapsulation, as described in RFC 6426, Section 3.2.

non-ip — Do not use an Associated Channel, as described in RFC 6426, Section 3.1.

none — Use the Associated Channel mechanism described in RFC 6426, Section 3.3.

Default non-ip

global-id

Specifies the MPLS-TP global ID for the far end node of the LSP under test. If this is not entered, then the dest-global-id is taken from the LSP context.

Values 0 to 4294967295

Default 0

node-id

Specifies the MPLS-TP global ID for the far end node of the LSP under test. If this is not entered, then the dest-global-id is taken from the LSP context.

Values a.b.c.d, 1 to 4294967295

Default 0

force

Allows LSP ping to test a path that is operationally down, including cases where MPLS-TP BFD CC/V is enabled and has taken a path down. This parameter is only allowed in the OAM context; it is not allowed for a test configured as a part of an SAA.

Default disabled

path-type {active | working | protect}

The LSP path to test.

Values **active** — The currently active path. If MPLS-TP linear protection is configured on the LSP, then this is the path that is selected by the

MPLS-TP PSC protocol for sending user plane traffic. If MPLS-TP linear protection is not configured, then this is the working path.

working — The working path of the MPLS-TP LSP.

protect — The protect path of the MPLS-TP LSP.

Default active

fc-name

Specifies the FC and profile parameters that are used to indicate the forwarding class and profile of the MPLS echo request packet.

When an MPLS echo request packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the specified fc and profile parameter values. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface.

When the MPLS echo request packet is received on the responding node, The FC and profile parameter values are dictated by the LSP-EXP mappings of the incoming interface.

When an MPLS echo reply packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the FC and profile parameter values determined by the classification of the echo request packet, which is being replied to, at the incoming interface. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface. The ToS byte is not modified. Table 50: lsp-ping Request Packet and Behavior summarizes this behavior.

Table 50: lsp-ping Request Packet and Behavior

CPM (sender node)	Echo request packet: • packet {tos=1, fc1, profile1} • fc1 and profile1 are as entered by user in OAM command or default values • tos1 as per mapping of {fc1, profile1} to IP precedence in network egress QoS policy of outgoing interface
Outgoing interface (sender node)	Echo request packet: • packet queued as {fc1, profile1} • ToS field=tos1 not remarked • EXP=exp1, as per mapping of {fc1, profile1} to EXP in network egress QoS policy of outgoing interface
Incoming interface (responder node)	Echo request packet: • packet {tos1, exp1} • exp1 mapped to {fc2, profile2} as per classification in network QoS policy of incoming interface
CPM (responder node)	Echo reply packet:

	• packet{tos=1, fc2, profile2}
Outgoing interface (responder node)	Echo reply packet: • packet queued as {fc2, profile2} • ToS filed= tos1 not remarked (reply inband or out-of-band) • EXP=exp2, if reply is inband, remarked as per mapping of {fc2, profile2} to EXP in network egress QoS policy of outgoing interface
Incoming interface (sender node)	Echo reply packet: • packet {tos1, exp2} • exp2 mapped to {fc1, profile1} as per classification in network QoS policy of incoming interface

The LSP-EXP mappings on the receive network interface controls the mapping of the message reply at the originating router**.**

Values be, l2, af, l1, h2, ef, h1, nc

Default be

profile {in | out}

Specifies the profile state of the MPLS echo request packet.

Default out

flex-algo *flex-algo-num*

Specifies the Segment Routing Flexible Algorithm for the test. This option is only supported for **oam lsp-ping sr-isis** and **oam lsp-ping sr-ospf**. This option is not supported for SAA. If this option is not set, then the system looks up the prefix without flex-algo awareness.

Values 128 to 255

Default **none**

interval

Specifies the time, in seconds, used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

Values 1 to 10

Default 1

send-count

Specifies the number of messages to send, expressed as a decimal integer. The **send-count** parameter is used to override the default number of message requests sent. Each message request must either time out or receive a reply before the next message request

is sent. The message **interval** value must be expired before the next message request is sent.

Values 1 to 100

Default 1

octets

Specifies the MPLS echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeros to the specified size.

Values 1 to 9786

Default 1

src-ip-address *ip-address*

Specifies the source IP address. This option is used when an OAM packet must be generated from a different address than the node's system interface address. An example is when the OAM packet is sent over an LDP LSP and the LDP LSR-ID of the corresponding LDP session to the next-hop is set to an address other than the system interface address.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

timeout

Specifies number, in seconds, used to override the default *timeout* value and is the amount of time that the router waits for a message reply after sending the last probe for a specific test. Upon the expiration of the time out, the test is marked complete and no more packets are processed for any of those request probes.

Values 1 to 10

Default 5

label-ttl

Specifies the TTL value for the MPLS label, expressed as a decimal integer.

Values 1 to 255

Default 255

Platforms

7705 SAR Gen 2

Output

Output Example

The following output is an example of LDP IPv4 and IPv6 prefix FECs.

```
A:Dut-C# oam lsp-ping prefix 4.4.4.4/32 detail
LSP-PING 4.4.4.4/32: 80 bytes MPLS payload
Seq=1, send from intf dut1_to_dut3, reply from 4.4.4.4
       udp-data-len=32 ttl=255 rtt=5.23ms rc=3 (EgressRtr)

---- LSP 4.4.4.4/32 PING Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 5.23ms, avg = 5.23ms, max = 5.23ms, stddev = 0.000ms

===============================================================================
LDP LSR ID: 1.1.1.1
===============================================================================
Legend: U - Label In Use,  N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
===============================================================================


LDP Prefix Bindings
===============================================================================
Prefix              IngLbl       EgrLbl     EgrIntf/         EgrNextHop
   Peer                                     LspId
-------------------------------------------------------------------------------
4.4.4.4/32          131069N      131067     1/1/1            1.3.1.2
   3.3.3.3
4.4.4.4/32          131069U      131064       --               --
   6.6.6.6
-------------------------------------------------------------------------------
No. of Prefix Bindings: 2
===============================================================================
A:Dut-C#

*A:Dut-A# oam lsp-ping prefix fc00::a14:106/128

LSP-PING fc00::a14:106/128: 116 bytes MPLS payload

Seq=1, send from intf A_to_B, reply from fc00::a14:106

udp-data-len=32 ttl=255 rtt=7.16ms rc=3 (EgressRtr)



---- LSP fc00::a14:106/128 PING Statistics ----

1 packets sent, 1 packets received, 0.00% packet loss

round-trip min = 7.16ms, avg = 7.16ms, max = 7.16ms, stddev = 0.000ms

*A:Dut-A#
```

lsp-ping over SR-ISIS

```
*A:Dut-A# oam lsp-ping sr-isis prefix 10.20.1.6/32 igp-instance 0 detail
LSP-PING 10.20.1.6/32: 80 bytes MPLS payload
Seq=1, send from intf int_to_B, reply from 10.20.1.6
       udp-data-len=32 ttl=255 rtt=1220324ms rc=3 (EgressRtr)
---- LSP 10.20.1.6/32 PING Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
```

```
round-trip min = 1220324ms, avg = 1220324ms, max = 1220324ms, stddev = 0.000ms
```

lsp-ping with SR-TE

```
*A:Dut-A# oam lsp-ping sr-te "srteABCEDF" detail
LSP-PING srteABCEDF: 96 bytes MPLS payload
Seq=1, send from intf int_to_B, reply from 10.20.1.6
       udp-data-len=32 ttl=255 rtt=1220325ms rc=3 (EgressRtr)
---- LSP srteABCEDF PING Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 1220325ms, avg = 1220325ms, max = 1220325ms, stddev = 0.000ms
```

```
*A:Dut-A# oam lsp-ping sr-te "srteABCE_loose" detail
LSP-PING srteABCE_loose: 80 bytes MPLS payload
Seq=1, send from intf int_to_B, reply from 10.20.1.5
       udp-data-len=32 ttl=255 rtt=1220324ms rc=3 (EgressRtr)
---- LSP srteABCE_loose PING Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 1220324ms, avg = 1220324ms, max = 1220324ms, stddev = 0.000ms
```

```
*A:Dut-F# oam lsp-ping sr-te "srteFECBA_eth" detail
LSP-PING srteFECBA_eth: 116 bytes MPLS payload
Seq=1, send from intf int_to_E, reply from fc00::a14:101
       udp-data-len=32 ttl=255 rtt=1220326ms rc=3 (EgressRtr)
---- LSP srteFECBA_eth PING Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 1220326ms, avg = 1220326ms, max = 1220326ms, stddev = 0.000ms
```

lsp-ping with SR-Policy

```
*A:Dut-A#
# ipv4 sr-policy lsp-ping
*A:Dut-A# oam lsp-ping sr-policy color 200 endpoint 10.20.1.6 LSP-PING color 200 endpoint
 10.20.1.6: 76 bytes MPLS payload Seq=1, send from intf int_to_C, reply from 10.20.1.6
       udp-data-len=32 ttl=255 rtt=1220325ms rc=3 (EgressRtr)
---- LSP color 200 endpoint 10.20.1.6 PING Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss round-trip min = 1220325ms, avg =
 1220325ms, max = 1220325ms, stddev = 0.000ms

# ipv6 sr-policy lsp-ping
*A:Dut-A# oam lsp-ping sr-policy color 200 endpoint fc00::a14:106 LSP-PING color 200 endpoint
 fc00::a14:106: 76 bytes MPLS payload Seq=1, send from intf int_to_C, reply from 10.20.1.6
       udp-data-len=32 ttl=255 rtt=1220324ms rc=3 (EgressRtr)
---- LSP color 200 endpoint fc00::a14:106 PING Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss round-trip min = 1220324ms, avg =
 1220324ms, max = 1220324ms, stddev = 0.000ms
```

lsp-ping with sr-ospf3

```
# sr-ospf3 lsp-ping
*A:Dut-A# oam lsp-ping sr-ospf3 prefix fc00::a14:106/128 LSP-PING fc00::a14:106/128: 116 bytes
 MPLS payload Seq=1, send from intf int_to_B, reply from fc00::a14:106
       udp-data-len=32 ttl=255 rtt=3.17ms rc=3 (EgressRtr)
---- LSP fc00::a14:106/128 PING Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss round-trip min = 3.17ms, avg = 3.17ms,
 max = 3.17ms, stddev = 0.000ms *A:Dut-A#
```

lsp-ping

Syntax

lsp-ping

Context

[Tree] (config>saa>test>type-multi-line lsp-ping)

Full Context

configure saa test type-multi-line lsp-ping

Description

Commands in this context configure the lsp-ping OAM probe type.

Platforms

7705 SAR Gen 2

16.157 lsp-ping-trace

lsp-ping-trace

Syntax

lsp-ping-trace [{**tx** | **rx** | **both**}] [{**raw** | **detail**}]

no lsp-ping-trace

Context

[Tree] (debug>oam lsp-ping-trace)

Full Context

debug oam lsp-ping-trace

Description

This command enables debugging for lsp-ping.

Parameters

tx | rx | both

Specifies to enable LSP ping debugging for TX, RX, or both RX and TX for the for debug direction.

raw | detail

Displays output for the for debug mode.

Platforms

7705 SAR Gen 2

16.158 lsp-refresh-interval

lsp-refresh-interval

Syntax

lsp-refresh-interval [*seconds*] [**half-lifetime** {**enable** | **disable**}]

no lsp-refresh-interval

Context

[Tree] (config>service>vprn>isis lsp-refresh-interval)

Full Context

configure service vprn isis lsp-refresh-interval

Description

This command configures the IS-IS LSP refresh timer interval for the VPRN instance. When configuring the LSP refresh interval, the value that is specified for **lsp-lifetime** must also be considered. The LSP refresh interval cannot be greater than 90% of the LSP lifetime.

The **no** form of this command reverts to the default (600 seconds), unless this value is greater than 90% of the LSP lifetime. For example, if the LSP lifetime is 400, then the no lsp-refresh-interval command will be rejected.

Default

lsp-refresh-interval 600 half-lifetime enable

Parameters

seconds

Specifies the refresh interval.

Values 150 to 65535

half-lifetime

Sets the refresh interval to always be half the **lsp-lifetime** value. When this parameter is set to **enable**, the configured refresh interval is ignored.

Values enable, disable

Platforms

7705 SAR Gen 2

lsp-refresh-interval

Syntax

lsp-refresh-interval [*seconds*] [**half-lifetime** {**enable** | **disable**}

no lsp-refresh-interval

Context

[Tree] (config>router>isis lsp-refresh-interval)

Full Context

configure router isis lsp-refresh-interval

Description

This command configures the IS-IS LSP refresh timer interval. When configuring the LSP refresh interval, the value that is specified for **lsp-lifetime** must also be considered. The LSP refresh interval cannot be greater than 90% of the LSP lifetime.

The **no** form of this command reverts to the default (600 seconds), unless this value is greater than 90% of the LSP lifetime. For example, if the LSP lifetime is 400, then the no lsp-refresh-interval command will be rejected.

Default

lsp-refresh-interval 600 half-lifetime enable

Parameters

seconds

Specifies the refresh interval.

Values 150 to 65535

half-lifetime

Sets the refresh interval to always be half the **lsp-lifetime** value. When this parameter is set to **enable**, the configured refresh interval is ignored.

Values enable, disable

Platforms

7705 SAR Gen 2

16.159 lsp-self-ping

lsp-self-ping

Syntax

lsp-self-ping {**enable** | **disable** | **inherit**}

no lsp-self-ping

Context

[Tree] (config>router>mpls>lsp-template lsp-self-ping)

Full Context

configure router mpls lsp-template lsp-self-ping

Description

This command enables LSP Self-ping on a given RSVP-TE LSP or LSP template. If set to **disable**, then LSP Self-ping is disabled irrespective of the setting of **lsp-self-ping**>**rsvp-te** under the **mpls** context. By default, each LSP and LSP template inherits this value.

The **no** form of this command reverts to the default.

Default

lsp-self-ping inherit

Parameters

enable

Enables LSP Self-ping on this RSVP LSP or RSVP LSPs (one-hop-p2p or mesh-p2p) using this LSP template.

disable

Disables LSP Self-ping on this RSVP LSP or RSVP LSPs using this LSP template.

inherit

Inherits the value configured under **config**>**router**>**mpls**>**lsp-self-ping**>**rsvp-te**.

Platforms

7705 SAR Gen 2

16.160 lsp-setup

lsp-setup

Syntax

lsp-setup [**detail**]

no lsp-setup

Context

[Tree] (debug>router>mpls>event lsp-setup)

Full Context

debug router mpls event lsp-setup

Description

This command debugs LSP setup events.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about LSP setup events.

Platforms

7705 SAR Gen 2

16.161 lsp-template

lsp-template

Syntax

lsp-template *template-name* [**mesh-p2p** | **mesh-p2p-srte** | **one-hop-p2p** | **on-demand-p2p-srte** | **one-hop-p2p-srte** | **p2mp** | **pce-init-p2p-srte template-id** {**default** | *template-id*}]

no lsp-template *template-name*

Context

[Tree] (config>router>mpls lsp-template)

Full Context

configure router mpls lsp-template

Description

This command creates a template that can be referenced by a client application where dynamic LSP creation is required. The LSP template type (**p2mp**, **one-hop-p2p**, **mesh-p2p**, **one-hop-p2p-srte**, **mesh-p2p-srte**, **pce-init-p2p-srte**, or **on-demand-p2p-srte**) is mandatory.

The **no** form of this command deletes the LSP template. An LSP template cannot be deleted if a client application is using it.

Parameters

template-name

Specifies the name of the LSP template, up to 32 characters. An LSP template name and LSP name must not be the same.

mesh-p2p | mesh-p2p-srte | one-hop-p2p | one-hop-p2p-srte | p2mp | pce-init-p2p-srte | on-demand-p2p-srte

Identifies the type of LSP this template will signal.

default

Sets the template to be the default LSP template for PCE-initiated SR-TE LSPs.

template-id

Specifies the value that is signaled in the PCE to identify the LSP template.

Platforms

7705 SAR Gen 2

lsp-template

Syntax

lsp-template *template-name*

no lsp-template

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls>rsvp-te-auto lsp-template)

Full Context

configure oam-pm session ip tunnel mpls rsvp-te-auto lsp-template

Description

This command configures the name of the LSP template used to identify the unique LSP. Configure the following three commands to identify an RSVP-TE Auto LSP: **from**, **to**, and **lsp-template**. When all three of these values are configured, the specific RSVP LSP can be identified and the test packets can be carried across the tunnel

The **no** form of this command removes the LSP template name from the configuration.

Parameters

template-name

Specifies the LSP template name, up to 32 characters.

Platforms

7705 SAR Gen 2

16.162 lsp-trace

lsp-trace

Syntax

lsp-trace *lsp-name* [**path** *path-name*] [**detail**]

lsp-trace bgp-label prefix *ip-prefix/prefix-length* [**path-destination** *ip-address* [{**interface** *if-name* | **next-hop** *ip-address*}]]

lsp-trace ldp prefix *ip-prefix/length* [**path-destination** *ip-address* [{**interface** *if-name* | **next-hop** *ip-address*}]]

lsp-trace prefix *ip-prefix/length* [**path-destination** *ip-address* [{**interface** *if-name* | **next-hop** *ip-address*}]]

lsp-trace rsvp-te *lsp-name* [**path** *path-name*]

lsp-trace sr-isis prefix *ip-prefixIprefix-length* [**igp-instance** *igp-instance*] [**path-destination** *ip-address* [{**interface** *if-name* | **next-hop** *ip-address*}]]

lsp-trace sr-ospf prefix *ip-prefixIprefix-length* [**igp-instance** *igp-instance*] [**path-destination** *ip-address* [{**interface** *if-name* | **next-hop** *ip-address*}]]

lsp-trace sr-ospf3 prefix *ip-prefixIprefix-length* [**igp-instance** *igp-instance*] [**path-destination** *ip-address* [{**interface** *if-name*| **next-hop** *ip-address*}]]

lsp-trace sr-policy color *color-id* **endpoint** *ip-address* [**segment-list** *segment-list-id*] [**path-destination** *ip-address* [{**interface** *if-name* | **next-hop** *ip-address*}]]

lsp-trace sr-te *lsp-name* [**path** *path-name*] [**path-destination** *ip-address* [{**interface** *if-name* | **next-hop** *ip-address*}]]

lsp-trace static *lsp-name* [**assoc-channel** {**ipv4** | **non-ip** | **none**}] [**path-type** {**active** | **working** | **protect**}]

NOTE: Options common to all **lsp-trace** cases: [**detail**] [**downstream-map-tlv** *downstream-map-tlv*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**flex-algo** *flex-algo-num*] [**interval** *interval*] [**max-fail** *no-response-count*] [**max-ttl** *max-label-ttl*] [**min-ttl** *min-label-ttl*] [**probe-count** *probes-per-hop*] [**size** *octets*] [**src-ip-address** *ip-address*] [**timeout** *timeout*]

Context

[Tree] (oam lsp-trace)

Full Context

oam lsp-trace

Description

This command performs an LSP traceroute using the protocol and data structures defined in IETF RFC 8029.

The LSP trace operation is modeled after the IP traceroute utility, which uses ICMP echo request and reply packets with increasing TTL values to determine the hop-by-hop route to a destination IP.

In an LSP trace, the originating device creates an MPLS echo request packet for the LSP to be tested with increasing values of the TTL in the outermost label. The MPLS echo request packet is sent through the data plane and awaits a TTL exceeded response or the MPLS echo reply packet from the device terminating the LSP. The devices that reply to the MPLS echo request packets with the TTL exceeded and the MPLS echo reply are displayed.

The downstream mapping TLV is used in **lsp-trace** to provide a mechanism for the sender and responder nodes to exchange and validate interface and label stack information for each downstream hop in the path of the LDP FEC an RSVP LSP, or a BGP IPv4 label route.

Two downstream mapping TLVs are supported. The original Downstream Mapping (DSMAP) TLV defined in RFC 4379 (obsoleted by RFC 8029) and the new Downstream Detailed Mapping (DDMAP) TLV defined in RFC 6424 AND RFC 8029. More details are provided in the DDMAP TLV sub-section below.

In addition, when the responder node has multiple equal cost next hops for an LDP FEC, a BGP label IPv4 prefix, an SR-ISIS node SID, an SR-OSPF node SID, or an SR-TE LSP, it replies in the Downstream Mapping TLV with the downstream information for each outgoing interface which is part of the ECMP next-hop set for the prefix. The downstream mapping TLV can further be used to exercise a specific path of the ECMP set using the **path-destination** option.

This command, when used with the **static** option, performs in-band on-demand LSP traceroute tests for static MPLS-TP LSPs. For other LSP types, the **static** option should be excluded and these are described elsewhere in this user guide.

The **lsp-trace static** command performs an LSP trace using the protocol and data structures defined in the RFC 8029, Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures, as extended by RFC 6426, MPLS On-Demand Connectivity Verification and Route Tracing.

In MPLS-TP, the echo request and echo reply messages are always sent in-band over the LSP, either in a G-ACh channel or encapsulated as an IP packet below the LSP label.

The timestamp format to be sent, and to be expected when received in a PDU, is as configured by the **configure test-oam mpls-time-stamp-format** command. If RFC 4379 (obsoleted by RFC 8029) is selected, the timestamp is in seconds and microseconds since 1900, otherwise it is in seconds and microseconds since 1970.

Parameters

lsp-name

Specifies the name of the target RSVP-TE LSP, up to 64 characters.

rsvp-te lsp-name

Specifies the name of the target RSVP-TE LSP, up to 64 characters.

Note:
The **rsvp-te** explicit target FEC type is not supported under the SAA context.

path-name

Specifies the LSP path name along which to send the LSP trace request.

Values Any path name associated with the LSP.

Default The active LSP path.

bgp-label prefix *ip-prefix/prefix-length*

Specifies the address prefix and subnet mask of the target BGP IPv4 /32 label route or the target IPv6 /128 label route.

Values <ipv4-prefix>/32 | <ipv6-prefix>/128

ipv4-prefix a.b.c.d

ipv6-prefix x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

path-destination *ip-address*

Specifies the IP address of the path destination from the range 127/8. When the LDP FEC prefix is IPv6, the user must enter a 127/8 IPv4 mapped IPv6 address, that is, in the range ::ffff:127/104.

if-name

Specifies the name of an IP interface, up 32 characters, to send the MPLS echo request to. The name must already exist in the **config>router>interface** context.

next-hop *ip-address*

Specifies the next hop to send the MPLS echo request message to.

Values ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

prefix *ip-prefix/prefix-length*

Specifies the address prefix and subnet mask of the target LDP FEC.

Values <ipv4-prefix>/32 | <ipv6-prefix>/128

ipv4-prefix - a.b.c.d

ipv6-prefix - x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

ldp prefix *ip-prefix/prefix-length*

Specifies the address prefix and subnet mask of the target LDP FEC.

Values <ipv4-prefix>/32 | <ipv6-prefix>/128

ipv4-prefix a.b.c.d

ipv6-prefix x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

sr-isis prefix *ip-prefix/prefix-length*

Specifies the address prefix and subnet mask of the target node SID of the SR-ISIS tunnel.

Values <ipv4-prefix>/32 | <ipv6-prefix>/128

ipv4-prefix a.b.c.d

ipv6-prefix x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

igp-instance

Specifies the IGP instance of the node SID prefix.

Values isis-inst: 0 to 127

ospf3-inst: 0 to 31, 64 to 95

ospf-inst: 0 to 31

sr-ospf prefix *ip-prefix/prefix-length*

Specifies the address prefix and subnet mask of the target node SID of the SR-OSPF tunnel.

Values <ipv4-prefix>/32 | <ipv6-prefix>/128

ipv4-prefix - a.b.c.d

ipv6-prefix - x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

sr-ospf3 prefix *ip-prefix/prefix-length*

Specifies the address prefix and subnet mask of the target node SID of the SR-OSPF3 tunnel. Only IPv6 prefixes in OSPFv3 instance ID 0-31 are supported.

Values

ipv6-prefix - x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

sr-policy color *color-id* endpoint *ip-address* segment-list *segment-list-id*

Specifies the name of the target IPv4 or IPv6 SR policy.

Note:
The **sr-policy** target FEC type is supported under the OAM context and under **type-multi-line node** in the SAA context.

color *color-id* — Specifies the color ID.

Values 0 to 4294967295

endpoint *ip-address* — Specifies the endpoint address.

Values

ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

segment-list *segment-list-id* — Specifies the segment list ID.

Values 1 to 32

detail

Displays detailed information and allows the user to display hop 0 (that is, ingress) information. When this parameter is applied to static LSPs, the next hop 0 information is not displayed. This information is also not displayed if the **min-ttl** *min-label-ttl* value is greater than 1.

sr-te *lsp-name*

Specifies the name of the target SR-TE LSP, up to 64 characters.

static

Specifies the selection of the target FEC Stack sub-type "Static LSP”.

assoc-channel {ipv4 | non-ip | none}

Specifies the launched echo request’s usage of the Associated Channel (ACH) mechanism, when testing an MPLS-TP LSP.

Values **ipv4** — Use the Associated Channel mechanism with IP encapsulation, as described in RFC 6426, Section 3.2.

non-ip — Do not use an Associated Channel, as described in RFC 6426, Section 3.1.

none — Use the Associated Channel mechanism described in RFC 6426, Section 3.3.

path-type {active | working | protect}

Specifies the LSP path to test.

Values **active** — Specifies the currently active path. If MPLS-TP linear protection is configured on the LSP, then this is the path that is selected by the MPLS-TP PSC protocol for sending user plane traffic. If MPLS-TP linear protection is not configured, then this is the working path.

working — Specifies the working path of the MPLS-TP LSP.

protect — Specifies the protect path of the MPLS-TP LSP.

Default active

downstream-map-tlv

Specifies which format of the downstream mapping TLV to use in the LSP trace packet. The DSMAP TLV is the original format in RFC 4379 (obsoleted by RFC 8029). The DDMAP is the new enhanced format specified in RFC 6424 and RFC 8029. The user can also choose not to include the downstream mapping TLV by entering the value none. When lsp-trace is used on a MPLS-TP LSP (static option), it can only be executed if the control-channel is set to none. In addition, the DSMAP/DDMAP TLV is only included in the echo request message if the egress interface is either a numbered IP interface, or an unnumbered IP interface. The TLV is not included if the egress interface is of type **unnumbered-mpls-tp**.

Values ddmap: Sends a detailed downstream mapping TLV.

dsmap: Sends a downstream mapping TLV.

none: No mapping TLV is sent.

Default Inherited from global configuration of downstream mapping TLV in option **mpls-echo-request-downstream-map {dsmap | ddmap}**.

fc-name

Specifies the FC and profile parameters are used to indicate the forwarding class and profile of the MPLS echo request packet.

When an MPLS echo request packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the specified FC and profile parameter values. The marking of the packet EXP is dictated by the LSP-EXP mappings on the outgoing interface.

When the MPLS echo request packet is received on the responding node, The FC and profile parameter values are dictated by the LSP-EXP mappings of the incoming interface.

When an MPLS echo reply packet is generated in CPM and is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to the fc and profile parameter values determined by the classification of the echo request packet, which is being replied to, at the incoming interface. The marking of the packet's EXP is dictated by the LSP-EXP mappings on the outgoing interface. The ToS byte is not modified. Table 51: lsp-trace Request Packet and Behavior summarizes this behavior.

Table 51: lsp-trace Request Packet and Behavior

CPM (sender node)	Echo request packet: • packet {tos=1, fc1, profile1} • fc1 and profile1 are as entered by user in OAM command or default values • tos1 as per mapping of {fc1, profile1} to IP precedence in network egress QoS policy of outgoing interface
Outgoing interface (sender node)	Echo request packet: • pkt queued as {fc1, profile1} • ToS field=tos1 not remarked • EXP=exp1, as per mapping of {fc1, profile1} to EXP in network egress QoS policy of outgoing interface
Incoming interface (responder node)	Echo request packet: • packet {tos1, exp1} • exp1 mapped to {fc2, profile2} as per classification in network QoS policy of incoming interface
CPM (responder node)	Echo reply packet: • packet {tos=1, fc2, profile2}
Outgoing interface (responder node)	Echo reply packet: • pkt queued as {fc2, profile2} • ToS filed= tos1 not remarked (reply inband or out-of-band) • EXP=exp2, if reply is inband, remarked as per mapping of {fc2, profile2} to EXP in network egress QoS policy of outgoing interface

Incoming interface (sender node)	Echo reply packet: • packet {tos1, exp2} • exp2 mapped to {fc1, profile1} as per classification in network QoS policy of incoming interface

Values be, l2, af, l1, h2, ef, h1, nc

Default be

profile {in | out}

Specifies the profile state of the MPLS echo request packet.

Default out

interval

Specifies the number of seconds to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the *timeout* value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Values 1 to 10

Default **1**

no-response-count

Specifies the maximum number of consecutive MPLS echo requests, expressed as a decimal integer that do not receive a reply before the trace operation fails for a given TTL.

Values 1 to 255

Default 5

max-label-ttl

Specifies the maximum TTL value in the MPLS label for the LDP treetrace test, expressed as a decimal integer.

Values 1 to 255

Default 30

min-label-ttl

Specifies the minimum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Values 1 to 255

Default 1

probes-per-hop

Specifies the probes per hop.

Values 1 to 10

Default 1

octets

Specifies the size in octets, expressed as a decimal integer, of the MPLS echo request packet, including the IP header but not the label stack. The request payload is padded with zeros to the specified size. Note that an OAM command is not failed if the user entered a size lower than the minimum required to build the packet for the echo request message. The payload is automatically padded to meet the minimum size.

Values 1 to 9786

Default 1

src-ip-address *ip-address*

Specifies the source IP address. This option is used when an OAM packet must be generated from a different address than the node's system interface address. An example is when the OAM packet is sent over an LDP LSP and the LDP LSR-ID of the corresponding LDP session to the next-hop is set to an address other than the system interface address.

Values

ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

timeout

Specifies the time, in seconds, used to override the default *timeout* value and is the amount of time that the router waits for a message reply after sending the message request. Upon the expiration of the message time out, the requesting router assumes that the message response is not received. A **request timeout** message is displayed by the CLI for each message request sent that expires. Any response received after the request times out is silently discarded.

Values 1 to 60

Default 3

Platforms

7705 SAR Gen 2

Output

Output Example

```
*A:Dut-A# oam lsp-trace prefix 10.20.1.6/32 downstream-map-tlv ddmap path-
destination 127.0.0.1 detail lsp-trace to 10.20.1.6/
32: 0 hops min, 0 hops max, 152 byte packets
1  10.20.1.2  rtt=3.44ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=127.0.0.1 ifaddr=0 iftype=ipv4Unnumbered MRU=1500
           label[1]=131070 protocol=3(LDP)
2  10.20.1.4  rtt=4.65ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=127.0.0.1 ifaddr=0 iftype=ipv4Unnumbered MRU=1500
           label[1]=131071 protocol=3(LDP)
3  10.20.1.6  rtt=7.63ms rc=3(EgressRtr) rsc=1 *A:Dut-A#


*A:Dut-C# oam lsp-trace "p_1" detail
lsp-trace to p_1: 0 hops min, 0 hops max, 116 byte packets
1  10.20.1.2  rtt=3.46ms rc=8(DSRtrMatchLabel)
     DS 1: ipaddr 10.20.1.4 ifaddr 3 iftype 'ipv4Unnumbered' MRU=1500 label=131071
proto=4(RSVP-TE)
2  10.20.1.4  rtt=3.76ms rc=8(DSRtrMatchLabel)
     DS 1: ipaddr 10.20.1.6 ifaddr 3 iftype 'ipv4Unnumbered' MRU=1500 label=131071
proto=4(RSVP-TE)
3  10.20.1.6  rtt=5.68ms rc=3(EgressRtr)
*A:Dut-C#
```

lsp-trace over a numbered IP interface

```
A:Dut-C#
A:Dut-C# oam lsp-trace prefix 5.5.5.5/32 detail
lsp-trace to 5.5.5.5/32: 0 hops min, 0 hops max, 104 byte packets
1  6.6.6.6  rtt=2.45ms rc=8(DSRtrMatchLabel)
     DS 1: ipaddr=5.6.5.1 ifaddr=5.6.5.1 iftype=ipv4Numbered MRU=1564 label=131071
proto=3(LDP)
2  5.5.5.5  rtt=4.77ms rc=3(EgressRtr)
A:Dut-C#
```

lsp-trace over an unnumbered IP interface

```
*A:Dut-A# oam lsp-trace prefix 10.20.1.6/32 downstream-map-tlv ddmap path-
destination 127.0.0.1 detail lsp-trace to 10.20.1.6/
32: 0 hops min, 0 hops max, 152 byte packets
1  10.20.1.2  rtt=3.44ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=127.0.0.1 ifaddr=0 iftype=ipv4Unnumbered MRU=1500
           label[1]=131070 protocol=3(LDP)
2  10.20.1.4  rtt=4.65ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=127.0.0.1 ifaddr=0 iftype=ipv4Unnumbered MRU=1500
           label[1]=131071 protocol=3(LDP)
3  10.20.1.6  rtt=7.63ms rc=3(EgressRtr) rsc=1 *A:Dut-A#

*A:Dut-A# oam ldp-treetrace prefix 10.20.1.6/32

ldp-treetrace for Prefix 10.20.1.6/32:

         127.0.0.1, ttl =   3 dst =      127.1.0.255 rc = EgressRtr status = Done
    Hops:         127.0.0.1         127.0.0.1

         127.0.0.1, ttl =   3 dst =      127.2.0.255 rc = EgressRtr status = Done
    Hops:         127.0.0.1         127.0.0.1

ldp-treetrace discovery state: Done
ldp-treetrace discovery status: ' OK '
Total number of discovered paths: 2
```

```
Total number of failed traces: 0

lsp-trace of a LDP IPv6 prefix FEC



*A:Dut-A# oam lsp-trace prefix fc00::a14:106/128 path-destination ::ffff:127.0.0.1

lsp-trace to fc00::a14:106/128: 0 hops min, 0 hops max, 224 byte packets

1  fc00::a14:102  rtt=1.61ms rc=8(DSRtrMatchLabel) rsc=1

2  fc00::a14:103  rtt=3.51ms rc=8(DSRtrMatchLabel) rsc=1

3  fc00::a14:104  rtt=4.65ms rc=8(DSRtrMatchLabel) rsc=1

4  fc00::a14:106  rtt=7.02ms rc=3(EgressRtr) rsc=1



*A:Dut-A# oam lsp-trace prefix fc00::a14:106/128 path-destination ::ffff:127.0.0.2

lsp-trace to fc00::a14:106/128: 0 hops min, 0 hops max, 224 byte packets

1  fc00::a14:102  rtt=1.90ms rc=8(DSRtrMatchLabel) rsc=1

2  fc00::a14:103  rtt=3.10ms rc=8(DSRtrMatchLabel) rsc=1

3  fc00::a14:105  rtt=4.61ms rc=8(DSRtrMatchLabel) rsc=1

4  fc00::a14:106  rtt=6.45ms rc=3(EgressRtr) rsc=1
```

lsp-trace over SR-ISIS

```
*A:Dut-A# oam lsp-trace sr-isis prefix 10.20.1.6/32 igp-instance 0 detail
lsp-trace to 10.20.1.6/32: 0 hops min, 0 hops max, 108 byte packets
1  10.20.1.2  rtt=1220323ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.4.4 ifaddr=10.10.4.4 iftype=ipv4Numbered MRU=1496
           label[1]=26406 protocol=6(ISIS)
2  10.20.1.4  rtt=1220323ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.9.6 ifaddr=10.10.9.6 iftype=ipv4Numbered MRU=1496
           label[1]=26606 protocol=6(ISIS)
3  10.20.1.6  rtt=1220324ms rc=3(EgressRtr) rsc=1
```

```
*A:Dut-E# oam lsp-trace prefix 10.20.1.2/32 detail downstream-map-tlv ddmap
lsp-trace to 10.20.1.2/32: 0 hops min, 0 hops max, 108 byte packets
1  10.20.1.3  rtt=3.25ms rc=15(LabelSwitchedWithFecChange) rsc=1
     DS 1: ipaddr=10.10.3.2 ifaddr=10.10.3.2 iftype=ipv4Numbered MRU=1496
           label[1]=26202 protocol=6(ISIS)
           fecchange[1]=POP  fectype=LDP IPv4 prefix=10.20.1.2 remotepeer=0.0.0.0 (Unknown)
           fecchange[2]=PUSH fectype=SR Ipv4 Prefix prefix=10.20.1.2 remotepeer=10.10.3.2
2  10.20.1.2  rtt=4.32ms rc=3(EgressRtr) rsc=1
*A:Dut-E#
```

```
*A:Dut-B# oam lsp-trace prefix 10.20.1.5/32 detail downstream-map-tlv ddmap sr-isis
lsp-trace to 10.20.1.5/32: 0 hops min, 0 hops max, 108 byte packets
1  10.20.1.3  rtt=2.72ms rc=15(LabelSwitchedWithFecChange) rsc=1
     DS 1: ipaddr=10.11.5.5 ifaddr=10.11.5.5 iftype=ipv4Numbered MRU=1496
           label[1]=262143 protocol=3(LDP)
           fecchange[1]=POP  fectype=SR Ipv4 Prefix prefix=10.20.1.5 remotepeer=0.0.0.0
 (Unknown)
           fecchange[2]=PUSH fectype=LDP IPv4 prefix=10.20.1.5 remotepeer=10.11.5.5
```

```
2  10.20.1.5  rtt=4.43ms rc=3(EgressRtr) rsc=1
```

lsp-trace over SR policy

```
# ipv4 sr-policy lsp-trace
*A:Dut-A# oam lsp-trace sr-policy color 2 endpoint 10.20.1.6 downstream-map-tlv ddmap path-
destination 127.1.1.1 detail lsp-trace to color 2 endpoint 10.20.1.6: 0 hops min, 0 hops max,
 188 byte packets
1  10.20.1.2  rtt=1220323ms rc=3(EgressRtr) rsc=4
1  10.20.1.2  rtt=1220323ms rc=8(DSRtrMatchLabel) rsc=3
     DS 1: ipaddr=10.10.3.3 ifaddr=10.10.3.3 iftype=ipv4Numbered MRU=1496
           label[1]=28303 protocol=6(ISIS)
           label[2]=28305 protocol=0(Unknown)
           label[3]=28506 protocol=0(Unknown)
     DS 2: ipaddr=10.10.12.3 ifaddr=10.10.12.3 iftype=ipv4Numbered MRU=1496
           label[1]=28303 protocol=6(ISIS)
           label[2]=28305 protocol=0(Unknown)
           label[3]=28506 protocol=0(Unknown)
2  10.20.1.3  rtt=1220323ms rc=3(EgressRtr) rsc=3
2  10.20.1.3  rtt=1220324ms rc=8(DSRtrMatchLabel) rsc=2
     DS 1: ipaddr=10.10.5.5 ifaddr=10.10.5.5 iftype=ipv4Numbered MRU=1496
           label[1]=28505 protocol=6(ISIS)
           label[2]=28506 protocol=0(Unknown)
     DS 2: ipaddr=10.10.11.5 ifaddr=10.10.11.5 iftype=ipv4Numbered MRU=1496
           label[1]=28505 protocol=6(ISIS)
           label[2]=28506 protocol=0(Unknown)
3  10.20.1.5  rtt=1220325ms rc=3(EgressRtr) rsc=2
3  10.20.1.5  rtt=1220325ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.10.6 ifaddr=10.10.10.6 iftype=ipv4Numbered MRU=1496
           label[1]=28606 protocol=6(ISIS)
4  10.20.1.6  rtt=1220325ms rc=3(EgressRtr) rsc=1


# ipv6 sr-policy lsp-trace
*A:Dut-A# oam lsp-trace sr-policy color 500 endpoint fc00::a14:106 lsp-trace to color 500
 endpoint fc00::a14:106: 0 hops min, 0 hops max, 204 byte packets
1  fc00::a14:102  rtt=1220323ms rc=3(EgressRtr) rsc=4
1  fc00::a14:102  rtt=1220323ms rc=8(DSRtrMatchLabel) rsc=3
2  fc00::a14:103  rtt=1220323ms rc=3(EgressRtr) rsc=3 ^C *A:Dut-A# oam lsp-trace sr-policy
 color 500 endpoint fc00::a14:106 downstream-map-tlv ddmap path-destination ::ffff:127.1.1.1
 detail lsp-trace to color 500 endpoint fc00::a14:106: 0 hops min, 0 hops max, 260 byte packets
1  fc00::a14:102  rtt=1220323ms rc=3(EgressRtr) rsc=4
1  fc00::a14:102  rtt=1220323ms rc=8(DSRtrMatchLabel) rsc=3
     DS 1: ipaddr=fe80::c617:1ff:fe01:2 ifaddr=fe80::c617:1ff:fe01:2 iftype=ipv6Numbered MRU=
1496
           label[1]=28363 protocol=6(ISIS)
           label[2]=28365 protocol=0(Unknown)
           label[3]=28566 protocol=0(Unknown)
     DS 2: ipaddr=fe80::c415:ffff:fe00:141 ifaddr=fe80::c415:ffff:fe00:141 iftype=ipv6Numbered
 MRU=1496
           label[1]=28363 protocol=6(ISIS)
           label[2]=28365 protocol=0(Unknown)
           label[3]=28566 protocol=0(Unknown)
2  fc00::a14:103  rtt=1220323ms rc=3(EgressRtr) rsc=3
2  fc00::a14:103  rtt=1220324ms rc=8(DSRtrMatchLabel) rsc=2
     DS 1: ipaddr=fe80::c61e:1ff:fe01:1 ifaddr=fe80::c61e:1ff:fe01:1 iftype=ipv6Numbered MRU=
1496
           label[1]=28565 protocol=6(ISIS)
           label[2]=28566 protocol=0(Unknown)
     DS 2: ipaddr=fe80::c61e:1ff:fe01:5 ifaddr=fe80::c61e:1ff:fe01:5 iftype=ipv6Numbered MRU=
1496
           label[1]=28565 protocol=6(ISIS)
           label[2]=28566 protocol=0(Unknown)
```

```
3  fc00::a14:105  rtt=1220325ms rc=3(EgressRtr) rsc=2
3  fc00::a14:105  rtt=1220325ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=fe80::c420:1ff:fe01:2 ifaddr=fe80::c420:1ff:fe01:2 iftype=ipv6Numbered MRU=
1496
           label[1]=28666 protocol=6(ISIS)
4  fc00::a14:106  rtt=1220326ms rc=3(EgressRtr) rsc=1 *A:Dut-A#
```

lsp-trace over SR-TE

```
*A:Dut-A# oam lsp-trace sr-te "srteABCEDF" downstream-map-tlv ddmap detail
lsp-trace to srteABCEDF: 0 hops min, 0 hops max, 252 byte packets
1  10.20.1.2  rtt=1220323ms rc=3(EgressRtr) rsc=5
1  10.20.1.2  rtt=1220322ms rc=8(DSRtrMatchLabel) rsc=4
     DS 1: ipaddr=10.10.33.3 ifaddr=10.10.33.3 iftype=ipv4Numbered MRU=1520
           label[1]=3 protocol=6(ISIS)
           label[2]=262135 protocol=6(ISIS)
           label[3]=262134 protocol=6(ISIS)
           label[4]=262137 protocol=6(ISIS)
2  10.20.1.3  rtt=1220323ms rc=3(EgressRtr) rsc=4
2  10.20.1.3  rtt=1220323ms rc=8(DSRtrMatchLabel) rsc=3
     DS 1: ipaddr=10.10.5.5 ifaddr=10.10.5.5 iftype=ipv4Numbered MRU=1496
           label[1]=3 protocol=6(ISIS)
           label[2]=262134 protocol=6(ISIS)
           label[3]=262137 protocol=6(ISIS)
3  10.20.1.5  rtt=1220325ms rc=3(EgressRtr) rsc=3
3  10.20.1.5  rtt=1220325ms rc=8(DSRtrMatchLabel) rsc=2
     DS 1: ipaddr=10.10.11.4 ifaddr=10.10.11.4 iftype=ipv4Numbered MRU=1496
           label[1]=3 protocol=6(ISIS)
           label[2]=262137 protocol=6(ISIS)
4  10.20.1.4  rtt=1220324ms rc=3(EgressRtr) rsc=2
4  10.20.1.4  rtt=1220325ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.9.6 ifaddr=10.10.9.6 iftype=ipv4Numbered MRU=1496
           label[1]=3 protocol=6(ISIS)
5  10.20.1.6  rtt=1220325ms rc=3(EgressRtr) rsc=1
```

```
*A:Dut-A# oam lsp-trace sr-te "srteABCE_loose" downstream-map-tlv ddmap detail
lsp-trace to srteABCE_loose: 0 hops min, 0 hops max, 140 byte packets
1  10.20.1.2  rtt=1220323ms rc=3(EgressRtr) rsc=3
1  10.20.1.2  rtt=1220322ms rc=8(DSRtrMatchLabel) rsc=2
     DS 1: ipaddr=10.10.3.3 ifaddr=10.10.3.3 iftype=ipv4Numbered MRU=1496
           label[1]=26303 protocol=6(ISIS)
           label[2]=26305 protocol=6(ISIS)
     DS 2: ipaddr=10.10.12.3 ifaddr=10.10.12.3 iftype=ipv4Numbered MRU=1496
           label[1]=26303 protocol=6(ISIS)
           label[2]=26305 protocol=6(ISIS)
     DS 3: ipaddr=10.10.33.3 ifaddr=10.10.33.3 iftype=ipv4Numbered MRU=1496
           label[1]=26303 protocol=6(ISIS)
           label[2]=26305 protocol=6(ISIS)
2  10.20.1.3  rtt=1220323ms rc=3(EgressRtr) rsc=2
2  10.20.1.3  rtt=1220323ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.5.5 ifaddr=10.10.5.5 iftype=ipv4Numbered MRU=1496
           label[1]=26505 protocol=6(ISIS)
     DS 2: ipaddr=10.10.11.5 ifaddr=10.10.11.5 iftype=ipv4Numbered MRU=1496
           label[1]=26505 protocol=6(ISIS)
3  10.20.1.5  rtt=1220324ms rc=3(EgressRtr) rsc=1
```

```
*A:Dut-F# oam lsp-trace sr-te "srteFECBA_eth" path-destination ::ffff:127.1.1.1 detail
lsp-trace to srteFECBA_eth: 0 hops min, 0 hops max, 336 byte packets
1  fc00::a14:105  rtt=1220323ms rc=3(EgressRtr) rsc=4
1  fc00::a14:105  rtt=1220323ms rc=8(DSRtrMatchLabel) rsc=3
     DS 1: ipaddr=fe80::c618:2ff:fe01:1 ifaddr=fe80::c618:2ff:fe01:1 iftype=ipv6Numbered MRU=
1496
```

```
           label[1]=28363 protocol=6(ISIS)
           label[2]=74032 protocol=6(ISIS)
           label[3]=28261 protocol=6(ISIS)
     DS 2: ipaddr=fe80::c618:2ff:fe01:2 ifaddr=fe80::c618:2ff:fe01:2 iftype=ipv6Numbered MRU=
1496
           label[1]=28363 protocol=6(ISIS)
           label[2]=74032 protocol=6(ISIS)
           label[3]=28261 protocol=6(ISIS)
2  fc00::a14:103  rtt=1220324ms rc=3(EgressRtr) rsc=3
2  fc00::a14:103  rtt=1220324ms rc=8(DSRtrMatchLabel) rsc=2
     DS 1: ipaddr=fe80::c613:1ff:fe01:3 ifaddr=fe80::c613:1ff:fe01:3 iftype=ipv6Numbered MRU=
1496
           label[1]=3 protocol=6(ISIS)
           label[2]=28261 protocol=6(ISIS)
3  fc00::a14:102  rtt=1220325ms rc=3(EgressRtr) rsc=2
3  fc00::a14:102  rtt=1220325ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=fe80::c0ea:1ff:fe01:1 ifaddr=fe80::c0ea:1ff:fe01:1 iftype=ipv6Numbered MRU=
1496
           label[1]=28161 protocol=6(ISIS)
4  fc00::a14:101  rtt=1220325ms rc=3(EgressRtr) rsc=1
```

lsp-trace with sr-ospf3

```
# sr-ospf3 lsp-trace
*A:Dut-A# oam lsp-trace sr-ospf3 prefix fc00::a14:106/128 detail lsp-trace to fc00::a14:106/
128: 0 hops min, 0 hops max, 164 byte packets
1  fc00::a14:102  rtt=1.33ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=fe80::c61c:1ff:fe01:1 ifaddr=fe80::c61c:1ff:fe01:1 iftype=ipv6Numbered MRU=
1496
           label[1]=29466 protocol=5(OSPF)
2  fc00::a14:104  rtt=2.27ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=fe80::c420:1ff:fe01:1 ifaddr=fe80::c420:1ff:fe01:1 iftype=ipv6Numbered MRU=
1496
           label[1]=29666 protocol=5(OSPF)
3  fc00::a14:106  rtt=2.50ms rc=3(EgressRtr) rsc=1
```

First egress label with lsp-trace

```
lsp-trace to srteABCEDF_loose: 0 hops min, 0 hops max, 216 byte packets 0  10.20.1.1
     DS 1: ipaddr=10.101.1.2 ifaddr=10.101.1.2 iftype=ipv4Numbered MRU=1496
           label[1]=26202 protocol=6(ISIS)
           label[2]=26203 protocol=6(ISIS)
           label[3]=26305 protocol=6(ISIS)
           label[4]=26504 protocol=6(ISIS)
           label[5]=26406 protocol=6(ISIS)
1  10.20.1.2  rtt=1220323ms rc=3(EgressRtr) rsc=5
1  10.20.1.2  rtt=1220323ms rc=8(DSRtrMatchLabel) rsc=4
     DS 1: ipaddr=10.10.3.3 ifaddr=10.10.3.3 iftype=ipv4Numbered MRU=1496
           label[1]=26303 protocol=6(ISIS)
           label[2]=26305 protocol=6(ISIS)
           label[3]=26504 protocol=6(ISIS)
           label[4]=26406 protocol=6(ISIS)
     DS 2: ipaddr=10.10.12.3 ifaddr=10.10.12.3 iftype=ipv4Numbered MRU=1496
           label[1]=26303 protocol=6(ISIS)
           label[2]=26305 protocol=6(ISIS)
           label[3]=26504 protocol=6(ISIS)
           label[4]=26406 protocol=6(ISIS)
2  10.20.1.3  rtt=1220323ms rc=3(EgressRtr) rsc=4
2  10.20.1.3  rtt=1220323ms rc=8(DSRtrMatchLabel) rsc=3
     DS 1: ipaddr=10.10.5.5 ifaddr=10.10.5.5 iftype=ipv4Numbered MRU=1496
           label[1]=26505 protocol=6(ISIS)
           label[2]=26504 protocol=6(ISIS)
           label[3]=26406 protocol=6(ISIS)
```

```
    DS 2: ipaddr=10.10.11.5 ifaddr=10.10.11.5 iftype=ipv4Numbered MRU=1496
          label[1]=26505 protocol=6(ISIS)
          label[2]=26504 protocol=6(ISIS)
          label[3]=26406 protocol=6(ISIS)
```

16.163 lsp-wait

lsp-wait

Syntax

lsp-wait *lsp-wait* **lsp-initial-wait** [*initial-wait*] [**lsp-second-wait** *second-wait*]

Context

[Tree] (config>service>vprn>isis>timers lsp-wait)

Full Context

configure service vprn isis timers lsp-wait

Description

This command configures the throttling of IS-IS LSP-generation. Timers that determine when to generate the first, second, and subsequent LSPs are controlled with this command. Subsequent LSPs are generated at increasing intervals of the second **lsp-wait** timer until a maximum value is reached.

Note: The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity; for example, a configured value of 550 ms is internally rounded down to 500 ms.

The **no** form of this command reverts to the default value.

Default

lsp-wait 5000 lsp-initial-wait 10 lsp-second-wait 1000

Parameters

lsp-wait

Specifies the maximum interval, in milliseconds, between two consecutive occurrences of an LSP being generated.

Values 10 to 120000

initial-wait

Specifies the initial LSP generation delay, in milliseconds. Values less than 100 ms are internally rounded down to 0, so that there is no added initial LSP generation delay.

Values 10 to 100000

second-wait

Specifies the hold time, in milliseconds, between the first and second LSP generation.

Values 10 to 100000

Platforms

7705 SAR Gen 2

lsp-wait

Syntax

lsp-wait *lsp-wait* [**lsp-initial-wait** *initial-wait*] [**lsp-second-wait** *second-wait*]

Context

[Tree] (config>router>isis>timers lsp-wait)

Full Context

configure router isis timers lsp-wait

Description

This command configures the throttling of IS-IS LSP-generation. Timers that determine when to generate the first, second, and subsequent LSPs are controlled with this command. Subsequent LSPs are generated at increasing intervals of the second **lsp-wait** timer until a maximum value is reached.

Note: The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity; for example, a configured value of 550 ms is internally rounded down to 500 ms.

The **no** form of this command reverts to the default value.

Default

lsp-wait 5000 lsp-initial-wait 10 lsp-second-wait 1000

Parameters

lsp-wait

Specifies the maximum interval, in milliseconds, between two consecutive occurrences of an LSP being generated.

Values 10 to 120000

initial-wait

Specifies the initial LSP generation delay, in milliseconds. Values less than 100 ms are internally rounded down to 0, so that there is no added initial LSP generation delay.

Values 10 to 100000

second-wait

Specifies the hold time, in milliseconds, between the first and second LSP generation.

Values 10 to 100000

Platforms

7705 SAR Gen 2

16.164 lsr-label-route

lsr-label-route

Syntax

lsr-label-route [{**none** | **all**}]

Context

[Tree] (config>router>ttl-propagate lsr-label-route)

Full Context

configure router ttl-propagate lsr-label-route

Description

This command configures the TTL propagation for transit packets at a router acting as an LSR for a BGP label route.

When an LSR swaps the BGP label for a ipv4 prefix packet, therefore acting as a ABR, ASBR, or data-path Route-Reflector (RR) in the base routing instance, or swaps the BGP label for a vpn-ipv4 or vpn-ipv6 prefix packet, therefore acting as an inter-AS Option B VPRN ASBR or VPRN data path Route-Reflector (RR), the **all** value of this command enables TTL propagation of the decremented TTL of the swapped BGP label into all outgoing LDP or RSVP transport labels.

When an LSR swaps a label or stitches a label, it always writes the decremented TTL value into the outgoing swapped or stitched label. What this feature controls is whether this decremented TTL value is also propagated to the transport label stack pushed on top of the swapped or stitched label.

The none value reverts to the default mode which disables TTL propagation. This changes the existing default behavior which propagates the TTL to the transport label stack. When a customer upgrades, the new default becomes in effect. This command does not have a no version.

This feature also controls the TTL propagation at an LDP-BGP stitching LSR in the LDP to BGP stitching direction. It also controls the TTL propagation in Carrier Supporting Carrier (CsC) VPRN at both the CsC CE and CsC PE.

SR OS does not support ASBR or data path RR functionality for labeled IPv6 routes in the global routing instance (6PE). As such the CLI command of this feature has no impact on prefix packets forwarded in this context.

Default

lsr-label-route none

Parameters

none

Specifies that the TTL of the swapped label is not propagated into the transport label stack.

all

Specifies that the TTL of the swapped label is propagated into all labels of the transport label stack.

Platforms

7705 SAR Gen 2

16.165 lsr-load-balancing

lsr-load-balancing

Syntax

lsr-load-balancing *hashing-algorithm*

no lsr-load-balancing

Context

[Tree] (config>service>vprn>nw-if>load-balancing lsr-load-balancing)

Full Context

configure service vprn network-interface load-balancing lsr-load-balancing

Description

This command specifies whether the IP header is used in the LAG and ECMP LSR hashing algorithm. This is the per interface setting.

Default

no lsr-load-balancing

Parameters

lbl-only

Only the label is used in the hashing algorithm.

lbl-ip

The IP header is included in the hashing algorithm.

ip-only

The IP header is used exclusively in the hashing algorithm.

eth-encap-ip

The hash algorithm parses down the label stack and once it hits the bottom, the stack assumes Ethernet II non-tagged/dot1q or qinq header follows. At the expected Ethertype offset location, algorithm checks whether the value present is IPv4/v6 (0x0800 or 0x86DD). If the check passes, the hash algorithm checks the first nibble at the expected IP header location for IPv4/IPv6 (0x0100/0x0110). If the secondary check passes, the hash is performed using IP SA/DA fields in the expected IP header; otherwise (if any of the checks failed) label-stack hash is performed.

Platforms

7705 SAR Gen 2

lsr-load-balancing

Syntax

lsr-load-balancing {**lbl-only** | **lbl-ip** | **ip-only** | **eth-encap-ip** | **lbl-ip-l4-teid**}

no lsr-load-balancing

Context

[Tree] (config>router>if>load-balancing lsr-load-balancing)

Full Context

configure router interface load-balancing lsr-load-balancing

Description

This command specifies whether the IP header is used in the LAG and ECMP LSR hashing algorithm. This is the per interface setting.

Default

no lsr-load-balancing

Parameters

lbl-only

Specifies that only the label is used in the hashing algorithm

lbl-ip

Specifies that only the IP header is included in the hashing algorithm.

ip-only

Specifies that only the IP header is used exclusively in the hashing algorithm

eth-encap-ip

Specifies that the hash algorithm parses down the label stack and once it hits the bottom, the stack assumes Ethernet II non-tagged/dot1q or qinq header follows. At the expected

Ethertype offset location, algorithm checks whether the value present is IPv4/v6 (0x0800 or0x86DD). If the check passes, the hash algorithm checks the first nibble at the expected IP header location for IPv4/IPv6 (0x0100/0x0110). If the secondary check passes, the hash is performed using IP SA/DA fields in the expected IP header; otherwise (any of the check failed) label-stack hash is performed.

lbl-ip-l4-teid

Specifies that this hashing algorithm hashes based on label, IP header, Layer 4 header and GTP header (TEID) in order. The algorithm uses all the supported headers that are found in the header fragment of incoming traffic.

Platforms

7705 SAR Gen 2

lsr-load-balancing

Syntax

lsr-load-balancing *hashing-algorithm*

no lsr-load-balancing

Context

[Tree] (config>system>load-balancing lsr-load-balancing)

Full Context

configure system load-balancing lsr-load-balancing

Description

This command configures system-wide LSR load balancing. Hashing can be enabled on the label stack, IP header, or both. The hashing can be at an LSR for spraying labeled IP packets over multiple equal-cost paths, or over multiple links of a LAG group.

For IPv4 packets, the LSR hash routine operates on the label stack and the IP header. An LSR considers a packet to be IPv4 if the first nibble following the bottom of the label stack is 4. You can enable or disable hashing on the label stack and IPv4 and IPv6 headers at the system level or incoming network IP interface level.

Default

no lsr-load-balancing

Parameters

lbl-only

Specifies that only the label is used in the hashing algorithm.

lbl-ip

Specifies that the IP header is included in the hashing algorithm.

ip-only

Specifies that the IP header is used exclusively in the hashing algorithm.

eth-encap-ip

Specifies that the hash algorithm parses down the label stack and after it reaches the bottom, the stack assumes the Ethernet II non-tagged, dot1q, or QinQ header follows. At the expected Ethertype offset location, the algorithm checks whether the value present is IPv4/IPv6 (0x0800/0x86DD). If the check passes, the hash algorithm checks the first nibble at the expected IP header location for IPv4/IPv6 (0x0100/0x0110). If the secondary check passes, the algorithm performs the hash using the IP SA/DA fields in the expected IP header. If any of the checks fail, the label-stack hash is performed.

lbl-ip-l4-teid

Specifies that the hashing applies as follows for Layer 2 and Layer 3 encapsulated traffic:

- If an IPv4 or IPv6 header is found immediately after the MPLS label stack, the hashing includes label stack, source and destination IP addresses, TCP/UDP port numbers, and, if present, TEID values.
- If an IPv4 or IPv6 header is not found immediately after the MPLS label stack, the data plane searches for a valid Ethertype value for the IPv4 and IPv6 payload. If a valid Ethertype value is found and an IP header follows the Ethernet header, hashing includes the source and destination IP addresses, TCP/UDP port numbers, and, if present, TEID values.

Platforms

7705 SAR Gen 2

17 m Commands – Part I

17.1 mac

mac

Syntax

mac *ieee-address*

no mac

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>host-ident mac)

Full Context

configure subscriber-mgmt local-user-db ipoe host host-identification mac

Description

This command specifies the MAC address to match for a host lookup.

Note:
This command is only used when **mac** is configured as one of the **match-list** parameters.

The **no** form of this command removes the MAC address from the configuration.

Parameters

ieee-address

Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

Platforms

7705 SAR Gen 2

mac

Syntax

[**no**] **mac** *ieee-mac-address*

Context

[Tree] (config>service>vprn>if>ipv6>vrrp mac)

[Tree] (config>service>vprn>if mac)

[Tree] (config>service>vprn>if>vrrp mac)

[Tree] (config>service>vprn>nw-if mac)

Full Context

configure service vprn interface ipv6 vrrp mac

configure service vprn interface mac

configure service vprn interface vrrp mac

configure service vprn network-interface mac

Description

This command assigns a specific MAC address to a VPRN IP interface.

The **no** form of this command returns the MAC address of the IP interface to the default value.

Default

The physical MAC address associated with the Ethernet interface on which the SAP is configured.

Parameters

ieee-mac-address

Specifies the 48-bit MAC address for the static ARP in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

Platforms

7705 SAR Gen 2

mac

Syntax

[**no**] **mac** *ieee-address*

Context

[Tree] (config>service>vpls>mac-protect mac)

Full Context

configure service vpls mac-protect mac

Description

This command specifies the 48-bit IEEE 802.3 MAC address.

The **no** form of the command reverts to the default.

Parameters

ieee-address

Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

Platforms

7705 SAR Gen 2

mac

Syntax

mac *ieee-address*

no mac [*ieee-address*]

Context

[Tree] (config>service>ies>if mac)

Full Context

configure service ies interface mac

Description

This command assigns a specific MAC address to an IES IP interface.

For Routed Central Office (CO), a group interface has no IP address explicitly configured but inherits an address from the parent subscriber interface when needed. For example, a MAC will respond to an ARP request when an ARP is requested for one of the IPs associated with the subscriber interface through the group interface.

The **no** form of this command returns the MAC address of the IP interface to the default value.

Default

The physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system).

Parameters

ieee-address

Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff, where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

Platforms

7705 SAR Gen 2

mac

Syntax

mac *ieee-address*

no mac

Context

[Tree] (config>service>vpls>mcr-default-gtw mac)

Full Context

configure service vpls mcr-default-gtw mac

Description

This command relates to a system configured for Dual Homing in L2-TPSDA. It defines the MAC address used when the system sends out a gratuitous ARP on an active SAP after a ring heals or fails in order to attract traffic from subscribers on the ring with connectivity to that SAP.

The **no** form of this command reverts to the default.

Default

no mac

Parameters

ieee-address

Specifies the address in xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format (cannot be all zeros).

Platforms

7705 SAR Gen 2

mac

Syntax

mac *ieee-address*

no mac

Context

[Tree] (config>port>ethernet mac)

[Tree] (config>lag mac)

Full Context

configure port ethernet mac

configure lag mac

Description

This command assigns a specific MAC address to an Ethernet port, Link Aggregation Group (LAG), Ethernet tunnel, or BCP-enabled port or sub-port.

Only one MAC address can be assigned to a port. When multiple **mac** commands are entered, the last command overwrites the previous command. When the command is issued while the port is operational, IP will issue an ARP, if appropriate, and BPDUs are sent with the new MAC address.

The **no** form of this command returns the MAC address to the default value.

By default, a MAC address is assigned by the system from the chassis MAC address pool. The use of an all-zeroes MAC address indicates that an operational MAC address should be assigned from the chassis MAC address pool.

Default

mac 00:00:00:00:00:00

Parameters

ieee-address

Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

Platforms

7705 SAR Gen 2

mac

Syntax

[no] mac *ieee-address*

Context

[Tree] (config>service>proxy-arp-nd>mac-list mac)

Full Context

configure service proxy-arp-nd mac-list mac

Description

This command configures the proxy ARP or ND MAC address information.

The **no** form of the command deletes the MAC address.

Parameters

ieee-address

Specifies the MAC address added to the list. The MAC list can be empty or contain up to 10 addresses.

Values xx:xx:xx:xx:xx:xx
xx-xx-xx-xx-xx-xx

Platforms

7705 SAR Gen 2

mac

Syntax

mac *ieee-address* [**create**] **black-holemac** *ieee-address* [**create**] **sap** *sap-id* **monitor** {**fwd-status**}

mac *ieee-address* [**create**] **spoke-sdp** *sdp-id:vc-id* **monitor** {**fwd-status**}

no mac ieee-address

Context

[Tree] (config>service>vpls>static-mac mac)

Full Context

configure service vpls static-mac mac

Description

This command assigns a conditional static MAC address entry to an SPBM B-VPLS SAP/spoke-SDP allowing external MACs for single and multi-homed operation.

For the 7705 SAR Gen 2, this command also assigns a conditional static MAC address entry to an EVPN VPLS SAP/spoke-SDP.

Static MACs are used for PBB Epipe and I-VPLS services that may terminate external to SPBM. If this is configured under a Control B-VPLS the interface referenced will not use IS-IS for this neighbor. This may also be configured under a User B-VPLS where the corresponding interface is not supported under the Control B-VPLS.

Parameters

ieee-address

Specifies the static MAC address to an SPBM/sdp-binding interface.

Values 6-byte mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx) Cannot be all zeros.

sap-id

Specifies the SAP identifier.

sdp-id

Specifies the SDP identifier.

Values 1 to 17407

vc-id

Specifies the virtual circuit identifier.

Values 1 to 4294967295

create

Mandatory keyword used to create a static MAC.

fwd-status

Specifies that this static mac is based on the forwarding status of the SAP or spoke-SDP for multi-homed operation.

black-hole

Specifies for TLS FDB entries defined on a local SAP the value 'sap', remote entries defined on an SDP have the value 'sdp'.

Platforms

7705 SAR Gen 2

mac

Syntax

mac *ieee-address* [**mask** *six-byte-mask*]

no mac *ieee-address*

Context

[Tree] (config>service>mac-list mac)

Full Context

configure service mac-list mac

Description

This command adds a protected MAC address entry.

The **no** form of this command removes the protected MAC address entry.

Parameters

ieee-address

Specifies the address in xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format (cannot be all zeros), up to 30 characters.

six-byte-mask

Specifies the mask address in xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format (cannot be all zeros), up to 30 characters.

Platforms

7705 SAR Gen 2

mac

Syntax

mac *ieee-address*

no mac

Context

[Tree] (config>service>vpls>interface mac)

Full Context

configure service vpls interface mac

Description

This command assigns a specific MAC address to a VPLS IP interface.

For Routed Central Office (CO), a group interface has no IP address explicitly configured but inherits an address from the parent subscriber interface when needed. For example, a MAC will respond to an ARP request when an ARP is requested for one of the IPs associated with the subscriber interface through the group interface.

The **no** form of this command returns the MAC address of the IP interface to the default value.

Default

mac

Parameters

ieee-address

Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

Default The system chassis MAC address.

Platforms

7705 SAR Gen 2

mac

Syntax

mac *name*

no mac

Context

[Tree] (config>service>template>vpls-sap-template>egress>filter-name mac)

[Tree] (config>service>template>vpls-sap-template>ingress>filter-name mac)

Full Context

configure service template vpls-sap-template egress filter-name mac

configure service template vpls-sap-template ingress filter-name mac

Description

This command associates an existing IP filter policy with the template.

Parameters

name

Specifies the MAC filter policy name, up to 64 characters.

Platforms

7705 SAR Gen 2

mac

Syntax

[**no**] **mac** *ieee-address*

Context

[Tree] (debug>service>id>igmp-snooping mac)

Full Context

debug service id igmp-snooping mac

Description

This command shows IGMP packets for the specified MAC address.

The **no** form of this command disables the MAC debugging.

Platforms

7705 SAR Gen 2

mac

Syntax

[**no**] **mac** *ieee-address*

Context

[Tree] (debug>service>id>mld mac)

Full Context

debug service id mld-snooping mac

Description

This command shows MLD packets for the specified MAC address.

The **no** form of this command disables the MAC debugging.

Platforms

7705 SAR Gen 2

mac

Syntax

mac *mac-address*

no mac

Context

[Tree] (config>service>ies>if>ipv6>vrrp mac)

Full Context

configure service ies interface ipv6 vrrp mac

Description

This command assigns a specific MAC address to an IES IP interface.

The **no** form of this command returns the MAC address of the IP interface to the default value.

Default

The physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system).

Parameters

mac-address

Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

Platforms

7705 SAR Gen 2

mac

Syntax

mac *mac-address*

no mac

Context

[Tree] (config>service>ies>if>vrrp mac)

Full Context

configure service ies interface vrrp mac

Description

This command assigns a specific MAC address to an IES IP interface.

The **no** form of this command returns the MAC address of the IP interface to the default value.

Default

The physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system).

Parameters

mac-address

Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff, where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

Platforms

7705 SAR Gen 2

mac

Syntax

mac *ieee-address*

no mac

Context

[Tree] (config>router>if mac)

Full Context

configure router interface mac

Description

This command assigns a specific MAC address to an IP interface. Only one MAC address can be assigned to an IP interface. When multiple **mac** commands are entered, the last command overwrites the previous command.

The **no** form of this command returns the MAC address of the IP interface to the default value.

Default

no mac

Parameters

ieee-address

Specifies the 48-bit MAC address for the IP interface in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff,* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

Platforms

7705 SAR Gen 2

mac

Syntax

mac *mac-address*

no mac

Context

[Tree] (config>router>if>ipv6>vrrp mac)

[Tree] (config>router>if>vrrp mac)

Full Context

configure router interface ipv6 vrrp mac

configure router interface vrrp mac

Description

This command sets an explicit MAC address used by the virtual router instance overriding the VRRP default derived from the VRID.

Changing the default MAC address is useful when an existing HSRP or other non-VRRP default MAC is in use by the IP hosts using the virtual router IP address. Many hosts do not monitor unessential ARPs and continue to use the cached non-VRRP MAC address after the virtual router becomes master of the host's gateway address.

The **mac** command sets the MAC address used in ARP responses when the virtual router instance is master. Routing of IP packets with *mac-address* as the destination MAC is also enabled. The **mac** setting must be the same for all virtual routers participating as a virtual router or indeterminate connectivity by the

attached IP hosts will result. All VRRP advertisement messages are transmitted with *mac-address* as the source MAC.

The command can be configured in both non-owner and owner **vrrp** nodal contexts.

The **mac** command can be executed at any time and takes effect immediately. When the virtual router MAC on a master virtual router instance changes, a gratuitous ARP is immediately sent with a VRRP advertisement message. If the virtual router instance is disabled or operating as backup, the gratuitous ARP and VRRP advertisement message is not sent.

The **no** form of the command restores the default VRRP MAC address to the virtual router instance.

Default

no mac

Parameters

mac-address

The 48-bit MAC address for the virtual router instance in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC, and non-IEEE reserved MAC addresses.

Platforms

7705 SAR Gen 2

mac

Syntax

mac *index* **name** *mac-name*

no mac *index*

Context

[Tree] (config>system>security>ssh>client-mac-list mac)

[Tree] (config>system>security>ssh>server-mac-list mac)

Full Context

configure system security ssh client-mac-list mac

configure system security ssh server-mac-list mac

Description

This command configures SSH MAC algorithms for SR OS as an SSH server or an SSH client.

The **no** form of this command removes the specified **mac** *index*.

Default

no mac *index*

Parameters

index

Specifies the index of the algorithm in the list.

Values 1 to 255

mac-name

Specifies the algorithm for calculating the message authentication code.

Values The following table lists the default client and server algorithms used for SSHv2.

Table 52: SSHv2 Default client and server algorithms

index	mac-name
200	hmac-sha2-512
210	hmac-sha2-256
215	hmac-sha1
220	hmac-sha1-96
225	hmac-md5
240	hmac-md5-96

Platforms

7705 SAR Gen 2

17.2 mac-address

mac-address

Syntax

mac-address *ieee-address*

no mac-address *ieee-address*

Context

[Tree] (config>port>ethernet>dot1x>per-host-authentication>allowed-source-macs mac-address)

Full Context

configure port ethernet dot1x per-host-authentication allowed-source-macs mac-address

Description

This command configures the host MAC address on the allowed MAC list.

The **no** form of the command deletes the MAC address from the list.

Default

no mac

Parameters

ieee-address

Specifies the MAC address.

Values xx:xx:xx:xx:xx:xx

Platforms

7705 SAR Gen 2

17.3 mac-advertisement

mac-advertisement

Syntax

[**no**] **mac-advertisement**

Context

[Tree] (config>service>vpls>bgp-evpn mac-advertisement)

Full Context

configure service vpls bgp-evpn mac-advertisement

Description

This command enables the advertisement in BGP of the learned macs on SAPs and SDP bindings. When the mac-advertisement is disabled, the local macs will be withdrawn in BGP.

Default

mac-advertisement

Platforms

7705 SAR Gen 2

17.4 mac-criteria

mac-criteria

Syntax

[**no**] **mac-criteria**

Context

[Tree] (config>qos>sap-ingress mac-criteria)

Full Context

configure qos sap-ingress mac-criteria

Description

This command is used to enter the node to create or edit policy entries that specify MAC criteria.

The **mac-criteria** based SAP ingress policies are used to select the appropriate ingress queue and corresponding forwarding class for matched traffic.

Router implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason, entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. When mac-criteria entries are removed from a SAP ingress policy, the mac-criteria is removed from all services where that policy is applied.

Platforms

7705 SAR Gen 2

17.5 mac-duplication

mac-duplication

Syntax

mac-duplication

Context

[Tree] (config>service>vpls>bgp-evpn mac-duplication)

Full Context

configure service vpls bgp-evpn mac-duplication

Description
Commands in this context configure the BGP EVPN MAC duplication parameters.

Platforms
7705 SAR Gen 2

17.6 mac-filter

mac-filter

Syntax
[**no**] **mac-filter**

Context
[Tree] (config>system>security>mgmt-access-filter mac-filter)

Full Context
configure system security management-access-filter mac-filter

Description
This command configures a management access MAC-filter.

Platforms
7705 SAR Gen 2

17.7 mac-list

mac-list

Syntax
mac-list *name* [**create**]

no mac-list *name*

Context
[Tree] (config>service>proxy-arp-nd mac-list)

Full Context
configure service proxy-arp-nd mac-list

Description

This command creates a list of MAC addresses that can be pointed at from the service for a specified IP. The list may contain up to 10 MAC addresses; an empty list is also allowed.

The MAC list allows on-the-fly changes, but a change in the list deletes the proxy entries for all the IPs using that list.

The **no** form of the command deletes the entire MAC-list. Deleting a MAC list is only possible if it is not referenced in the configuration.

Parameters

name

Specifies the name of the MAC address list, which can be up to 32 characters.

create

Mandatory keyword to create a MAC list.

Platforms

7705 SAR Gen 2

mac-list

Syntax

mac-list *name*

no mac-list

Context

[Tree] (config>service>vpls>proxy-nd>dynamic mac-list)

[Tree] (config>service>vpls>proxy-arp>dynamic mac-list)

Full Context

configure service vpls proxy-nd dynamic mac-list

configure service vpls proxy-arp dynamic mac-list

Description

This command associates a previously created MAC list to a dynamic IP. The MAC list is created using the **configure service proxy-arp-nd mac-list** command.

The **no** form of the command deletes the association of the MAC list and the dynamic IP.

Parameters

name

Specifies the name of the MAC list previously created using the **configure service proxy-arp-nd mac-list** command.

Platforms

7705 SAR Gen 2

mac-list

Syntax

mac-list *name* [**create**]

no mac-list *name*

Context

[Tree] (config>service mac-list)

Full Context

configure service mac-list

Description

This command configures a MAC list name. The MAC list is composed of a list of MAC addresses and masks, which along with Auto-Learn Mac Protect (ALMP) can be used to exclude certain MACs from being protected in a given object. This is typically used on SAPs and spoke SDPs configured with ALMP where certain MACs must be able to move to other objects (for example, VRRP virtual MACs).

The **no** form of this command removes the MAC list name.

Parameters

name

Specifies the MAC list name, up to 32 characters.

create

Keyword used to create the MAC list.

Platforms

7705 SAR Gen 2

17.8 mac-move

mac-move

Syntax

[**no**] **mac-move**

Context

[Tree] (config>service>vpls mac-move)

[Tree] (config>service>template>vpls-template mac-move)

Full Context

configure service vpls mac-move

configure service template vpls-template mac-move

Description

Commands in this context configure MAC move attributes. A sustained high re-learn rate can be a sign of a loop somewhere in the VPLS topology. Typically, STP detects loops in the topology, but for those networks that do not run STP, the mac-move feature is an alternative way to protect your network against loops.

When enabled in a VPLS, **mac-move** monitors the re-learn rate of each MAC. If the rate exceeds the configured maximum allowed limit, it disables the SAP where the source MAC was last seen. The SAP can be disabled permanently (until a **shutdown**/**no shutdown** command is executed) or for a length of time that grows linearly with the number of times the specified SAP was disabled. You have the option of marking a SAP as non-blockable in the **config>service>vpls>sap>limit-mac-move** or **config>service>vpls>spoke-sdp>limit-mac-move** contexts. This means that when the re-learn rate has exceeded the limit, another (blockable) SAP will be disabled instead.

The **mac-move** command enables the feature at the service level for SAPs and spoke-SDPs, as only those objects can be blocked by this feature. Mesh SDPs are never blocked, but their re-learn rates (sap-to-mesh/spoke-to-mesh or vice versa) are still measured.

The operation of this feature is the same on the SAP and spoke-SDP. For example, if a MAC address moves from SAP to SAP, from SAP to spoke-SDP, or between spoke-SDPs, one will be blocked to prevent thrashing. If the MAC address moves between a SAP and mesh SDP or spoke-SDP and mesh SDP combinations, the respective SAP or spoke-SDP will be blocked.

mac-move will disable a VPLS port when the number of relearns detected has reached the number of relearns needed to reach the move-frequency in the 5-second interval. For example, when the move-frequency is configured to 1 (relearn per second) mac-move will disable one of the VPLS ports when 5 relearns were detected during the 5-second interval because then the average move-frequency of 1 relearn per second has been reached. This can already occur in the first second if the real relearn rate is 5 relearns per second or higher.

The **no** form of this command disables MAC move.

Platforms

7705 SAR Gen 2

17.9 mac-move-level

mac-move-level

Syntax

mac-move-level {**primary** | **secondary**| **tertiary**}

Context

[Tree] (config>service>template>vpls-sap-template mac-move-level)

Full Context

configure service template vpls-sap-template mac-move-level

Description

When a SAP is instantiated using vpls-sap-template, if the MAC move feature is enabled at VPLS level, the command mac-move-level indicates whether the sap should be populated as primary-port, secondary-port, or tertiary-port in the instantiated VPLS.

If configured to the default, SAP is populated as a tertiary-port.

Default

no mac-move-level

Platforms

7705 SAR Gen 2

17.10 mac-pinning

mac-pinning

Syntax

[**no**] **mac-pinning**

Context

[Tree] (config>service>vpls>spoke-sdp mac-pinning)

[Tree] (config>service>vpls>mesh-sdp mac-pinning)

[Tree] (config>service>vpls>sap mac-pinning)

[Tree] (config>service>vpls>endpoint mac-pinning)

Full Context

configure service vpls spoke-sdp mac-pinning

configure service vpls mesh-sdp mac-pinning

configure service vpls sap mac-pinning

configure service vpls endpoint mac-pinning

Description

This command disables re-learning of MAC addresses on other SAPs within the VPLS. The MAC address will remain attached to a given SAP for duration of its age-timer.

The age of the MAC address entry in the FDB is set by the age timer. If **mac-aging** is disabled on a given VPLS service, any MAC address learned on a SAP or SDP with **mac-pinning** enabled will remain in the FDB on this SAP or SDP forever.

Every event that would otherwise result in re-learning is logged (MAC address; original-SAP; new-SAP).

When a SAP or spoke SDP is part of a Residential Split Horizon Group (RSHG), MAC pinning is activated at creation of the SAP. Otherwise MAC pinning is not enabled by default.

The **no** form of the command enables re-learning of MAC addresses.

Note:
MAC addresses learned during DHCP address assignment (DHCP snooping enabled) are not impacted by this command. MAC-pinning for such addresses is implicit.

Default

no mac-pinning

Platforms

7705 SAR Gen 2

mac-pinning

Syntax

[**no**] **mac-pinning**

Context

[Tree] (config>service>pw-template mac-pinning)

Full Context

configure service pw-template mac-pinning

Description

Enabling this command will disable re-learning of MAC addresses on other SAPs within the service. The MAC address will remain attached to a given SAP for duration of its age-timer.

The age of the MAC address entry in the FDB is set by the age timer. If **mac-aging** is disabled on a given VPLS service, any MAC address learned on a SAP or SDP with **mac-pinning** enabled will remain in the FDB on this SAP or SDP forever. Every event that would otherwise result in re-learning will be logged (MAC address; original-SAP; new-SAP).

When a SAP or spoke SDP is part of a Residential Split Horizon Group (RSHG), MAC pinning is activated at creation of the SAP. Otherwise MAC pinning is not enabled by default.

Note:
For 7705 SAR Gen 2, MAC addresses learned during DHCP address assignment (DHCP snooping enabled) are not impacted by this command. MAC-pinning for such addresses is implicit.

Default

no mac-pinning

Platforms

7705 SAR Gen 2

17.11 mac-policy

mac-policy

Syntax

mac-policy *mac-policy-id* [**create**]

no mac-policy *mac-policy-id*

Context

[Tree] (config>macsec mac-policy)

Full Context

configure macsec mac-policy

Description

This command configures MAC address policy groups.

The **no** form of this command removes the MAC address policy group configuration.

Parameters

mac-policy-id

Specifies the value of the MAC address policy.

Values 0 to 4294967295

create

Mandatory keyword used to create the configuration.

Platforms

7705 SAR Gen 2

17.12 mac-protect

mac-protect

Syntax

[**no**] **mac-protect**

Context

[Tree] (config>service>vpls mac-protect)

Full Context

configure service vpls mac-protect

Description

This command indicates if this MAC is protected on the MAC protect list. When enabled, the agent will protect the MAC from being learned or re-learned on a SAP, spoke SDP or mesh SDP that has restricted learning enabled. The MAC protect list is used in conjunction with **restrict-protected-src**, **restrict-unprotected-dst** and **auto-learn-mac-protect**.

The **no** form of the command reverts to the default.

Platforms

7705 SAR Gen 2

17.13 mac-subnet-length

mac-subnet-length

Syntax

mac-subnet-length *subnet-length*

no mac-subnet-length

Context

[Tree] (config>service>vpls mac-subnet-length)

Full Context

configure service vpls mac-subnet-length

Description

This command specifies the number of bits to be considered when performing MAC learning (MAC source) and MAC switching (MAC destination). Specifically, this value identifies how many bits, starting from the beginning of the MAC address are used. For example, if the mask-value of 28 is used, MAC learning only performs a lookup for the first 28 bits of the source MAC address when comparing with existing FDB entries. Then, it installs the first 28 bits in the FDB while zeroing out the last 20 bits of the MAC address. When performing switching in the reverse direction, only the first 28 bits of the destination MAC address are used to perform a FDB lookup to determine the next hop.

The **no** form of this command switches back to full MAC lookup.

Default

mac-subnet-length 48

Parameters

subnet-length

Specifies the number of bits to be considered when performing MAC learning or MAC switching.

Values 24 to 48

Platforms

7705 SAR Gen 2

17.14 macsec

macsec

Syntax

macsec

Context

[Tree] (config macsec)

Full Context

configure macsec

Description

Commands in this context configure MACsec, including the MACsec MKA profile.

Platforms

7705 SAR Gen 2

macsec

Syntax

[**no**] **macsec**

Context

[Tree] (config>port>ethernet>dot1x macsec)

Full Context

configure port ethernet dot1x macsec

Description

This command configures MACsec under this port.

Platforms

7705 SAR Gen 2

17.15 macsec-encrypt

macsec-encrypt

Syntax

[**no**] **macsec-encrypt**

Context

[Tree] (config>macsec>connectivity-association macsec-encrypt)

Full Context

configure macsec connectivity-association macsec-encrypt

Description

This command specifies that all PDUs are encrypted and authenticated (ICV payload).

The **no** form of this command specifies that all PDUs are transmitted with cleartext, but still authenticated and have the trailing ICV.

Default

macsec-encrypt

Platforms

7705 SAR Gen 2

17.16 main-ct-retry-limit

main-ct-retry-limit

Syntax

main-ct-retry-limit *number*

no main-ct-retry-limit

Context

[Tree] (config>router>mpls>lsp main-ct-retry-limit)

[Tree] (config>router>mpls>lsp-template main-ct-retry-limit)

Full Context

configure router mpls lsp main-ct-retry-limit

configure router mpls lsp-template main-ct-retry-limit

Description

This command configures the maximum number of retries the LSP primary path should be retried with the LSP Diff-Serv main Class Type (CT).

When an unmapped LSP primary path goes into retry, it uses the main CT until the number of retries reaches the value of the new main-ct-retry-limit parameter. If the path did not come up, it must start using the backup CT at that point in time. By default, this parameter is set to infinite value. The new main-ct-retry-limit parameter has no effect on an LSP primary path which retries due to a failure event.

An unmapped LSP primary path is a path which has never received a Resv in response to the first Path message sent. This can occur when performing a "shut/no-shut" on the LSP or LSP primary path or when the node reboots. An unmapped LSP primary path goes into retry if the retry timer expired or the head-end node received a PathErr message before the retry timer expired.

If the user entered a value of the main-ct-retry-limit parameter that is greater than the value of the LSP retry-limit, the number of retries will still stop when the LSP primary path reaches the value of the LSP retry-limit. In other words, the meaning of the LSP retry-limit parameter is not changed and always represents the upper bound on the number of retries. The unmapped LSP primary path behavior applies to both CSPF and non-CSPF LSPs.

The **no** form of this command sets the parameter to the default value of zero (0) which means the LSP primary path will retry forever.

Default

no main-ct-retry-limit

Parameters

number

Specifies the number of times MPLS will attempt to re-establish the LSP primary path using the Diff-Serv main CT. Allowed values are integers in the range of zero (0) to 10,000, where zero indicates to retry infinitely.

Values 0 to 1000, integer

Platforms

7705 SAR Gen 2

17.17 maintenance-policy

maintenance-policy

Syntax

[**no**] **maintenance-policy** *maintenance-policy-name*

Context

[Tree] (config>router>segment-routing maintenance-policy)

Full Context

configure router segment-routing maintenance-policy

Description

This command configures a named maintenance policy that can be applied to SR Policy candidate paths that are either statically configured or imported via BGP. A maintenance policy is used to configure seamless BFD and protection for an SR Policy candidate path.

A maintenance policy must be administratively disabled in order to change any of the parameters.

A maintenance policy cannot be enabled unless a **mode**, **bfd-enable**, and **bfd-template** are configured.

If a maintenance-template is administratively disabled, then all candidate paths to which it is applied are deprogrammed from the data path.

The **no** form of this command removes the specified maintenance policy.

Parameters

maintenance-policy-name

Specifies the name of the maintenance policy, up to 32 characters and cannot start with a space or underscore.

Platforms

7705 SAR Gen 2

maintenance-policy

Syntax

[**no**] **maintenance-policy** *maintenance-policy-name*

Context

[Tree] (conf>router>segment-routing>sr-policies>policy maintenance-policy)

Full Context

configure router segment-routing sr-policies static-policy maintenance-policy

Description

This command applies a named maintenance policy to the static SR policy path. The maintenance policy must exist under the **configure router segment-routing** context.

The **no** form of this command removes the specified maintenance policy.

Parameters

maintenance-policy-name

Specifies the name of the maintenance policy, up to 32 characters and cannot start with a space or underscore.

Platforms

7705 SAR Gen 2

17.18 managed-configuration

managed-configuration

Syntax

[**no**] **managed-configuration**

Context

[Tree] (config>router>router-advert>if managed-configuration)

[Tree] (config>service>vprn>router-advert>if managed-configuration)

Full Context

configure router router-advertisement interface managed-configuration

configure service vprn router-advertisement interface managed-configuration

Description

This command sets or resets managed address configuration flag for this group-interface. This flag indicates that DHCPv6 is available for address configuration in addition to any address auto-configured using stateless address auto-configuration. See RFC 3315 for additional details.

The **no** form of this command reverts to the default.

Default

no managed-configuration

Platforms

7705 SAR Gen 2

17.19 managed-vlan-list

managed-vlan-list

Syntax

managed-vlan-list

Context

[Tree] (config>service>vpls>sap managed-vlan-list)

Full Context

configure service vpls sap managed-vlan-list

Description

Commands in this context configure VLAN ranges to be managed by a management VPLS. The list indicates, for each SAP, the ranges of associated VLANs that will be affected when the SAP changes state. This managed-vlan-list is not used when STP mode is MSTP in which case the vlan-range is taken from the **config>service>vpls>stp>msti** configuration.

This command is only valid when the VPLS in which it is entered was created as a management VPLS.

Platforms

7705 SAR Gen 2

17.20 management

management

Syntax

management [**create**]

no management

Context

[Tree] (config>service>vprn management)

Full Context

configure service vprn management

Description

Commands in this context configure node management within the VPRN.

Parameters

create

Keyword used to create a management server entry.

Platforms

7705 SAR Gen 2

management

Syntax

management

Context

[Tree] (config>system>security management)

Full Context

configure system security management

Description

Commands in this context allow access to management servers.

Platforms

7705 SAR Gen 2

17.21 management-access-filter

management-access-filter

Syntax

[**no**] **management-access-filter**

Context

[Tree] (config>system>security management-access-filter)

Full Context

configure system security management-access-filter

Description

This command creates the context to edit management access filters and to reset match criteria.

Management access filters control all traffic in and out of the CPM. They can be used to restrict management of the router by other nodes outside either specific (sub)networks or through designated ports.

Management filters, as opposed to other traffic filters, are enforced by system software.

The **no** form of this command removes management access filters from the configuration.

Platforms

7705 SAR Gen 2

17.22 management-interface

management-interface

Syntax

management-interface

Context

[Tree] (config>system management-interface)

Full Context

configure system management-interface

Description

Commands in this context configure the capabilities of router management interfaces such as CLI and NETCONF.

Platforms

7705 SAR Gen 2

management-interface

Syntax

management-interface

Context

[Tree] (config>system>security management-interface)

Full Context

configure system security management-interface

Description

Commands in this context configure the selection of a management interface for hash configuration. The management interfaces are **classic-cli**, **md-cli**, **netconf**, or **grpc**.

Platforms

7705 SAR Gen 2

17.23 manager

manager

Syntax

manager *manager-name* [**create**]

no manager *manager-name*

Context

[Tree] (config>system>management-interface>remote-management manager)

Full Context

configure system management-interface remote-management manager

Description

Commands configured in this context take precedence over command values specified directly in the **configure management-interface remote-management** context.

If a command is not configured in this context, the command setting is inherited from the higher level context.

The **no** form of this command removes the remote manager configuration.

Default

system-name

Parameters

manager-name

Specifies the name of the remote manager, up to 32 characters.

Platforms

7705 SAR Gen 2

17.24 manager-address

manager-address

Syntax

manager-address *ip-address* | *fqdn*

no manager-address

Context

[Tree] (config>system>management-interface>remote-management>manager manager-address)

Full Context

configure system management-interface remote-management manager manager-address

Description

This command configures the destination IP address or FQDN of the manager.

The no form of this command removes the configured IP address or FQDN of the configured manager.

Parameters

ip-address

Specifies the IP address, up to 255 characters.

fqdn

Specifies the FQDN, up to 255 characters.

Platforms

7705 SAR Gen 2

17.25 manager-port

manager-port

Syntax

manager-port *port*

no manager-port

Context

[Tree] (config>system>management-interface>remote-management>manager manager-port)

Full Context

configure system management-interface remote-management manager manager-port

Description

This command assigns a destination TCP port to be used for opening gRPC connections to the specified remote manager.

The **no** form of this command reverts the destination TCP port for the remote manager to the default gRPC port (57400).

Parameters

port

Specifies the TCP destination port.

Values 1 to 65535

Default 57400

Platforms

7705 SAR Gen 2

17.26 manual-keying

manual-keying

Syntax

[**no**] **manual-keying**

Context

[Tree] (config>service>vprn>if>sap>ipsec-tunnel manual-keying)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel manual-keying)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel manual-keying)

[Tree] (config>router>if>ipsec>ipsec-tunnel manual-keying)

Full Context

configure service vprn interface sap ipsec-tunnel manual-keying

configure service vprn interface ipsec ipsec-tunnel manual-keying

configure service ies interface ipsec ipsec-tunnel manual-keying

configure router interface ipsec ipsec-tunnel manual-keying

Description

This command configures Security Association (SA) for manual keying. When enabled, the command specifies whether this SA entry is created manually, by the user, or dynamically by the IPsec sub-system.

Platforms

7705 SAR Gen 2

17.27 mapping-server

mapping-server

Syntax

[**no**] **mapping-server**

Context

[Tree] (config>router>isis>segment-routing mapping-server)

Full Context

configure router isis segment-routing mapping-server

Description

Commands in this context configures the Segment Routing mapping server feature in an IS-IS instance.

SR mapping server enables the configuration and advertisement, via IS-IS, of the node SID index for IS-IS prefixes of routers which are in the LDP domain. This is performed in the router acting as a mapping server, which uses a prefix-SID sub-TLV within the SID/Label binding TLV in IS-IS.

The **no** form of this command deletes all node SID entries in the IS-IS instance.

Platforms

7705 SAR Gen 2

mapping-server

Syntax

[**no**] **mapping-server**

Context

[Tree] (config>router>ospf>segm-rtng mapping-server)

Full Context

configure router ospf segment-routing mapping-server

Description

Commands in this context configure the Segment Routing mapping server feature in an OSPF instance.

The mapping server feature allows the configuration and advertisement in OSPF of the node SID index for OSPF prefixes of routers which are in the LDP domain. This is performed in the router acting as a mapping server and using a prefix-SID sub-TLV within the Extended Prefix Range TLV in OSPF.

The **no** form of this command deletes all node SID entries in the OSPF instance.

Platforms

7705 SAR Gen 2

17.28 mask

mask

Syntax

mask type *ppp-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}

no mask type *ppp-match-type*

mask type *ipoe-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}

no mask type *ipoe-match-type*

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe mask)

Full Context

configure subscriber-mgmt local-user-db ipoe mask

Description

This command configures a mask for the specified match type. The masking is applied on the parameter when performing an LUDB lookup to identify a host.

The **no** form of this command removes the mask from the configuration.

Parameters

ppp-match-type

Specifies the parameter on which the mask should be applied for an LUDB lookup to identify a PPP host.

Values circuit-id, mac, remote-id, sap-id, service-name, username

ipoe-match-type

Specifies the parameter on which the mask should be applied for an LUDB lookup to identify an IPoE host.

The *prefix-string* and *suffix-string* command options are not supported when the *ipoe-match-type* value is set to **duid-en** or **duid-ll-llt**.

Values circuit-id, duid-en, duid-ll-llt, option60, remote-id, sap-id, string, system-id

prefix-string

Specifies a substring that is stripped of the start of the incoming parameter value before it is matched against the value configured in the LUDB host identification.

This string can only contain printable ASCII characters. The "*” character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

This command option is unsupported when the *ppp-match-type* equals mac.

Values up to 127 characters, "*”

prefix-length

Specifies the number of characters to remove from the start of the incoming parameter value before it is matched against the value configured in the LUDB host identification.

When used with the **mac** or **duid-ll-llt** parameter, it specifies the number of bits to remove from the start of the MAC address. For example, if the MAC address is 0a:0b:0c:0d:0e:0f, to obtain the last bit for matching purposes (match an odd or even MAC address), the prefix length is 47. The result in this example would be a binary number of 1 (0xf = 1111).

Values 1 to 127

suffix-string

Specifies a substring that is stripped of the end of the incoming parameter value before it is matched against the value configured in the LUDB host identification.

This string can only contain printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

This command option is unsupported when the *ppp-match-type* equals **mac**.

Values up to 127 characters

suffix-length

Specifies the number of characters to remove from the end of the incoming parameter value before it is matched against the value configured in the LUDB host identification.

When used with the **mac** or **duid-ll-llt** command option, the number of bits to remove from the end of the MAC address is specified.

Values 1 to 127

Platforms

7705 SAR Gen 2

mask

Syntax

mask *mask-value* [**type** {**included** | **excluded**}]

no mask

Context

[Tree] (config>system>security>snmp>view mask)

Full Context

configure system security snmp view mask

Description

The mask value and the mask type, along with the *oid-value* configured in the **view** command, determines the access of each sub-identifier of an object identifier (MIB subtree) in the view.

Each bit in the mask corresponds to a sub-identifier position. For example, the most significant bit for the first sub-identifier, the next most significant bit for the second sub-identifier, and so on. If the bit position on the sub-identifier is available, it can be included or excluded.

For example, the MIB subtree that represents MIB-II is 1.3.6.1.2.1. The mask that catches all MIB-II would be 0xfc or 0b11111100.

Only a single mask may be configured per view and OID value combination. If more than one entry is configured, each subsequent entry overwrites the previous entry.

The **no** form of this command removes the mask from the configuration.

Parameters

mask-value

The mask value associated with the OID value determines whether the sub-identifiers are included or excluded from the view. (Default: all 1s)

The mask can be entered either:

- In hex. For example, 0xfc.
- In binary. For example, 0b11111100.

Note:
If the number of bits in the bit mask is less than the number of sub-identifiers in the MIB subtree, then the mask is extended with ones until the mask length matches the number of sub-identifiers in the MIB subtree.

type

Specifies to include or exclude MIB subtree objects.

Values **included** - All MIB subtree objects that are identified with a 1 in the mask are available in the view.

excluded - All MIB subtree objects that are identified with a 1 in the mask are denied access in the view.

Default **included**

Platforms

7705 SAR Gen 2

17.29 mask-reply

mask-reply

Syntax

[**no**] **mask-reply**

Context

[Tree] (config>service>ies>if>icmp mask-reply)

[Tree] (config>service>vprn>nw-if>icmp mask-reply)

[Tree] (config>service>vprn>if>icmp mask-reply)

Full Context

configure service ies interface icmp mask-reply

configure service vprn network-interface icmp mask-reply

configure service vprn interface icmp mask-reply

Description

This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance replies to mask requests.

The **no** form of this command disables replies to ICMP mask requests on the router interface.

Default

mask-reply — Specifies to reply to ICMP mask requests.

Platforms

7705 SAR Gen 2

mask-reply

Syntax

[**no**] **mask-reply**

Context

[Tree] (config>router>if>icmp mask-reply)

Full Context

configure router interface icmp mask-reply

Description

This command enables responses to ICMP mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

The **no** form of this command disables replies to ICMP mask requests on the router interface.

Default

mask-reply — Replies to ICMP mask requests.

Platforms

7705 SAR Gen 2

17.30 master-int-inherit

master-int-inherit

Syntax

[**no**] **master-int-inherit**

Context

[Tree] (config>service>ies>if>ipv6>vrrp master-int-inherit)

Full Context

configure service ies interface ipv6 vrrp master-int-inherit

Description

This command allows the master instance to dictate the master down timer (non-owner context only).

Default

no master-int-inherit

Platforms

7705 SAR Gen 2

master-int-inherit

Syntax

[**no**] **master-int-inherit**

Context

[Tree] (config>service>ies>if>vrrp master-int-inherit)

Full Context

configure service ies interface vrrp master-int-inherit

Description

This command allows the master instance to dictate the master down timer (non-owner context only).

Default

no master-int-inherit

Platforms

7705 SAR Gen 2

master-int-inherit

Syntax

[**no**] **master-int-inherit**

Context

[Tree] (config>service>vprn>if>ipv6>vrrp master-int-inherit)

[Tree] (config>service>vprn>if>vrrp master-int-inherit)

Full Context

configure service vprn interface ipv6 vrrp master-int-inherit

configure service vprn interface vrrp master-int-inherit

Description

This command allows the master instance to dictate the master down timer (non-owner context only).

Default

no master-int-inherit

Platforms

7705 SAR Gen 2

master-int-inherit

Syntax

[**no**] **master-int-inherit**

Context

[Tree] (config>router>if>vrrp master-int-inherit)

[Tree] (config>router>if>ipv6>vrrp master-int-inherit)

Full Context

configure router interface vrrp master-int-inherit

configure router interface ipv6 vrrp master-int-inherit

Description

This command enables the virtual router instance to inherit the master VRRP router's advertisement interval timer which is used by backup routers to calculate the master down timer.

The **master-int-inherit** command is only available in the non-owner nodal context and is used to allow the current virtual router instance master to dictate the master down timer for all backup virtual routers. The **master-int-inherit** command has no effect when the virtual router instance is operating as master.

If **master-int-inherit** is not enabled, the locally configured **message-interval** must match the master's VRRP advertisement message advertisement interval field value or the message is discarded.

The **no** form of the command restores the default operating condition which requires the locally configured **message-interval** to match the received VRRP advertisement message advertisement interval field value. The virtual router instance does not inherit the master VRRP router's advertisement interval timer and uses the locally configured message interval.

Default

no master-int-inherit

Platforms

7705 SAR Gen 2

17.31 match

match

Syntax

match

Context

[Tree] (config>service>vprn>bgp>group>dynamic-neighbor match)

Full Context

configure service vprn bgp group dynamic-neighbor match

Description

This command configures match conditions for the dynamic neighbors.

Platforms

7705 SAR Gen 2

match

Syntax

[**no**] **match**

Context

[Tree] (config>service>vprn>log>filter>entry match)

Full Context

configure service vprn log filter entry match

Description

This command creates context to enter/edit match criteria for a filter entry. When the match criteria is satisfied, the action associated with the entry is executed.

If more than one match parameter (within one match statement) is specified, then all the criteria must be satisfied (AND functional) before the action associated with the match is executed.

Use the **match** command to display a list of the valid applications.

Match context can consist of multiple match parameters (application, event-number, severity, subject), but multiple **match** statements cannot be entered per entry.

The **no** form of this command removes the match criteria for the *entry-id*.

Default

no match

Platforms

7705 SAR Gen 2

match

Syntax

match [**protocol** *protocol-id*]

no match

Context

[Tree] (config>qos>sap-egress>ip-criteria>entry match)

[Tree] (config>qos>sap-ingress>ip-criteria>entry match)

Full Context

configure qos sap-egress ip-criteria entry match

configure qos sap-ingress ip-criteria entry match

Description

This command creates a context to configure match criteria for SAP QoS policy match criteria. When the match criteria have been satisfied, the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured, all criteria must be satisfied (AND function) before the action associated with the match is executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

It is possible that a SAP policy includes the **dscp** map command, the **dot1p** map command, and an IP match criteria. When multiple matches occur for the traffic, the order of precedence is used to arrive at the final action. The order of precedence is as follows:

1. 802.1p bits
2. DSCP
3. IP quintuple or MAC headers

The **no** form of this command removes the match criteria for the *entry-id*.

Parameters

protocol *protocol-id*

Specifies an IP protocol to be used as a SAP QoS policy match criterion.

The protocol type such as TCP / UDP / OSPF is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17)

Table 53: IP Protocol Names lists the IP protocols and their respective IDs and descriptions.

Values protocol-id: 0 to 255 protocol numbers accepted in decimal, hexadecimal, or binary

keywords: none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-frag, ipv6-icmp, ipv6-no-nxt, ipv6-opts, ipv6-route, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp

Table 53: IP Protocol Names

Protocol	Protocol ID	Description
icmp	1	Internet Control Message
igmp	2	Internet Group Management
ip	4	IP in IP (encapsulation)
tcp	6	Transmission Control
egp	8	Exterior Gateway Protocol
igp	9	Any private interior gateway (used by Cisco for their IGRP)
udp	17	User Datagram
rdp	27	Reliable Data Protocol
ipv6	41	IPv6
ipv6-route	43	Routing Header for IPv6

Protocol	Protocol ID	Description
ipv6-frag	44	Fragment Header for IPv6
idrp	45	Inter-Domain Routing Protocol
rsvp	46	Reservation Protocol
gre	47	General Routing Encapsulation
ipv6-icmp	58	ICMP for IPv6
ipv6-no-nxt	59	No Next Header for IPv6
ipv6-opts	60	Destination Options for IPv6
iso-ip	80	ISO Internet Protocol
eigrp	88	EIGRP
ospf-igp	89	OSPFIGP
ether-ip	97	Ethernet-within-IP Encapsulation
encap	98	Encapsulation Header
pnni	102	PNNI over IP
pim	103	Protocol Independent Multicast
vrrp	112	Virtual Router Redundancy Protocol
l2tp	115	Layer Two Tunneling Protocol
stp	118	Schedule Transfer Protocol
ptp	123	Performance Transparency Protocol
isis	124	ISIS over IPv4
crtp	126	Combat Radio Transport Protocol
crudp	127	Combat Radio User Datagram

Platforms

7705 SAR Gen 2

match

Syntax

match [**next-header** *next-header*]

no match

Context

[Tree] (config>qos>sap-ingress>ipv6-criteria>entry match)

[Tree] (config>qos>sap-egress>ipv6-criteria>entry match)

Full Context

configure qos sap-ingress ipv6-criteria entry match

configure qos sap-egress ipv6-criteria entry match

Description

This command creates a context to configure match criteria for ingress SAP QoS policy match IPv6 criteria. When the match criteria have been satisfied, the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured, all criteria must be satisfied (logical AND) before the action associated with the match is executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be created per entry.

A SAP ingress policy may include the **dscp** map command, the **dot1p** map command, and an IPv6 match criteria. When multiple matches occur for the traffic, the following order of precedence is used to arrive at the final action.

1. 802.1p bits
2. DSCP
3. IP quintuple or MAC headers

The **no** form of this command removes the match criteria for the *entry-id*.

Parameters

next-header

protocol-number, *protocol-name*

Specifies the IPv6 next header to match.

On the 7705 SAR Gen 2, the protocol type such as TCP, UDP, or OSPF is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6) and UDP(17).

protocol-number

Specifies the protocol number value to be configured as a match criterion.

Values [0 to 255]D

[0x0 to 0xFF]H

[0b0 to 0b11111111]B

protocol-name

Specifies the protocol name to be configured as a match criterion.

Values none, icmp, igmp, ip, tcp, egp, igp, udp, rdp, ipv6, ipv6-route, ipv6-frag, idrp, rsvp, gre, ipv6-icmp, ipv6-no-nxt, ipv6-opts, iso-ip, eigrp, ospf-igp, ether-ip, encap, pnni, pim, vrrp, l2tp, stp, ptp, isis, crtp, crudp, sctp

* - udp/tcp wildcard

Platforms

7705 SAR Gen 2

match

Syntax

match [**frame-type** {**802dot3** | **802dot2-llc** | **802dot2-snap** | **ethernet-II** | **atm**}]

no match

Context

[Tree] (config>qos>sap-ingress>mac-criteria>entry match)

Full Context

configure qos sap-ingress mac-criteria entry match

Description

This command creates a context for entering/editing match MAC criteria for ingress SAP QoS policy match criteria. When the match criteria have been satisfied, the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured, all criteria must be satisfied (AND function) before the action associated with the match will be executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

The **no** form of this command removes the match criteria for the *entry-id*.

Parameters

frame-type

The **frame-type** keyword configures an Ethernet frame type or an ATM frame type to be used for the MAC filter match criteria.

Values 802dot3, 802dot2-llc, 802dot2-snap, ethernet_II, atm

Default 802dot3

802dot3

Specifies the frame type is Ethernet IEEE 802.3.

802dot2-llc

Specifies the frame type is Ethernet IEEE 802.2 LLC.

802dot2-snap

Specifies the frame type is Ethernet IEEE 802.2 SNAP.

ethernet-II

Specifies the frame type is Ethernet Type II.

atm

Specifies the frame type as ATM cell. The user is not allowed to configure entries with frame type of atm and a frame type of other supported values in the same QoS policy.

Platforms

7705 SAR Gen 2

match

Syntax

match [**protocol** *protocol-id*]

no match

Context

[Tree] (config>qos>network>ingress>ip-criteria>entry match)

[Tree] (config>qos>network>egress>ip-criteria>entry match)

Full Context

configure qos network ingress ip-criteria entry match

configure qos network egress ip-criteria entry match

Description

This command creates a context to configure match criteria for a network QoS policy. When the match criteria have been satisfied, the action associated with it is executed.

If more than one match criteria (within one match statement) are configured, then all criteria must be satisfied before the associated action with the match is executed.

A match context can consist of multiple match criteria, but multiple match statements cannot be entered per entry.

A network QoS policy can include the DSCP map command, the dot1p map command (ingress only), the prec map command (egress only), and an IP match criteria. When multiple matches occur for the traffic, the order of precedence is used to arrive at the final action. The order of precedence is as follows:

- 802.1p bits (ingress only)
- DSCP
- prec (egress only)
- IP quintuple

The **no** form of this command removes the match criteria for the entry identifier.

Parameters

protocol *protocol-id*

Specifies an IP protocol to be used as an ingress or egress network QoS policy match criterion.

The protocol type is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), and UDP(17).

Values *protocol-id*: 0 to 255 protocol numbers accepted in decimal, hexadecimal, or binary

keywords: none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-frag, ipv6-icmp, ipv6-no-nxt, ipv6-opts, ipv6-route, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp

* — udp/tcp wildcard

Table 54: Protocol ID Descriptions lists the protocols and their protocol IDs and descriptions.

Table 54: Protocol ID Descriptions

Protocol	Protocol ID	Description
icmp	1	Internet Control Message
igmp	2	Internet Group Management
ip	4	IP in IP (encapsulation)
tcp	6	Transmission Control
egp	8	Exterior Gateway Protocol
igp	9	Any private interior gateway (used by Cisco for their IGRP)
udp	17	User Datagram
rdp	27	Reliable Data Protocol
ipv6	41	IPv6
ipv6-route	43	Routing Header for IPv6
ipv6-frag	44	Fragment Header for IPv6
idrp	45	Inter-Domain Routing Protocol
rsvp	46	Reservation Protocol
gre	47	General Routing Encapsulation
ipv6-icmp	58	ICMP for IPv6

Protocol	Protocol ID	Description
ipv6-no-nxt	59	No Next Header for IPv6
ipv6-opts	60	Destination Options for IPv6
iso-ip	80	ISO Internet Protocol
eigrp	88	EIGRP
ospf-igp	89	OSPFIGP
ether-ip	97	Ethernet-within-IP Encapsulation
encap	98	Encapsulation Header
pnni	102	PNNI over IP
pim	103	Protocol Independent Multicast
vrrp	112	Virtual Router Redundancy Protocol
l2tp	115	Layer Two Tunneling Protocol
stp	118	Schedule Transfer Protocol
ptp	123	Performance Transparency Protocol
isis	124	ISIS over IPv4
crtp	126	Combat Radio Transport Protocol
crudp	127	Combat Radio User Datagram

Platforms

7705 SAR Gen 2

match

Syntax

match [**next-header** *next-header*]

no match

Context

[Tree] (config>qos>network>ingress>ipv6-criteria>entry match)

[Tree] (config>qos>network>egress>ipv6-criteria>entry match)

Full Context

configure qos network ingress ipv6-criteria entry match

configure qos network egress ipv6-criteria entry match

Description

This command creates a context to configure match criteria for a network QoS policy match IPv6 criteria. When the match criteria have been satisfied, the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured, all criteria must be satisfied (logical AND) before the action associated with the match is executed.

A match context can consist of multiple match criteria, but multiple match statements cannot be created per entry.

A network policy can include the DSCP map command, the dot1p map command (ingress only), the prec map command (egress only), and an IPv6 match criteria. When multiple matches occur for the traffic, the following order of precedence is used to arrive at the final action.

- 802.1p bits (ingress only)
- DSCP
- prec (egress only)
- IP quintuple

The **no** form of this command removes the match criteria for the entry identifier.

Parameters

next-header

protocol-number, *protocol-name*

Specifies the next header to match.

The protocol type is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), and UDP(17).

protocol-number

Specifies the protocol number value to be configured as a match criterion.

Values [0 to 255]D

[0x0 to 0xFF]H

[0b0 to 0b11111111]B

protocol-name

Specifies the protocol name to be configured as a match criterion.

Values none, icmp, igmp, ip, tcp, egp, igp, udp, rdp, ipv6, ipv6-route, ipv6-frag, idrp, rsvp, gre, ipv6-icmp, ipv6-no-nxt, ipv6-opts, iso-ip, eigrp, ospf-igp, ether-ip, encap, pnni, pim, vrrp, l2tp, stp, ptp, isis, crtp, crudp, sctp

* - udp/tcp wildcard

Platforms

7705 SAR Gen 2

match

Syntax

match [**protocol** *protocol-id*]

match protocol none

no match

Context

[Tree] (config>filter>ip-exception>entry match)

Full Context

configure filter ip-exception entry match

Description

Commands in this context enter match criteria for the filter entry. When the match criteria have been satisfied the action associated with the match criteria is executed.

A **match** context may consist of multiple match criteria, but multiple **match** statements cannot be entered per entry. More precisely, the command can be entered multiple times but this only results in modifying the *protocol-id*. and does not affect the underlying match criteria configuration.

The **no** form of the command removes all the match criteria from the filter entry and sets the *protocol-id* of the match command to **none** (keyword). As per above, **match protocol none** is however not equivalent to **no match**.

Default

match protocol none

Parameters

protocol-id

Sets an IP protocol to be used as an IP filter match criterion. The protocol type, such as TCP or UDP, is identified by its respective protocol number.

Values protocol-number: [0..255]D

[0x0..0xFF]H

[0b0..0b11111111]B

protocol-name:0 to 255 in decimal format. Values can also be specified in hexadecimal format, in binary format, or using the following keywords:

IPv4 filter keywords: none (default), icmp, igmp, ip, tcp, egp, igp, udp, rdp, ipv6, ipv6-route, ipv6-frag, idrp, rsvp, gre, ipv6-icmp, ipv6-no-nxt, ipv6-opts, iso-ip, eigrp, ospf-igp, ether-ip, encap, pnni, pim, vrrp, l2tp, stp, ptp, isis, crtp, crudp, sctp

IP exception filter keywords: none, icmp, igmp, ospf-igp, pim, rsvp, tcp, udp, vrrp

* — udp/tcp wildcard

Table 55: Protocol ID Descriptions

Protocol	Protocol ID	Description
icmp	1	Internet Control Message
igmp	2	Internet Group Management
ip	4	IP in IP (encapsulation)
tcp	6	Transmission Control
egp	8	Exterior Gateway Protocol
igp	9	Any private interior gateway (used by Cisco for IGRP)
udp	17	User Datagram
rdp	27	Reliable Data Protocol
ipv6	41	IPv6
ipv6-route	43	Routing Header for IPv6
ipv6-frag	44	Fragment Header for IPv6
idrp	45	Inter-Domain Routing Protocol
rsvp	46	Reservation Protocol
gre	47	General Routing Encapsulation
ipv6-icmp	58	ICMP for IPv6
ipv6-no-nxt	59	No Next Header for IPv6
ipv6-opts	60	Destination Options for IPv6
iso-ip	80	ISO Internet Protocol
eigrp	88	EIGRP
ospf-igp	89	OSPFIGP
ether-ip	97	Ethernet-within-IP Encapsulation
encap	98	Encapsulation Header
pnni	102	PNNI over IP
pim	103	Protocol Independent Multicast

Protocol	Protocol ID	Description
vrrp	112	Virtual Router Redundancy Protocol
l2tp	115	Layer Two Tunneling Protocol
stp	118	Spanning Tree Protocol
ptp	123	Performance Transparency Protocol
isis	124	ISIS over IPv4
crtp	126	Combat Radio Transport Protocol
crudp	127	Combat Radio User Datagram
sctp	132	Stream Control Transmission Protocol

Platforms

7705 SAR Gen 2

match

Syntax

match [{**protocol** *protocol-id* | **protocol-list** *protocol-list-name*}]

match protocol none

no match

Context

[Tree] (config>filter>ip-filter>entry match)

Full Context

configure filter ip-filter entry match

Description

Commands in this context enter match criteria for the filter entry. When the match criteria have been satisfied, the action associated with the match criteria is executed.

A match context may consist of multiple match criteria, but multiple match statements cannot be created per entry. More precisely, the **protocol** command can be entered multiple times but this only results in modifying the *protocol-id*. Matching on more than one protocol can be achieved using the protocol-list match criteria in an IP filter policy.

The **no** form of the command removes all the match criteria from the filter entry and sets the *protocol-id* of the match command to **none**. However, **match protocol none** is not equivalent to **no match**.

Default

match protocol none

Parameters

protocol-id

protocol-number | *protocol-name*

protocol-number

Specifies the protocol number value to be configured as a match criterion. The value can be expressed as a decimal integer, or in hexadecimal or binary format.

Values [0..255]D, [0x0..0xFF]H, [0b0..0b11111111]B

protocol-name

Specifies the protocol name to be configured as a match criterion.

Values IPv4 filter keywords: none (default), icmp, igmp, ip, tcp, egp, igp, udp, rdp, ipv6, ipv6-route, ipv6-frag, idrp, rsvp, gre, ipv6-icmp, ipv6-no-nxt, ipv6-opts, iso-ip, eigrp, ospf-igp, ether-ip, encap, pnni, pim, vrrp, l2tp, stp, ptp, isis, crtp, crudp, sctp

***** — udp/tcp

Table 56: Protocol ID Descriptions

Protocol	Protocol ID	Description
icmp	1	Internet Control Message
igmp	2	Internet Group Management
ip	4	IP in IP (encapsulation)
tcp	6	Transmission Control
egp	8	Exterior Gateway Protocol
igp	9	Any private interior gateway (used by Cisco for IGRP)
udp	17	User Datagram
rdp	27	Reliable Data Protocol
ipv6	41	IPv6
ipv6-route	43	Routing Header for IPv6
ipv6-frag	44	Fragment Header for IPv6
idrp	45	Inter-Domain Routing Protocol
rsvp	46	Reservation Protocol

Protocol	Protocol ID	Description
gre	47	General Routing Encapsulation
ipv6-icmp	58	ICMP for IPv6
ipv6-no-nxt	59	No Next Header for IPv6
ipv6-opts	60	Destination Options for IPv6
iso-ip	80	ISO Internet Protocol
eigrp	88	EIGRP
ospf-igp	89	OSPFIGP
ether-ip	97	Ethernet-within-IP Encapsulation
encap	98	Encapsulation Header
pnni	102	PNNI over IP
pim	103	Protocol Independent Multicast
vrrp	112	Virtual Router Redundancy Protocol
l2tp	115	Layer Two Tunneling Protocol
stp	118	Spanning Tree Protocol
ptp	123	Performance Transparency Protocol
isis	124	ISIS over IPv4
crtp	126	Combat Radio Transport Protocol
crudp	127	Combat Radio User Datagram
sctp	132	Stream Control Transmission Protocol

protocol-list-name

Specifies the name of the protocol list, up to 32 characters.

Platforms

7705 SAR Gen 2

match

Syntax

match [**next-header** *next-header*]

no match

Context

[Tree] (config>filter>ipv6-exception>entry match)

Full Context

configure filter ipv6-exception entry match

Description

Commands in this context enter match criteria for the IPv6 filter exception. When the match criteria have been satisfied, the action associated with the match criteria is executed.

The **no** form of the command removes all the match criteria from the IPv6 filter exception.

Parameters

next-header

protocol-number, *protocol-name*

Specifies the next header to match.

protocol-number

Specifies the protocol number value to be configured as a match criterion.

Values [0 to 255]D

[0x0 to 0xFF]H

[0b0 to 0b11111111]B

protocol-name

Specifies the protocol name to be configured as a match criterion.

Values none, icmp, igmp, ip, tcp, egp, igp, udp, rdp, ipv6, ipv6-route, ipv6-frag, idrp, rsvp, gre, ipv6-icmp, ipv6-no-nxt, ipv6-opts, iso-ip, eigrp, ospf-igp, ether-ip, encap, pnni, pim, vrrp, l2tp, stp, ptp, isis, crtp, crudp, sctp

* - udp/tcp wildcard

Platforms

7705 SAR Gen 2

match

Syntax

match [{**next-header** *protocol-id* | **next-header-list** *protocol-list-name*}]

match next-header none

no match

Context

[Tree] (config>filter>ipv6-filter>entry match)

Full Context

configure filter ipv6-filter entry match

Description

Commands in this context enter match criteria for the filter entry. When the match criteria have been satisfied, the action associated with the match criteria is executed.

A match context may consist of multiple match criteria, but multiple match statements cannot be created per entry. More precisely, the **next-header** command can be entered multiple times, but this only results in modifying the *protocol-id*. Matching on more than one protocol can be achieved using the **next-header-list** match criteria.

The **no** form of the command removes all the match criteria from the filter entry and sets the *protocol-id* of the match command to **none**. However, **match next-header none** is not equivalent to **no match**.

Default

match next-header none

Parameters

next-header

protocol-number, *protocol-name*

Specifies the IPv6 next header to match. This parameter is analogous to the protocol parameter used in IPv4 filter match command.

protocol-number

Specifies the protocol number value to be configured as a match criterion.

Values [0 to 255]D

[0x0 to 0xFF]H

[0b0 to 0b11111111]B

protocol-name

Specifies the protocol name to be configured as a match criterion.

Values none, icmp, igmp, ip, tcp, egp, igp, udp, rdp, ipv6, ipv6-route, ipv6-frag, idrp, rsvp, gre, ipv6-icmp, ipv6-no-nxt, ipv6-opts, iso-ip, eigrp, ospf-igp, ether-ip, encap, pnni, pim, vrrp, l2tp, stp, ptp, isis, crtp, crudp, sctp

* - udp/tcp wildcard

protocol-list-name

Specifies the name of the protocol list, up to 32 characters.

Platforms

7705 SAR Gen 2

match

Syntax

[**no**] **match**

Context

[Tree] (config>log>filter>entry match)

Full Context

configure log filter entry match

Description

This command creates the context to enter and edit match criteria for a filter entry. When the match criteria is satisfied, the action associated with the entry is executed.

If more than one match parameter (within one match statement) is specified, then all the criteria must be satisfied (AND functional) before the action associated with the match is executed.

Use the **application** command to display a list of the valid applications.

Match context can consist of multiple match parameters (application, event-number, severity, subject), but multiple **match** statements cannot be entered per entry.

The **no** form of this command removes the match criteria for the *entry-id*.

Platforms

7705 SAR Gen 2

match

Syntax

match [**frame-type** *frame-type*]

no match

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry match)

Full Context

configure system security management-access-filter mac-filter entry match

Description

This command configures math criteria for this MAC filter entry.

Parameters

frame-type

Specifies the type of MAC frame to use as match criteria.

Values 802dot3 | 802dot2-llc | 802dot2-snap | 802dot1ag | ethernet_II

Default 802dot3

Platforms

7705 SAR Gen 2

match

Syntax

match *command-string*

no match

Context

[Tree] (config>system>security>profile>entry match)

Full Context

configure system security profile entry match

Description

This command configures a command or subtree commands in subordinate command levels are specified.

Evaluation stops when the first match is found, so subordinate levels cannot be modified with subsequent action commands. More specific action commands should be entered with a lower entry number or in a profile that is evaluated prior to this profile.

All commands below the hierarchy level of the matched command are denied.

The **no** form of this command removes a match condition.

Parameters

command-string

Specifies the CLI command or CLI tree level that is the scope of the profile entry.

Platforms

7705 SAR Gen 2

match

Syntax

match

Context

[Tree] (config>router>bgp>group>dynamic-neighbor match)

Full Context

configure router bgp group dynamic-neighbor match

Description

This command configures match conditions for the dynamic neighbors.

Platforms

7705 SAR Gen 2

17.32 match-list

match-list

Syntax

match-list *ipoe-match-type-1* [*ipoe-match-type-2*]

no match-list

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe match-list)

Full Context

configure subscriber-mgmt local-user-db ipoe match-list

Description

This command specifies the type of matching done to identify a host. There are different match-types for IPoE hosts of which a maximum of four can be specified.

The **no** form of this command removes all match criteria.

Parameters

match-type-x

Specifies up to four matching types to identify a host.

Values For IPoE: circuit-id, derived-id, dual-stack-remote-id, duid-en, duid-ll-llt, encap-tag-range, encap-tag-separate-range, ip, mac, option60, remote-id, sap-id, service-id, string, system-id

Note:
The format of **remote-id** in IPv6 is different than the format of **remote-id** in IPv4; IPv6 **remote-id** contains *enterprise-id* field that is also honored in matching.

circuit-id — Specifies to use the circuit ID to match against.

derived-id — Specifies the value extracted by Python script during processing of DHCP Discover/Solicit/Request/Renew/Rebind Messages (client to server bound messages). The value is stored in the DHCP Transaction Cache (DTC) in a variable named alc.dtc.derivedId. This value has a lifespan of a DHCP transaction (a single pair of messages exchanged between the client and the server, for example DHCP Discover and DHCP Offer).

dual-stack-remote-id — Specifies the enterprise-id in IPv6 remote-id is stripped off before LUDB matching is performed. Processing of IPv4 remote ID remains unchanged. This will allow a single host entry in LUDB for dual-stack host where host identification is performed based on the remote ID field.

duid-en — Specifies to match against the concatenation of the enterprise number and identifier fields of DHCPv6 option CLIENTID (1) with DUID type = 2 (assigned by vendor based on the enterprise number) in the DHCPv6 client message.

duid-ll-llt — Specifies to match against the link-layer address field of DHCPv6 option CLIENTID (1) with DUID type = 3 (based on link-layer address) or DUID type = 1 (based on link-layer address plus time) and hardware type = 1 (Ethernet) in the DHCPv6 client message. For DUID type = 1, the time field is ignored.

encap-tag-separate-range — Specifies the match encapsulation inner and outer tag in two separate ranges.

encap-tag-range — Specifies to match tag ranges for inner and outer tags.

ip — Specifies the source IPv4/IPv6 address of a data-trigger packet.

mac — Specifies to use the MAC address to match against.

option-60 — Specifies to use Option60 to match against.

remote-id — Specifies to use the remote ID to match against.

sap-id — Specifies the SAP ID on which DHCPv4 packet are received. The SAP ID is inserted as ALU VSO (82,9,4) by the DHCPv4 relay in router. This is enabled via configuration under the vendor-specific-option CLI hierarchy of the DHCPv4 relay. Since the **dhcp-relay** configuration is enabled under the group interface CLI hierarchy, the group interface and the service ID must be known before the SAP ID can be used for LUDB match.

service-id — Specifies the service ID of the ingress SAP for DHCPv4 packets. The service ID is inserted as ALU VSO (82,9,3) by the DHCPv4 relay in router. This is enabled via configuration under the vendor-specific-option CLI hierarchy of the DHCPv4 relay.

string — Specifies the custom string configured under the vendor-specific-option CLI hierarchy of the DHCPv4 relay. The string is inserted as ALU VSO (82,9,5) by the DHCPv4 relay in router. Since the **dhcp-relay** configuration is enabled under the group-interface CLI hierarchy, the group-interface and the service ID must be known before the string can be used for LUDB match.

system-id — Specifies the system ID of the node name configured under the **system>name** CLI hierarchy. The system ID is inserted as ALU VSO (82,9,1) by the

DHCPv4 relay in router. This is enabled via configuration under the vendor-specific-option CLI hierarchy of the DHCPv4 relay. Since the **dhcp-relay** configuration is enabled under the group interface CLI hierarchy, the group interface and the service ID must be known before the system ID can be used for LUDB match.

Platforms

7705 SAR Gen 2

match-list

Syntax

match-list

Context

[Tree] (config>ipsec>client-db match-list)

Full Context

configure ipsec client-db match-list

Description

This command enables the match list context on a client database. The match list defines the match input used during IPsec's tunnel setup. If there are multiple inputs configured in the match list, then they all must have matches before the system considers a client entry is a match.

Platforms

7705 SAR Gen 2

match-list

Syntax

match-list

Context

[Tree] (config>qos match-list)

Full Context

configure qos match-list

Description

This command is used to enter the context to create or edit match lists used in QoS policies.

Platforms

7705 SAR Gen 2

match-list

Syntax

match-list

Context

[Tree] (config>filter match-list)

Full Context

configure filter match-list

Description

This command enables the configuration context for match lists to be used in filter policies (IOM/FP and CPM).

Platforms

7705 SAR Gen 2

17.33 match-peer-id-to-cert

match-peer-id-to-cert

Syntax

[**no**] **match-peer-id-to-cert**

Context

[Tree] (config>ipsec>ike-policy match-peer-id-to-cert)

Full Context

configure ipsec ike-policy match-peer-id-to-cert

Description

This command enables checking the IKE peer's ID matches the peer's certificate when performing certificate authentication.

Default

no match-peer-id-to-cert

Platforms

7705 SAR Gen 2

17.34 match-qinq-dot1p

match-qinq-dot1p

Syntax

match-qinq-dot1p {**top** | **bottom**}

no match-qinq-dot1p

Context

[Tree] (config>service>ies>if>sap>ingress match-qinq-dot1p)

[Tree] (config>service>vpls>sap>ingress match-qinq-dot1p)

Full Context

configure service ies interface sap ingress match-qinq-dot1p

configure service vpls sap ingress match-qinq-dot1p

Description

This command specifies which dot1Q tag position dot1P bits in a QinQ encapsulated packet should be used to evaluate dot1P QoS classification.

The **match-qinq-dot1p** command allows the top or bottom PBits to be used when evaluating the applied sap-ingress QoS policy's dot1P entries. The **top** and **bottom** keywords specify which position should be evaluated for QinQ encapsulated packets.

By default, the bottom-most service delineating dot1Q tag's dot1P bits are used. Table 57: Default QinQ and TopQ SAP Dot1P Evaluation defines the default behavior for dot1P evaluation when the **match-qinq-dot1p** command is not executed.

Table 57: Default QinQ and TopQ SAP Dot1P Evaluation

Port/SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits

Port/SAP Type	Existing Packet Tags	PBits Used for Match
QinQ/TopQ	TopQ	TopQ PBits
QinQ/TopQ	TopQ BottomQ	TopQ PBits
QinQ/TopQ	TopQ BottomQ	BottomQ PBits

The **no** form of this command restores the default dot1p evaluation behavior for the SAP.

Default

no match-qinq-dot1p (no filtering based on p-bits)

(top or bottom must be specified to override the default QinQ dot1p behavior)

Parameters

top

The top parameter is mutually exclusive to the bottom parameter. When the **top** parameter is specified, the topmost PBits are used (if existing) to match any dot1p dot1p-value entries. Table 58: Top Position QinQ and TopQ SAP Dot1P Evaluation defines the dot1p evaluation behavior when the **top** parameter is specified.

Table 58: Top Position QinQ and TopQ SAP Dot1P Evaluation

Port/SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ/TopQ	TopQ	TopQ PBits
QinQ/TopQ	TopQ BottomQ	TopQ PBits
QinQ/QinQ	TopQ BottomQ	TopQ PBits

bottom

The **bottom** parameter is mutually exclusive to the **top** parameter. When the bottom parameter is specified, the bottom most PBits are used (if existing) to match any dot1p

dot1p-value entries. Table 59: Bottom Position QinQ and TopQ SAP Dot1P Evaluation defines the dot1p evaluation behavior when the bottom parameter is specified.

Table 59: Bottom Position QinQ and TopQ SAP Dot1P Evaluation

Port/SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	BottomQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ/TopQ	TopQ	TopQ PBits
QinQ/TopQ	TopQ BottomQ	BottomQ PBits
QinQ/QinQ	TopQ BottomQ	BottomQ PBits

Platforms

7705 SAR Gen 2

match-qinq-dot1p

Syntax

match-qinq-dot1p {**top** | **bottom**}

no match-qinq-dot1p de

Context

[Tree] (config>service>epipe>sap>ingress match-qinq-dot1p)

Full Context

configure service epipe sap ingress match-qinq-dot1p

Description

This command specifies which Dot1Q tag position Dot1P bits in a QinQ encapsulated packet should be used to evaluate Dot1P QoS classification.

The **match-qinq-dot1p** command allows the top or bottom PBits to be used when evaluating the applied sap-ingress QoS policy's Dot1P entries. The **top** and **bottom** keywords specify which position should be evaluated for QinQ encapsulated packets.

The setting also applies to classification based on the DE indicator bit.

The **no** form of this command reverts the dot1p and de bits matching to the default tag.

By default, the bottom most service delineating Dot1Q tags Dot1P bits are used. Table 60: Default QinQ and TopQ SAP Dot1P Evaluation defines the default behavior for Dot1P evaluation. Top or bottom must be specified to override the default QinQ dot1p behavior.

Table 60: Default QinQ and TopQ SAP Dot1P Evaluation

Port/SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

Default

no match-qinq-dot1p (no filtering based on p-bits)

Parameters

top

The top parameter is mutually exclusive to the bottom parameter. When the top parameter is specified, the top most PBits are used (if existing) to match any dot1p dot1p-value entries. Table 61: Top Position QinQ dpt1p Evaluation Behavior defines the dot1p evaluation behavior when the top parameter is specified.

Table 61: Top Position QinQ dpt1p Evaluation Behavior

Port/SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None

Port/SAP Type	Existing Packet Tags	PBits Used for Match
Null	Dot1P (VLAN ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	TopQ PBits

bottom

The bottom parameter and the top parameter are mutually exclusive. When the bottom parameter is specified, the bottom most PBits are used (if existing) to match any dot1p dot1p-value entries. Table 62: Bottom Position QinQ and TopQ SAP Dot1P Evaluation defines the dot1p evaluation behavior when the bottom parameter is specified.

Table 62: Bottom Position QinQ and TopQ SAP Dot1P Evaluation

Port/SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	BottomQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

Table 63: Egress SAP Types

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
Null	No preserved Dot1P bits	None
Null	Preserved Dot1P bits	Preserved tag PBits remarked using dot1p-value
Dot1Q	No preserved Dot1P bits	New PBits marked using dot1p-value
Dot1Q	Preserved Dot1P bits	Preserved tag PBits remarked using dot1p-value
TopQ	No preserved Dot1P bits	TopQ PBits marked using dot1p-value
TopQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits marked using dot1p-value, BottomQ PBits preserved
QinQ	No preserved Dot1P bits	TopQ PBits and BottomQ PBits marked using dot1p-value
QinQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits and BottomQ PBits marked using dot1p-value

The QinQ and TopQ SAP PBit/DEI bit marking follows the default behavior defined in the preceding table when **qinq-mark-top-only** is not specified.

The dot1p *dot1p-value* command must be configured without the qinq-mark-top-only parameter to remove the TopQ PBits only marking restriction.

A QinQ-encapsulated Ethernet port can have two different sap types:

For a TopQ SAP type, only the outer (top) tag is explicitly specified. For example, **sap 1/1/1:10.***

For QinQ SAP type, both inner (bottom) and outer (top) tags are explicitly specified. For example, **sap 1/1/1:10.100**.

Platforms

7705 SAR Gen 2

match-qinq-dot1p

Syntax

match-qinq-dot1p {**top** | **bottom**}

no match-qinq-dot1p

Context

[Tree] (config>service>vprn>if>sap>ingress match-qinq-dot1p)

Full Context

configure service vprn interface sap ingress match-qinq-dot1p

Description

This command specifies which Dot1Q tag position Dot1P bits in a QinQ encapsulated packet should be used to evaluate Dot1P QoS classification.

The **match-qinq-dot1p** command allows the top or bottom PBits to be used when evaluating the applied sap-ingress QoS policy's Dot1P entries. The **top** and **bottom** keywords specify which position should be evaluated for QinQ encapsulated packets.

The **no** form of this command restores the default dot1p evaluation behavior for the SAP.

By default, the bottom most service delineating Dot1Q tags Dot1P bits are used. Table 64: Dot1P Default Behavior defines the default behavior for Dot1P evaluation when the **match-qinq-dot1p** command is not executed.

Table 64: Dot1P Default Behavior

Port / SAP Type	Existing Packet Tags	PBits Used for Match
null	none	none
null	Dot1P (VLAN-ID 0)	Dot1P PBits
null	—	Dot1Q PBits
null	TopQ BottomQ	TopQ PBits
null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	none (Default SAP)	none
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

Default

no match-qinq-dot1p - No filtering based on p-bits.

top or **bottom** must be specified to override the default QinQ dot1p behavior.

Parameters

top

The top parameter is mutually exclusive to the bottom parameter. When the top parameter is specified, the top most PBits are used (if existing) to match any dot1p dot1p-value entries. Table 65: Dot1P Evaluation Behavior defines the dot1p evaluation behavior when the top parameter is specified.

Table 65: Dot1P Evaluation Behavior

Port / SAP Type	Existing Packet Tags	PBits Used for Match
null	none	none
null	Dot1P (VLAN-ID 0)	Dot1P PBits
null	Dot1Q	Dot1Q PBits
null	TopQ BottomQ	TopQ PBits
null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	none (Default SAP)	none
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits

bottom

The bottom parameter is mutually exclusive to the top parameter. When the bottom parameter is specified, the bottom most PBits are used (if existing) to match any **dot1p** *dot1p-value* entries. The following tables define the bottom position QinQ and TopQ SAP dot1p evaluation and the default dot1p explicit marking actions.

Table 66: Bottom Position QinQ and TopQ SAP Dot1P Evaluation

Port / SAP Type	Existing Packet Tags	PBits Used for Match
null	none	none
null	Dot1P (VLAN-ID 0)	Dot1P PBits
null	Dot1Q	Dot1Q PBits
null	TopQ BottomQ	BottomQ PBits
null	TopQ (No BottomQ)	TopQ PBits

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Dot1Q	none (default SAP)	none
Dot1Q	Dot1P (default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	BottomQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

Table 67: Default Dot1P Explicit Marking Actions

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
null	no preserved Dot1P bits	none
null	preserved Dot1P bits	preserved tag PBits remarked using dot1p-value
Dot1Q	no preserved Dot1P bits	new PBits marked using dot1p-value
Dot1Q	preserved Dot1P bits	preserved tag PBits remarked using dot1p-value
TopQ	no preserved Dot1P bits	TopQ PBits marked using dot1p-value
TopQ	preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits marked using dot1p-value, BottomQ PBits preserved
QinQ	no preserved Dot1P bits	TopQ PBits and BottomQ PBits marked using dot1p-value
QinQ	preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits and BottomQ PBits marked using dot1p-value

The **dot1p** *dot1p-value* command must be configured without the **qinq-mark-top-only** parameter to remove the TopQ PBits only marking restriction.

Platforms

7705 SAR Gen 2

17.35 max

max

Syntax

max *num-sessions*

no max

Context

[Tree] (config>service>nat>nat-policy>session-limits max)

Full Context

configure service nat nat-policy session-limits max

Description

This command configures the session limit of this policy. The session limit is the maximum number of sessions allowed for a subscriber associated with this policy.

Default

max 65535

Parameters

num-sessions

Specifies the session limit.

Values 1 to 65535

Platforms

7705 SAR Gen 2

17.36 max-advertisement-interval

max-advertisement-interval

Syntax

[**no**] **max-advertisement-interval** *seconds*

Context

[Tree] (config>service>vprn>router-advert>if max-advertisement-interval)

[Tree] (config>router>router-advert>if max-advertisement-interval)

Full Context

configure service vprn router-advertisement interface max-advertisement-interval

configure router router-advertisement interface max-advertisement-interval

Description

This command configures the maximum interval between sending router advertisement messages.

Default

max-advertisement-interval 600

Parameters

seconds

Specifies the maximum interval in seconds between sending router advertisement messages.

Values 4 to 1800

Platforms

7705 SAR Gen 2

17.37 max-age

max-age

Syntax

max-age *max-age*

no max-age [*max-age*]

Context

[Tree] (config>service>vpls>stp max-age)

[Tree] (config>service>template>vpls-template>stp max-age)

Full Context

configure service vpls stp max-age

configure service template vpls-template stp max-age

Description

This command indicates how many hops a BPDU can traverse the network starting from the root bridge. The message age field in a BPDU transmitted by the root bridge is initialized to 0. Each other bridge will

take the message_age value from BPDUs received on their root port and increment this value by 1. The message_age therefore reflects the distance from the root bridge. BPDUs with a message age exceeding max-age are ignored.

STP uses the max-age value configured in the root bridge. This value is propagated to the other bridges via the BPDUs.

The **no** form of this command returns the max age to the default value.

Default

max-age 20

Parameters

max-age

The max info age for the STP instance in seconds. Allowed values are integers in the range 6 to 40.

Platforms

7705 SAR Gen 2

17.38 max-auth-req

max-auth-req

Syntax

max-auth-req *max-auth-request*

Context

[Tree] (config>port>ethernet>dot1x max-auth-req)

Full Context

configure port ethernet dot1x max-auth-req

Description

This command configures the maximum number of times that the router will send an access request RADIUS message to the RADIUS server. If a reply is not received from the RADIUS server after the specified number attempts, the 802.1x authentication procedure is considered to have failed.

The **no** form of this command returns the value to the default.

Default

max-auth-req 2

Parameters

max-auth-request

The maximum number of RADIUS retries.

Values 1 to 10

Platforms

7705 SAR Gen 2

17.39 max-bulk-duration

max-bulk-duration

Syntax

max-bulk-duration *milliseconds*

no max-bulk-duration

Context

[Tree] (config>system>snmp max-bulk-duration)

Full Context

configure system snmp max-bulk-duration

Description

This command sets the maximum duration to process an SNMP request before bulk responses are returned to avoid a timeout on the management system when a lot of information is returned in the response.

Default

no max-bulk-duration

Parameters

milliseconds

Specifies the maximum duration to process requests before bulk responses are returned.

Values 100 to 5000

Platforms

7705 SAR Gen 2

17.40 max-burst

max-burst

Syntax

max-burst *number*

no max-burst

Context

[Tree] (config>router>rsvp>msg-pacing max-burst)

Full Context

configure router rsvp msg-pacing max-burst

Description

This command specifies the maximum number of RSVP messages that are sent in the specified period under normal operating conditions.

Default

max-burst 650

Parameters

number

Specifies the maximum number of RSVP messages to be sent in increments of 10.

Values 100 to 1000

Platforms

7705 SAR Gen 2

17.41 max-bypass-associations

max-bypass-associations

Syntax

max-bypass-associations *integer*

no max-bypass-associations

Context

[Tree] (config>router>mpls max-bypass-associations)

Full Context

configure router mpls max-bypass-associations

Description

This command allows the user to set a maximum number of LSP primary path associations with each manual or dynamic bypass LSP that is created in the system.

By default, a Point of Local Repair (PLR) node will associate a maximum of 1000 primary LSP paths with a given bypass before using the next available manual bypass or signaling a new dynamic bypass.

Note that a new bypass LSP may need to be signaled if the constraint of a given primary LSP path is not met by an existing bypass LSP even if the max-bypass-associations for this bypass LSP has not been reached.

The **no** form of this command reinstates the default value of this parameter.

Default

max-bypass-associations 1000

Parameters

integer

Configures the number of LSP primary path associations

Values 100 to 131072

Platforms

7705 SAR Gen 2

17.42 max-bypass-plr-associations

max-bypass-plr-associations

Syntax

max-bypass-plr-associations *plr-value*

no max-bypass-plr-associations

Context

[Tree] (config>router>mpls max-bypass-plr-associations)

Full Context

configure router mpls max-bypass-plr-associations

Description

This command enables the configuration of the maximum number of Points of Local Repair (PLRs) per RSVP-TE bypass LSP.

A PLR summarizes the constraints applied to the computation of the path of the bypass LSP. It consists of the avoid link/node constraint, and potentially other TE constraints such as exclude SRLG, that are needed to protect against the failure of the primary path of the RSVP-TE LSP that is associated with this bypass LSP.

Additional PLRs with the same avoid link/node constraint are associated with the same bypass to minimize the number of bypass LSPs created. This command controls the maximum number of such PLRs.

Because MPLS saves only the PLR constraints of the first LSP that triggered the dynamic bypass creation, subsequent LSPs for the same avoid link/node and with the non-strict bypass SRLG disjointness enabled may be associated with the same bypass. This is even in cases where there exists a bypass LSP path that strictly satisfies the SRLG constraint.

When the maximum PLRs per bypass is configured with a value of 1, MPLS triggers the signaling of a new dynamic bypass LSP for each new PLR and saves each PLR constraint separately with its own bypass. As a result, when MPLS re-optimizes a bypass LSP it guarantees that SRLG disjointness of that PLR are checked and enforced.

The **no** form of this command returns the command to its default value.

Default

max-bypass-plr-associations 16

Parameters

plr-value

Configures the number of LSP primary path associations

Values 1 to 16

Default 16

Platforms

7705 SAR Gen 2

17.43 max-channels-per-connection

max-channels-per-connection

Syntax

max-channels-per-connection *number-of-channels*

no max-channels-per-connection

Context

[Tree] (config>system>login-control>ssh max-channels-per-connection)

Full Context

configure system login-control ssh max-channels-per-connection

Description

This command configures the maximum number of channels supported on an SSH connection.

The **no** form of this command configures this value to 5, which is the default.

Default

max-channels-per-connection 5

Parameters

number-of-channels

Specifies the number of channels.

Values 1 to 50

Platforms

7705 SAR Gen 2

17.44 max-cleared

max-cleared

Syntax

max-cleared *maximum*

Context

[Tree] (config>system>alarms max-cleared)

Full Context

configure system alarms max-cleared

Description

This command configures the maximum number of cleared alarms that the system will store and display.

Default

max-cleared 500

Parameters

maximum

Specifies the maximum number of cleared alarms, up to 500.

Platforms

7705 SAR Gen 2

17.45 max-completed

max-completed

Syntax

max-completed *unsigned*

Context

[Tree] (config>system>script-control>script-policy max-completed)

Full Context

configure system script-control script-policy max-completed

Description

This command is used to configure the maximum number of script run history status entries to keep.

Default

max-completed 1

Parameters

unsigned

Specifies the maximum number of script run history status entries to keep.

Values 1 to 1500

Default 1

Platforms

7705 SAR Gen 2

17.46 max-conn-prefix

max-conn-prefix

Syntax

max-conn-prefix *count*

no max-conn-prefix

Context

[Tree] (config>test-oam>twamp>server>prefix max-conn-prefix)

Full Context

configure test-oam twamp server prefix max-conn-prefix

Description

This command configures the maximum number of control connections by clients with an IP address in a specific prefix. A new control connection is rejected if accepting it would cause either the prefix limit defined by this command or the server limit (max-conn-server) to be exceeded.

The **no** form of this command returns the value to the default.

Default

max-conn-prefix 32

Parameters

count

Specifies the maximum number of control connections.

Values 0 to 64

Default 32

Platforms

7705 SAR Gen 2

17.47 max-conn-server

max-conn-server

Syntax

max-conn-server *count*

no max-conn-server

Context

[Tree] (config>test-oam>twamp>server max-conn-server)

Full Context

configure test-oam twamp server max-conn-server

Description

This command configures the maximum number of TWAMP control connections from all TWAMP clients. A new control connection is rejected if accepting it would cause either this limit or a prefix limit (max-conn-prefix) to be exceeded.

The **no** form of this command returns the value to the default.

Default

max-conn-server 32

Parameters

count

Specifies the maximum number of control connections.

Values 0 to 64

Default 32

Platforms

7705 SAR Gen 2

17.48 max-drop-count

max-drop-count

Syntax

max-drop-count *count*

no max-drop-count

Context

[Tree] (config>service>sdp>keep-alive max-drop-count)

Full Context

configure service sdp keep-alive max-drop-count

Description

This command configures the number of consecutive SDP keepalive failed request attempts or remote replies that can be missed after which the SDP is operationally downed. If the **max-drop-count** consecutive keepalive request messages cannot be sent or no replies are received, the SDP-ID will be brought operationally down by the keepalive SDP monitoring.

The **no** form of this command reverts the **max-drop-count** *count* value to the default settings.

Default

max-drop-count 3

Parameters

count

Specifies the number of consecutive SDP keepalive requests that are failed to be sent or replies missed, expressed as a decimal integer.

Values 1 to 5

Platforms

7705 SAR Gen 2

17.49 max-ecmp-routes

max-ecmp-routes

Syntax

max-ecmp-routes *max-routes*

no max-ecmp-routes

Context

[Tree] (config>router>ldp max-ecmp-routes)

Full Context

configure router ldp max-ecmp-routes

Description

This command sets the maximum number of ECMP routes that LDP may use to resolve the next hop for a FEC.

Note:
The system-wide maximum number of ECMP routes is limited by the **config>router>ecmp** command. This command, under the LDP context, simply allows LDP to use more than 32 routes, if they are available in RTM or TTM. When configured, the actual number of ECMP routes used by LDP is therefore min[**config>router>ecmp**, **config>router>ldp>max-ecmp-routes**].

The **no** form of this command reverts to the default value.

Default

max-ecmp-routes 32

Parameters

max-routes

Specifies the maximum number of routes.

Values 1 to 64

Platforms

7705 SAR Gen 2

17.50 max-groups

max-groups

Syntax

max-groups *max-groups*

no max-groups

Context

[Tree] (config>service>vprn>igmp>if max-groups)

Full Context

configure service vprn igmp interface max-groups

Description

This command configures the maximum number of groups for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed.

The **no** form of this command removes the value.

Parameters

max-groups

Specifies the maximum number of groups for this interface.

Values 1 to 16000

Platforms

7705 SAR Gen 2

max-groups

Syntax

max-groups *value*

no max-groups

Context

[Tree] (config>service>vprn>mld>if max-groups)

Full Context

configure service vprn mld interface max-groups

Description

This command specifies the maximum number of groups for which MLD can have local receiver information based on received MLD reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed.

Default

0 (no limit to the number of groups)

Parameters

value

Specifies the maximum number of groups for this interface.

Values 1 to 16000

Platforms

7705 SAR Gen 2

max-groups

Syntax

max-groups *value*

no max-groups

Context

[Tree] (config>service>vprn>pim>if max-groups)

Full Context

configure service vprn pim interface max-groups

Description

This command configures the maximum number of groups for which PIM can have downstream state based on received PIM Joins on this interface. This does not include IGMP local receivers on the interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed. When this object has a value of 0, there is no limit to the number of groups.

Parameters

value

Specifies the maximum number of groups for this interface.

Values 1 to 16000

Platforms

7705 SAR Gen 2

max-groups

Syntax

max-groups *value*

no max-groups

Context

[Tree] (config>router>igmp>if max-groups)

Full Context

configure router igmp interface max-groups

Description

This command specifies the maximum number of groups for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed. This command is applicable for IPv4 and IPv6.

The **no** form of the command sets no limit to the number of groups.

Default

no max-groups

Parameters

value

Specifies the maximum number of groups for this interface.

Values 1 to 16000

Platforms

7705 SAR Gen 2

max-groups

Syntax

max-groups [*1..16000*]

no max-groups

Context

[Tree] (config>router>mld>if max-groups)

Full Context

configure router mld interface max-groups

Description

This command specifies the maximum number of groups for which MLD can have local receiver information based on received MLD reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. New groups are not allowed.

The **no** form of this command reverts to the default value.

Default

max-groups 0 (no limit to the number of groups)

Parameters

1..16000

Specifies the maximum number of groups for this interface.

Values 1 to 16000

Platforms

7705 SAR Gen 2

max-groups

Syntax

max-groups [*value*]

no max-groups

Context

[Tree] (config>router>pim>interface max-groups)

Full Context

configure router pim interface max-groups

Description

This command specifies the maximum number of groups for which PIM can have local receiver information based on received PIM reports on this interface. When this configuration is changed dynamically to a value

lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed. This command is applicable for IPv4 and IPv6.

The **no** form of this command sets no limit to the number of groups.

Default

no max-groups

Parameters

value

Specifies the maximum number of groups for this interface.

Values 1 to 16000

Platforms

7705 SAR Gen 2

17.51 max-grp-sources

max-grp-sources

Syntax

max-grp-sources *max-group-sources*

no max-grp-sources

Context

[Tree] (config>service>vprn>igmp>if max-grp-sources)

[Tree] (config>service>vprn>mld>interface max-grp-sources)

Full Context

configure service vprn igmp interface max-grp-sources

configure service vprn mld interface max-grp-sources

Description

This command configures the maximum number of group sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed.

The **no** form of this command reverts to the default.

Default

max-grp-sources 0

Parameters

max-grp-sources

Specifies the maximum number of group source.

Values 1 to 32000

Platforms

7705 SAR Gen 2

max-grp-sources

Syntax

max-grp-sources *value*

no max-grp-sources

Context

[Tree] (config>router>igmp>if max-grp-sources)

Full Context

configure router igmp interface max-grp-sources

Description

This command configures the maximum number of group sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed.

The **no** form of the command reverts to the default.

Default

no max-grp-sources

Parameters

value

Specifies the maximum number of group sources.

Values 1 to 32000

Platforms

7705 SAR Gen 2

max-grp-sources

Syntax

max-grp-sources [*grp-source]*

no max-grp-sources

Context

[Tree] (config>router>mld>if max-grp-sources)

Full Context

configure router mld interface max-grp-sources

Description

This command configures the maximum number of group sources for which MLD can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed.

The **no** form of this command reverts to the default.

Default

max-grp-sources 0 (no limit to the number of sources)

Parameters

grp-source

Specifies the maximum number of group sources.

Values 1 to 32000

Platforms

7705 SAR Gen 2

17.52 max-history-esp-key-records

max-history-esp-key-records

Syntax

max-history-esp-key-records *max-records*

no max-history-esp-key-records

Context

[Tree] (config>router>if>ipsec>ipsec-tunnel max-history-esp-key-records)

[Tree] (config>service>vprn>if>sap>ipsec-gw max-history-esp-key-records)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel max-history-esp-key-records)

[Tree] (config>ipsec>trans-mode-prof max-history-esp-key-records)

[Tree] (config>service>ies>if>sap>ipsec-gw max-history-esp-key-records)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel max-history-esp-key-records)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel max-history-esp-key-records)

Full Context

configure router interface ipsec ipsec-tunnel max-history-esp-key-records

configure service vprn interface sap ipsec-gw max-history-esp-key-records

configure service vprn interface ipsec ipsec-tunnel max-history-esp-key-records

configure ipsec ipsec-transport-mode-profile max-history-esp-key-records

configure service ies interface sap ipsec-gw max-history-esp-key-records

configure service ies interface ipsec ipsec-tunnel max-history-esp-key-records

configure service vprn interface sap ipsec-tunnel max-history-esp-key-records

Description

This command enables the system to keep records of CHILD-SA keys. There is a system wide limit of maximum number of IPsec tunnels that save keys. If the number of tunnel exceeds that limit, the system does not save keys for the new tunnels. Contact Nokia support for details of the limitation.

This command is ignored if the **config>ipsec>no show-ipsec-keys** command is configured.

The **no** form of this command prevents the system from keeping records.

Default

no max-history-esp-key-records

Parameters

max-records

Specifies the maximum number of recent records.

Values 1 to 48

Platforms

7705 SAR Gen 2

17.53 max-history-ike-key-records

max-history-ike-key-records

Syntax

max-history-ike-key-records *max-records*

no max-history-ike-key-records

Context

[Tree] (config>service>ies>if>sap>ipsec-gw max-history-ike-key-records)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel max-history-ike-key-records)

[Tree] (config>ipsec>trans-mode-prof max-history-ike-key-records)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel max-history-ike-key-records)

[Tree] (config>service>vprn>if>sap>ipsec-gw max-history-ike-key-records)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel max-history-ike-key-records)

Full Context

configure service ies interface sap ipsec-gw max-history-ike-key-records

configure service vprn interface sap ipsec-tunnel max-history-ike-key-records

configure ipsec ipsec-transport-mode-profile max-history-ike-key-records

configure service ies interface ipsec ipsec-tunnel max-history-ike-key-records

configure service vprn interface sap ipsec-gw max-history-ike-key-records

configure service vprn interface ipsec ipsec-tunnel max-history-ike-key-records

Description

This command enables the system to keep records of IKE-SA keys for the corresponding **ipsec-gw**, **ipsec-tunnel**, or **ipsec-transport-mode-profile**.

This command is ignored if the **config>ipsec>no show-ipsec-keys** command is enabled. There is a system-wide limit for the maximum number of IPsec tunnels that save keys. If the number of tunnels exceeds that limit, the system does not save keys for the new tunnels. Contact Nokia support for details of the limitation.

The **no** form of this command prevents the system from keeping records.

Default

no max-history-ike-key-records

Parameters

max-records

Specifies the maximum number of recent records.

Values 1 to 3

Platforms

7705 SAR Gen 2

17.54 max-lease-time

max-lease-time

Syntax

max-lease-time [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no max-lease-time

Context

[Tree] (config>router>dhcp>server>pool max-lease-time)

Full Context

configure router dhcp local-dhcp-server pool max-lease-time

Description

This command configures the maximum lease time.

The **no** form of this command reverts to the default.

Default

max-lease-time days 10

Parameters

max-lease-time

Specifies the maximum lease time.

Values		
	days	0 to 3650
	hours	0 to 23
	minutes	0 to 59
	seconds	0 to 59

Platforms

7705 SAR Gen 2

17.55 max-msg-count

max-msg-count

Syntax

max-msg-count *count*

Context

[Tree] (config>system>telemetry>notification-bundling max-msg-count)

Full Context

configure system telemetry notification-bundling max-msg-count

Description

This command sets the maximum number of notifications that can be bundled in a single telemetry message.

The **no** form of this command returns the message count to the default value.

Default

max-msg-count 100

Parameters

count

Specifies the maximum of notifications that can be bundled in a single telemetry message.

Values 2 to 1000

Platforms

7705 SAR Gen 2

17.56 max-msg-size

max-msg-size

Syntax

max-msg-size *number*

no max-msg-size

Context

[Tree] (config>system>grpc max-msg-size)

Full Context

configure system grpc max-msg-size

Description

This command configures the maximum rx message size that can be received.

The **no** form of this command reverts to the default.

Default

max-msg-size 512

Parameters

number

Specifies the message size, in MB.

Values 1 to 1024

Default 512

Platforms

7705 SAR Gen 2

17.57 max-nbr-mac-addr

max-nbr-mac-addr

Syntax

max-nbr-mac-addr *table-size*

no max-nbr-mac-addr

Context

[Tree] (config>service>template>vpls-sap-template max-nbr-mac-addr)

[Tree] (config>service>vpls>spoke-sdp max-nbr-mac-addr)

[Tree] (config>service>vpls>sap max-nbr-mac-addr)

Full Context

configure service template vpls-sap-template max-nbr-mac-addr

configure service vpls spoke-sdp max-nbr-mac-addr

configure service vpls sap max-nbr-mac-addr

Description

This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this instance.

When the configured limit is reached, no new addresses are learned from the SAP or spoke SDP until at least one FDB entry is aged out or cleared.

When the configured limit is reached and the **discard-unknown-source** command is enabled for this instance, packets with unknown source MAC addresses are discarded. If **discard-unknown-source** is disabled, the packets are forwarded if their destination MAC addresses are known, or flooded if their destination MAC addresses are unknown.

However, if the **configure service vpls discard-unknown** command is enabled, packets with unknown destination MAC addresses are discarded, even if the limit of FDB entries on the specific VPLS instance is not reached.

The **no** form of this command restores the global MAC learning limitations for this instance.

Default

no max-nbr-mac-addr

Parameters

table-size

Specifies the maximum number of learned and static entries allowed in the FDB of this service.

Values 1 to 32767

Platforms

7705 SAR Gen 2

max-nbr-mac-addr

Syntax

max-nbr-mac-addr *table-size*

no max-nbr-mac-addr

Context

[Tree] (config>service>pw-template max-nbr-mac-addr)

Full Context

configure service pw-template max-nbr-mac-addr

Description

This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this instance.

When the configured limit is reached, no new addresses are learned from the SAP or spoke SDP until at least one FDB entry is aged out or cleared.

When the configured limit is reached and the **discard-unknown-source** command is enabled for this instance, packets with unknown source MAC addresses are discarded. If **discard-unknown-source** is disabled, the packets are forwarded if their destination MAC addresses are known, or flooded if their destination MAC addresses are unknown.

However, if the **configure service vpls discard-unknown** command is enabled, packets with unknown destination MAC addresses are discarded, even if the limit of FDB entries on the specific VPLS instance is not reached.

The **no** form of this command restores the global MAC learning limitations for this instance.

Default

no max-nbr-mac-addr

Parameters

table-size

Specifies the maximum number of learned and static entries allowed in the FDB of this service.

Values 1 to 32767

Platforms

7705 SAR Gen 2

max-nbr-mac-addr

Syntax

max-nbr-mac-addr *table-size*

no max-nbr-mac-addr

Context

[Tree] (config>service>vpls>endpoint max-nbr-mac-addr)

Full Context

configure service vpls endpoint max-nbr-mac-addr

Description

This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this instance.

When the configured limit is reached, no new addresses are learned from the SAP or spoke SDP until at least one FDB entry is aged out or cleared. Packets with unknown source MAC addresses are still forwarded if their destination MAC addresses are known, or flooded if their destination MAC addresses are unknown.

The **no** form of this command restores the global MAC learning limitations for this instance.

Default

no max-nbr-mac-addr

Parameters

table-size

Specifies the maximum number of learned and static entries allowed in the FDB of this service.

Values 1 to 32767

Platforms

7705 SAR Gen 2

17.58 max-num-groups

max-num-groups

Syntax

max-num-groups *count*

no max-num-groups

Context

[Tree] (config>service>vpls>spoke-sdp>mld-snooping max-num-groups)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping max-num-groups)

[Tree] (config>service>vpls>sap>igmp-snooping max-num-groups)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping max-num-groups)

[Tree] (config>service>vpls>sap>mld-snooping max-num-groups)

[Tree] (config>service>vpls>mesh-sdp>mld-snooping max-num-groups)

Full Context

configure service vpls spoke-sdp mld-snooping max-num-groups

configure service vpls mesh-sdp igmp-snooping max-num-groups

configure service vpls sap igmp-snooping max-num-groups

configure service vpls spoke-sdp igmp-snooping max-num-groups

configure service vpls sap mld-snooping max-num-groups

configure service vpls mesh-sdp mld-snooping max-num-groups

Description

This command defines the maximum number of multicast groups that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of groups, the request is ignored.

The **no** form of this command reverts to the default value.

Default

no max-num-groups

Parameters

count

Specifies the maximum number of groups that can be joined on this SAP or SDP.

Values 1 to 1000

Platforms

7705 SAR Gen 2

max-num-groups

Syntax

max-num-groups *count*

no max-num-groups

Context

[Tree] (config>service>pw-template>igmp-snooping max-num-groups)

Full Context

configure service pw-template igmp-snooping max-num-groups

Description

This command defines the maximum number of multicast groups that can be joined. If the router receives an IGMP join message that would exceed the configured number of groups, the request is ignored.

Default

no max-num-groups

Parameters

count

Specifies the maximum number of groups that can be joined.

Values 1 to 1000

Platforms

7705 SAR Gen 2

17.59 max-num-grp-sources

max-num-grp-sources

Syntax

max-num-grp-sources [1 to 32000]

no max-num-grp-sources

Context

[Tree] (config>service>vpls>sap>igmp-snooping max-num-grp-sources)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping max-num-grp-sources)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping max-num-grp-sources)

Full Context

configure service vpls sap igmp-snooping max-num-grp-sources

configure service vpls spoke-sdp igmp-snooping max-num-grp-sources

configure service vpls mesh-sdp igmp-snooping max-num-grp-sources

Description

This command defines the maximum number of multicast SGs that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of SGs, the request is ignored.

The **no** form of this command disables the check.

Default

no max-num-grp-sources

Parameters

1 to 32000

Specifies the maximum number of multicast sources allowed to be tracked per group.

Platforms

7705 SAR Gen 2

17.60 max-num-sources

max-num-sources

Syntax

max-num-sources *max-num-sources*

no max-num-sources

Context

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping max-num-sources)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping max-num-sources)

[Tree] (config>service>vpls>sap>igmp-snooping max-num-sources)

Full Context

configure service vpls mesh-sdp igmp-snooping max-num-sources

configure service vpls spoke-sdp igmp-snooping max-num-sources

configure service vpls sap igmp-snooping max-num-sources

Description

This command defines the maximum number of multicast sources that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of sources, the request is ignored.

The **no** form of this command removes the value from the configuration.

Parameters

max-num-sources

Specifies the maximum number of multicast sources allowed per group.

Values 1 to 1000

Platforms

7705 SAR Gen 2

17.61 max-peer

max-peer

Syntax

max-peer *max-peer*

no max-peer

Context

[Tree] (config>port>ethernet>dot1x>macsec>sub-port max-peer)

Full Context

configure port ethernet dot1x macsec sub-port max-peer

Description

This command configures the max peer allowed under this MACsec instance.

Note:
The peer establishment is a race condition and first come first serve. On any security zone, only 32 peers can be supported. See SA Exhaustion Behavior for more details.

The **no** form of this command returns the value to the default.

Default

no max-peer

Parameters

max-peer

The maximum number of peers supported on this port.

Values 0 to 32

Platforms

7705 SAR Gen 2

17.62 max-percent-rate

max-percent-rate

Syntax

max-percent-rate *percentage* [**local-limit** | **reference-port-limit**]

no max-percent-rate

Context

[Tree] (config>qos>plcr-ctrl-plcy>root max-percent-rate)

Full Context

configure qos policer-control-policy root max-percent-rate

Description

This command configures the maximum percentage rate for the policer control policy.

The **no** form of this command removes the configuration.

Parameters

percentage

Specifies the percentage.

Values 0.01 to 100.00

local-limit

Keyword used to specify the local limit.

reference-port-limit

Keyword used to specify the reference port limit.

Platforms

7705 SAR Gen 2

17.63 max-rate

max-rate

Syntax

max-rate {*rate* | **max**}

no max-rate

Context

[Tree] (config>card>fp>ingress>acc>qgrp>policer-ctrl-over max-rate)

[Tree] (config>card>fp>ingress>network>qgrp>policer-ctrl-over max-rate)

Full Context

configure card fp ingress access queue-group policer-control-override max-rate

configure card fp ingress network queue-group policer-control-override max-rate

Description

This command defines the parent policer's PIR leaky bucket's decrement rate. A parent policer is created for each time the policer-control-policy is applied to either a SAP or subscriber instance. Packets that are not discarded by the child policers associated with the SAP or subscriber instance are evaluated against the parent policer's PIR leaky bucket.

For each packet, the bucket is first decremented by the correct amount based on the decrement rate to derive the current bucket depth. The current depth is then compared to one of two discard thresholds associated with the packet. The first discard threshold (discard-unfair) is applied if the FIR (Fair Information Rate) leaky bucket in the packet's child policer is in the confirming state. The second discard threshold (discard-all) is applied if the child policer's FIR leaky bucket is in the exceed state. Only one of the two thresholds is applied per packet. If the current depth of the parent policer PIR bucket is less than the threshold value, the parent PIR bucket is in the conform state for that particular packet. If the depth is equal to or greater than the applied threshold, the bucket is in the violate state for the packet.

If the result is "conform," the bucket depth is increased by the size of the packet (plus or minus the per-packet-offset setting in the child policer) and the packet is not discarded by the parent policer. If the result is "violate," the bucket depth is not increased and the packet is discarded by the parent policer. When the parent policer discards a packet, any bucket depth increases (PIR, CIR and FIR) in the parent policer caused by the packet are canceled. This prevents packets that are discarded by the parent policer from consuming the child policers PIR, CIR and FIR bandwidth.

The **policer-control-policy root max-rate** setting may be overridden on each SAP or sub-profile where the policy is applied.

The **no** form of this command returns the policer-control-policy's parent policer maximum rate to **max**.

Default

max-rate max

Parameters

rate

Specifies that a kilobits-per-second value is mutually exclusive with the **max** keyword. The kilobits-per-second value must be defined as an integer that represents the number of kilobytes that the parent policer will be decremented per second. The actual decrement is performed per packet based on the time that has elapsed since the last packet associated with the parent policer.

Values 0 to 2000000000

max

The **max** keyword is mutually exclusive with defining a kilobits-per-second value. When max is specified, the parent policer does not enforce a maximum rate on the aggregate throughput of the child policers. This is the default setting when the **policer-control-policy** is first created and is the value that the parent policer returns to when no max-rate is executed. In order for the parent policer to be effective, a kilobits-per-second value should be specified.

Platforms

7705 SAR Gen 2

max-rate

Syntax

max-rate *pir-rate*

max-rate percent *percent-rate*

no max-rate

Context

[Tree] (config>port>ethernet>egr-scheduler-override max-rate)

Full Context

configure port ethernet egress-scheduler-override max-rate

Description

This command overrides the **max-rate** parameter found in the port-scheduler-policy associated with the port. When a max-rate is defined at the port or channel level, the port scheduler policies max-rate parameter is ignored.

The egress-scheduler-override **max-rate** command supports a parameter that allows the override command to restore the default of not having a rate limit on the port scheduler. This is helpful when the port scheduler policy has an explicit maximum rate defined and it is desirable to remove this limit at the port instance.

The **no** form of this command removes the maximum rate override from the egress port or channels port scheduler context. Once removed, the max-rate parameter from the port scheduler policy associated with the port or channel will be used by the local scheduler context.

Parameters

pir-rate

Specifies the explicit maximum frame based bandwidth limit, in kilobits per second. This value overrides the QoS scheduler policy rate.

Values **For Ethernet:** 1 to 6400000000, **max**

For SONET-SDH and TDM: 1 to 3200000000, **max**

percent-rate

Specifies the percent rate.

Values 0.01 to 100.00

Platforms

7705 SAR Gen 2

max-rate

Syntax

max-rate {*rate* | **max**}

Context

[Tree] (config>service>epipe>sap>ingress>policer-control-override max-rate)

[Tree] (config>service>epipe>sap>egress>policer-control-override max-rate)

Full Context

configure service epipe sap ingress policer-control-override max-rate

configure service epipe sap egress policer-control-override max-rate

Description

This command, within the SAP ingress and egress contexts, overrides the root arbiter parent policer max-rate that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy max-rate parameter have no effect on the SAP's parent policer until the override is removed using the **no max-rate** command within the SAP.

The **no** form of this command returns the policer-control-policy's parent policer maximum rate to **max**.

Parameters

rate

Specifies the rate override in kilobits per second.

Values 1 to 6400000000

max

Specifies the maximum rate override.

Platforms

7705 SAR Gen 2

max-rate

Syntax

max-rate {*rate* | **max**}

Context

[Tree] (config>service>vpls>sap>ingress>policer-ctrl-over max-rate)

[Tree] (config>service>vpls>sap>egress>policer-ctrl-over max-rate)

Full Context

configure service vpls sap ingress policer-control-override max-rate

configure service vpls sap egress policer-control-override max-rate

Description

This command, within the SAP ingress and egress contexts, overrides the root arbiter parent policer max-rate that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy max-rate parameter have no effect on the SAP's parent policer until the override is removed using the **no max-rate** command within the SAP.

The **no** form of this command removes an explicit rate value from the aggregate rate therefore returning it to its default value.

Parameters

***rate* | max**

Specifies the max rate override in kilobits per second or use the maximum

Values 1 to 6400000000, **max**

Platforms

7705 SAR Gen 2

max-rate

Syntax

max-rate {*rate* | **max**}

Context

[Tree] (config>service>ies>if>sap>ingress>policer-ctrl-over max-rate)

[Tree] (config>service>ies>if>sap>egress>policer-ctrl-over max-rate)

Full Context

configure service ies interface sap ingress policer-control-override max-rate

configure service ies interface sap egress policer-control-override max-rate

Description

This command, within the SAP ingress and egress contexts, overrides the root arbiter parent policer max-rate that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy max-rate parameter have no effect on the SAP's parent policer until the override is removed using the **no max-rate** command within the SAP.

Parameters

***rate* | max**

Specifies the rate override in kilobits per second or use the maximum override value.

Values 1 to 6400000000, max

Platforms

7705 SAR Gen 2

max-rate

Syntax

max-rate {*rate* | **max**}

Context

[Tree] (config>service>vprn>if>sap>egress>policer-ctrl-over max-rate)

[Tree] (config>service>vprn>if>sap>ingress>policer-ctrl-over max-rate)

Full Context

configure service vprn interface sap egress policer-control-override max-rate

configure service vprn interface sap ingress policer-control-override max-rate

Description

This command, within the SAP ingress and egress contexts, overrides the root arbiter parent policer max-rate that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy max-rate parameter have no effect on the SAP's parent policer until the override is removed using the **no max-rate** command within the SAP.

The **no** form of this command returns the policer-control-policy's parent policer maximum rate to **max**.

Parameters

***rate* | max**

Specifies the rate override in kilobits per second or use the maximum override value.

Values 1 to 6400000000, max

Platforms

7705 SAR Gen 2

max-rate

Syntax

max-rate *rate*

no max-rate

Context

[Tree] (config>qos>plcr-ctrl-plcy>root max-rate)

Full Context

configure qos policer-control-policy root max-rate

Description

The **max-rate** command defines the parent policer's PIR leaky bucket's decrement rate. A parent policer is created for each time the policer-control-policy is applied to either a SAP or subscriber instance. Packets that are not discarded by the child policers associated with the SAP or subscriber or multiservice site instances are evaluated against the parent policer's PIR leaky bucket.

For each packet, the bucket is first decremented by the correct amount based on the decrement rate to derive the current bucket depth. The current depth is then compared to one of two discard thresholds associated with the packet. The first discard threshold (discard-unfair) is applied if the FIR (Fair Information Rate) leaky bucket in the packet's child policer is in the confirming state. The second discard threshold (discard-all) is applied if the child policer's FIR leaky bucket is in the exceed state. Only one of the two thresholds is applied per packet. If the current depth of the parent policer PIR bucket is less than the threshold value, the parent PIR bucket is in the conform state for that particular packet. If the depth is equal to or greater than the applied threshold, the bucket is in the violate state for the packet.

If the result is "conform," the bucket depth is increased by the size of the packet (plus or minus the per-packet-offset setting in the child policer) and the packet is not discarded by the parent policer. If the result is "violate," the bucket depth is not increased and the packet is discarded by the parent policer. When the parent policer discards a packet, any bucket depth increases (PIR, CIR and FIR) in the parent policer caused by the packet are canceled. This prevents packets that are discarded by the parent policer from consuming the child policers PIR, CIR, and FIR bandwidth.

The **policer-control-policy root max-rate** setting may be overridden on each SAP or sub-profile where the policy is applied.

The **no** form of this command returns the policer-control-policy's parent policer maximum rate to **max**.

Parameters

rate

The kilobits-per-second value must be defined as an integer that represents the number of kilobytes that the parent policer will be decremented per second. The actual decrement is performed per packet, based on the time that has elapsed since the last packet associated with the parent policer.

Values 1 to 6400000000, **max**

max

When **max** is specified, the parent policer does not enforce a maximum rate on the aggregate throughput of the child policers. This is the default setting when the **policer-control-policy** is first created and is the value that the parent policer returns to when no max-rate is executed. In order for the parent policer to be effective, a kilobits-per-second value should be specified.

Platforms

7705 SAR Gen 2

max-rate

Syntax

max-rate *pir-rate*

max-rate percent *percent-rate*

no max-rate

Context

[Tree] (config>qos>port-scheduler-policy max-rate)

Full Context

configure qos port-scheduler-policy max-rate

Description

This command defines an explicit maximum frame-based bandwidth limit for the port scheduler policies scheduler context. By default, when a scheduler policy is associated with a port or channel, the instance of the scheduler on the port automatically limits the bandwidth to the lesser of port or channel line rate and a possible egress-rate value (for Ethernet ports). If a max-rate is defined that is smaller than the port or channel rate, the expressed kilobits per second value is used instead. The max-rate command is another way to sub-rate the port or channel.

The **max-rate** command may be executed at any time for an existing port-scheduler-policy. When a new max-rate is given for a policy, the system evaluates all instances of the policy to see if the configured rate is smaller than the available port or channel bandwidth. If the rate is smaller and the maximum rate is not currently overridden on the scheduler instance, the scheduler instance is updated with the new maximum rate value.

The max-rate value defined in the policy may be overridden on each scheduler instance. If the maximum rate is explicitly defined as an override on a port or channel, the policies max-rate value has no effect.

The **no** form of this command removes an explicit rate value from the port scheduler policy. When removed, all instances of the scheduler policy on egress ports or channel are allowed to run at the available line rate unless the instance has a max-rate override in place.

Parameters

pir-rate

Specifies the PIR rate, in kilobits per second.

Values 1 to 6400000000, **max**

percent *percent-rate*

Specifies the percent rate.

Values 0.01 to 100.00

Platforms

7705 SAR Gen 2

17.64 max-sess-prefix

max-sess-prefix

Syntax

max-sess-prefix *count*

no max-sess-prefix

Context

[Tree] (config>test-oam>twamp>server>prefix max-sess-prefix)

Full Context

configure test-oam twamp server prefix max-sess-prefix

Description

This command configures the maximum number of concurrent TWAMP-Test sessions by clients with an IP address in a specific prefix. A new test session (described by a Request-TW-Session message) is rejected if accepting it would cause either the limit defined by this command or the server limit (max-sess-server) to be exceeded.

The **no** form of this command returns the value to the default.

Default

max-sess-prefix 32

Parameters

count

Specifies the maximum number of concurrent test sessions.

Values 0 to 128

Default 32

Platforms

7705 SAR Gen 2

17.65 max-sess-server

max-sess-server

Syntax

max-sess-server *count*

no max-sess-server

Context

[Tree] (config>test-oam>twamp>server max-sess-server)

Full Context

configure test-oam twamp server max-sess-server

Description

This command configures the maximum number of concurrent TWAMP-Test sessions across all allowed clients. A new test session (described by a Request-TW-Session message) is rejected if accepting it would cause either the limit defined by this command or a prefix limit (max-sess-prefix) to be exceeded.

The **no** form of this command returns the value to the default.

Default

max-sess-server 32

Parameters

count

Specifies the maximum number of concurrent test sessions.

Values 0 to 128

Default 32

Platforms

7705 SAR Gen 2

17.66 max-sessions

max-sessions

Syntax

max-sessions *number*

Context

[Tree] (config>service>vprn>bgp>group>dynamic-neighbor>interface max-sessions)

[Tree] (config>router>bgp>group>dynamic-neighbor>interface max-sessions)

Full Context

configure service vprn bgp group dynamic-neighbor interface max-sessions

configure router bgp group dynamic-neighbor interface max-sessions

Description

This command configures the maximum number of dynamic sessions that are allowed to be set up on the interface as a result of accepting sessions from link-local addresses or initiating sessions by receiving IPv6 router advertisements.

Default

max-sessions 1

Parameters

number

Specifies the maximum number of sessions.

Values 1 to 255

Platforms

7705 SAR Gen 2

17.67 max-sources

max-sources

Syntax

max-sources *max-sources*

no max-sources

Context

[Tree] (config>service>vprn>mld>interface max-sources)

[Tree] (config>service>vprn>igmp>if max-sources)

Full Context

configure service vprn mld interface max-sources

configure service vprn igmp interface max-sources

Description

This command specifies the maximum number of sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of sources, the sources that are already accepted are not deleted. Only new sources will not be allowed.

Parameters

sources

Specifies the maximum number of sources for this interface.

Values 1 to 1000

Platforms

7705 SAR Gen 2

max-sources

Syntax

max-sources *value*

no max-sources

Context

[Tree] (config>router>igmp>if max-sources)

Full Context

configure router igmp interface max-sources

Description

This command configures the maximum number of group sources for this group-interface.

Parameters

value

Specifies the maximum number of group sources.

Values 1 to 1000

Platforms

7705 SAR Gen 2

max-sources

Syntax

max-sources [*grp-source*]

no max-sources

Context

[Tree] (config>router>mld>if max-sources)

Full Context

configure router mld interface max-sources

Description

This command configures the maximum number of group sources for this interface.

The **no** form of this command reverts to the default.

Default

no max-sources

Parameters

grp-source

Specifies the maximum number of group sources for this interface.

Platforms

7705 SAR Gen 2

17.68 max-sr-labels

max-sr-labels

Syntax

max-sr-labels *label-stack-size* [**additional-frr-labels** *labels*]

no max-sr-labels

Context

[Tree] (config>router>mpls>lsp-template max-sr-labels)

[Tree] (config>router>mpls>lsp max-sr-labels)

Full Context

configure router mpls lsp-template max-sr-labels

configure router mpls lsp max-sr-labels

Description

This command configures the maximum number of labels which the ingress LER can push for a given SR-TE LSP.

This command is used to allow room to insert additional transport, service, and other labels when packets are forwarded in a given context.

The **max-sr-labels** *label-stack-size* value should reflect the desired maximum label stack of the primary path of the SR-TE LSP.

The value in **additional-frr-labels** *labels* should reflect additional labels inserted by remote LFA for the backup next-hop of the SR-TE LSP.

The sum of both label values represents the worst case transport of SR label stack size for this SR-TE LSP and is populated by MPLS in the TTM such that services and shortcut applications can check it to decide if a service can be bound or a route can be resolved to this SR-TE LSP.

The maximum label stack supported by the router is always signaled by PCC in the PCEP Open object as part of the as SR-PCE-CAPABILITY TLV. It is referred to as the Maximum Stack Depth (MSD).

In addition, the per-LSP value for the max-sr-labels option, if configured, is signaled by PCC to PCE in the Segment-ID (SID) Depth value in a METRIC object for both a PCE computed LSP and a PCE controlled LSP. PCE will compute and provide the full explicit path with TE-links specified. If there is no path with the number of hops lower than the MSD value, or the Segment-ID (SID) Depth value if signaled, a reply with no path will be returned to PCC.

For a PCC controlled LSP, if the label stack returned by the TE-DB's hop-to-label translation exceeds the per LSP maximum SR label stack size, the LSP is brought down.

The **no** form of this command reverts to the default value.

Default

max-sr-labels 6 additional-frr-labels 1

Parameters

label-stack-size

Specifies the label stack size.

Values 1 to 11

additional-frr-labels ***labels***

Specifies the addition FRR labels.

Values 0 to 3

Platforms

7705 SAR Gen 2

17.69 max-srte-pce-init-lsps

max-srte-pce-init-lsps

Syntax

max-srte-pce-init-lsps *max-number*

no max-srte-pce-init-lsps

Context

[Tree] (config>router>pcep>pcc max-srte-pce-init-lsps)

Full Context

configure router pcep pcc max-srte-pce-init-lsps

Description

This command configures the maximum number of PCE-initiated SR-TE LSPs that can be created by the router.

The **no** form of the command sets this value to the default.

Default

max-srte-pce-init-lsps 8191

Parameters

max-number

Specifies the maximum number of SR-TE PCE-initiated LSPs.

Values 0 to 8191

Platforms

7705 SAR Gen 2

17.70 max-suppress

max-suppress

Syntax

max-suppress *minutes*

no max-suppress

Context

[Tree] (config>router>policy-options>damping max-suppress)

Full Context

configure router policy-options damping max-suppress

Description

This command configures the maximum suppression parameter for the route damping profile.

This value indicates the maximum time, expressed in minutes, that a route can remain suppressed.

The **no** form of this command removes the maximum suppression parameter from the damping profile.

Default

no max-suppress

Parameters

minutes

Specifies the maximum suppression time, in minutes, expressed as a decimal integer.

Values 1 to 720

Platforms

7705 SAR Gen 2

17.71 max-time-granularity

max-time-granularity

Syntax

[no] max-time-granularity *time*

Context

[Tree] (config>system>telemetry>notification-bundling max-time-granularity)

Full Context

configure system telemetry notification-bundling max-time-granularity

Description

This command sets the maximum time interval during which telemetry notifications are bundled. All bundled notifications will have the same timestamp, which is the timestamp of the bundle.

The **no** form of this command returns the time granularity to the default value.

Default

max-time-granularity 100

Parameters

time

Specifies the maximum time interval during which telemetry notifications are bundled, in milliseconds.

Values 1 to 1000

Platforms

7705 SAR Gen 2

17.72 max-ve-id

max-ve-id

Syntax

max-ve-id *value*

no max-ve-id

Context

[Tree] (config>service>vpls>bgp-vpls max-ve-id)

Full Context

configure service vpls bgp-vpls max-ve-id

Description

This command configures the allowed range for the VE-id value: locally configured and received in a NLRI. Configuration of a VE-id higher than the value specified in this command is not allowed.

Also upon reception of a higher VE-id in an NLRI imported in this VPLS instance (RT is the configured import RT) the following action must be taken:

- a trap must be generated informing the operator of the mismatch.
- NLRI must be dropped
- no service labels are to be installed for this VE-id
- no new NLRI must be generated if a new offset is required for VE-id.

The **no** form of this command sets the max-ve-id to un-configured. The BGP VPLS status should be administratively down for "no max-ve-id" to be used.

The max-ve-id value can be changed without shutting down bgp-vpls if the newly provisioned value does not conflict with the already configured local VE-ID. If the value of the local-VE-ID is higher than the new

max-ve-id value the command is rejected. The operator needs to decrease first the VE-ID before running the command.

The actions taken for other max-ve-id values are as follows:

- max-ve-id value higher than all VE-IDs (local and received) is allowed and there are no effects.
- max-ve-id higher than the local VE-ID but smaller than the remote VE-IDs:
  - Provisioning is allowed
  - A warning message will be generated stating that "Higher VE-ID values were received in the BGP VPLS context. Related pseudowires will be removed."
  - The pseudowires associated with the higher VE-IDs will be removed locally.
  - This is a situation that should be corrected by the operator as the pseudowire may be down just at the local PE, consuming unnecessarily core bandwidth. The higher VE-IDs should be removed or lowered.

If the max-ve-id has increased a BGP route refresh is sent to the VPLS community to get the routes which might have been rejected earlier due to max-ve-id check. A max-ve-id value needs to be provisioned for BGP VPLS to be in "no shutdown" state.

Default

no max-ve-id

Parameters

value

Specifies the allowed range of [1-value] for the VE-id. The configured value must be bigger than the existing VE-ids

Values 1 to 65535

Platforms

7705 SAR Gen 2

17.73 max-wait-to-advertise

max-wait-to-advertise

Syntax

max-wait-to-advertise *seconds*

no max-wait-to-advertise

Context

[Tree] (config>service>vprn>bgp>convergence>family max-wait-to-advertise)

Full Context

configure service vprn bgp convergence family max-wait-to-advertise

Description

This command configures the maximum amount of time that BGP waits until it starts advertising IPv4-unicast or IPv6-unicast routes to its BGP peers. For IPv4-unicast routes, *seconds* is measured from the time when the first peer that supports the IPv4-unicast address family comes up. For IPv6-unicast routes *seconds* is measured from the time when the first peer that negotiates the IPv6-unicast address family comes up.

The time limit configured by this command should allow sufficient time for all important peers to re-establish their sessions with the restarting router and advertise their complete set of IPv4-unicast or IPv6-unicast routes (followed by the applicable End of RIB marker).

The **no** form of this command implements the default value, which is three times the value of the **min-wait-to-advertise** time limit.

Default

no max-wait-to-advertise

Parameters

seconds

Specifies the maximum amount of time, in seconds, that BGP waits until IPv4-unicast or IPv6-unicast routes are advertised to peers.

Values 0 to 3600

Platforms

7705 SAR Gen 2

max-wait-to-advertise

Syntax

max-wait-to-advertise *seconds*

no max-wait-to-advertise

Context

[Tree] (config>router>bgp>convergence>family max-wait-to-advertise)

Full Context

configure router bgp convergence family max-wait-to-advertise

Description

This command configures the maximum amount of time that BGP waits until it starts advertising IPv4-unicast or IPv6-unicast routes to its BGP peers. For IPv4-unicast routes, the time limit value is measured from the time when the first peer that supports the IPv4-unicast address family comes up. For IPv6-unicast

routes the time limit value is measured from the time when the first peer that negotiates the IPv6-unicast address family comes up.

The time limit configured by this command should allow sufficient time for all important peers to re-establish their sessions with the restarting router and advertise their complete set of IPv4-unicast or IPv6-unicast routes (followed by the applicable End of RIB marker).

The **no** form of this command implements the default value, which is three times the value of the **min-wait-to-advertise** time-limit.

Default

no max-wait-to-advertise

Parameters

seconds

Specifies the maximum amount of time, in seconds, that BGP waits until IPv4-unicast or IPv6-unicast routes are advertised to peers.

Values 0 to 3600

Platforms

7705 SAR Gen 2

17.74 maximum-cert-chain-depth

maximum-cert-chain-depth

Syntax

maximum-cert-chain-depth *level*

no maximum-cert-chain-depth

Context

[Tree] (config>system>security>pki maximum-cert-chain-depth)

Full Context

configure system security pki maximum-cert-chain-depth

Description

This command defines the maximum depth of certificate chain verification. This number is applied system wide.

The **no** form of this command reverts to the default.

Default

maximum-cert-chain-depth 7

Parameters

level

Specifies the maximum depth level of certificate chain verification, range from 1 to 7. the certificate under verification is not counted in. for example, if this parameter is set to 1, then the certificate under verification must be directly signed by trust anchor CA.

Values 1 to 7

Platforms

7705 SAR Gen 2

17.75 maximum-client-lead-time

maximum-client-lead-time

Syntax

maximum-client-lead-time [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no maximum-client-lead-time

Context

[Tree] (config>router>dhcp>server>pool>failover maximum-client-lead-time)

[Tree] (config>router>dhcp6>server>failover maximum-client-lead-time)

[Tree] (config>router>dhcp6>server>pool>failover maximum-client-lead-time)

[Tree] (config>router>dhcp>server>failover maximum-client-lead-time)

[Tree] (config>service>vprn>dhcp>server>failover maximum-client-lead-time)

[Tree] (config>service>vprn>dhcp>server>pool>failover maximum-client-lead-time)

[Tree] (config>service>vprn>dhcp6>server>pool>failover maximum-client-lead-time)

[Tree] (config>service>vprn>dhcp6>server>failover maximum-client-lead-time)

Full Context

configure router dhcp local-dhcp-server pool failover maximum-client-lead-time

configure router dhcp6 local-dhcp-server failover maximum-client-lead-time

configure router dhcp6 local-dhcp-server pool failover maximum-client-lead-time

configure router dhcp local-dhcp-server failover maximum-client-lead-time

configure service vprn dhcp local-dhcp-server failover maximum-client-lead-time

configure service vprn dhcp local-dhcp-server pool failover maximum-client-lead-time

configure service vprn dhcp6 local-dhcp-server pool failover maximum-client-lead-time

configure service vprn dhcp6 local-dhcp-server failover maximum-client-lead-time

Description

The command configures the maximum time that a DHCP server can extend client's lease time beyond the lease time currently known by the DHCP partner node. In dual-homed environment, the initial lease time for all DHCP clients is by default restricted to MCLT. Consecutive DHCP renews can extend the lease time beyond the MCLT.

The maximum client lead time (MCLT) is a safeguard against IP address/prefix duplication in cases of a lease synchronization failure when local-remote failover model is deployed.

Once the intercommunication link failure between the redundant DHCP servers is detected, the DHCP IP address range configured as remote will not be allowed to start delegating new leases until the MCLT + partner-down-delay intervals expire. This is to ensure that the new lease that was delegated from the local IP address-range/prefix on one node but was never synchronized due to the intercommunication link failure, will expire before the same IP address/prefix is allocated from the remote IP address-range/prefix on the other node.

However, the already existing (and synchronized) lease times can be renewed from the remote IP address range at any time, regardless of the state of the intercommunication link (operational or failed).

Lease synchronization failure can be caused either by a node failure, or a failure of the link over which the DHCP leases are synchronized (intercommunication link). Synchronization failure detection can take up to 3 seconds.

During the failure, the DHCP lease time for the new clients is restricted to MCLT while for the existing clients the lease time will over time (by consecutive DHCP renews) be gradually reduced to the MCLT.

The **no** form of this command reverts to the default.

Default

maximum-client-lead-time min 10

Parameters

maximum-client-lead-time

Specifies the maximum client lead time.

Values	**hrs** *hours*	1 to 23
	min *minutes*	1 to 59
	sec *seconds*	1 to 59

Platforms

7705 SAR Gen 2

17.76 maximum-declined

maximum-declined

Syntax

maximum-declined *maximum-declined*

no maximum-declined

Context

[Tree] (config>router>dhcp>server>pool>subnet maximum-declined)

[Tree] (config>service>vprn>dhcp>server>pool>subnet maximum-declined)

Full Context

configure router dhcp local-dhcp-server pool subnet maximum-declined

configure service vprn dhcp local-dhcp-server pool subnet maximum-declined

Description

This command configures the maximum number of declined addresses allowed.

The **no** form of the reverts to the default.

Default

maximum-declined 64

Parameters

maximum-declined

Specifies the maximum number of declined addresses allowed.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

17.77 maximum-ipv6-routes

maximum-ipv6-routes

Syntax

maximum-ipv6-routes *number* [**log-only**] [**threshold** *percentage*]

no maximum-ipv6-routes

Context

[Tree] (config>service>vprn maximum-ipv6-routes)

Full Context

configure service vprn maximum-ipv6-routes

Description

This command specifies the maximum number of remote IPv6 routes that can be held within a VPN routing/ forwarding (VRF) context. The **local**, **host**, **static** and **aggregate** routes are not counted.

The VPRN service ID must be in a shutdown state in order to modify maximum-routes command parameters.

If the **log-only** parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then the offending RIP peer (if applicable) is brought down (but the VPRN instance remains up). BGP peering will remain up but the exceeding BGP routes will not be added to the VRF.

The maximum route threshold can dynamically change to increase the number of supported routes even when the maximum has already been reached. Protocols will resubmit their routes which were initially rejected.

The **no** form of this command disables any limit on the number of routes within a VRF context. The threshold will not be raised. Issue the **no** form of this command only when the VPRN instance is shutdown.

Default

0 or disabled

Parameters

number

Specifies an integer that specifies the maximum number of routes to be held in a VRF context.

Values 1 to 2147483647

log-only

Specifies that if the maximum limit is reached, only log the event. **log-only** does not disable the learning of new routes.

threshold percentage

Specifies the percentage at which a warning log message and SNMP trap should be set. There are two warnings, the first is a mid-level warning at the threshold value set and the second is a high-level warning at level between the maximum number of routes and the mid-level rate ([mid+max] / 2).

Values 0 to 100

Platforms

7705 SAR Gen 2

17.78 maximum-number-upas

maximum-number-upas

Syntax

maximum-number-upas *maximum-number-upas*

no maximum-number-upas

Context

[Tree] (config>router>isis>upa maximum-number-upas)

Full Context

configure router isis prefix-unreachable maximum-number-upas

Description

This command configures a limit for the number of UPAs the router can advertise. If overrun, a system log is generated and additional UPAs are not advertised.

The **no** form of this command reverts to the default.

Default

32

Parameters

maximum-number-upas

Specifies the maximum number of UPAs that the router can advertise.

Values 1 to 8192

Platforms

7705 SAR Gen 2

17.79 maximum-original-datagram

maximum-original-datagram

Syntax

[**no**] **maximum-original-datagram**

Context

[Tree] (config>test-oam>icmp>ipv6 maximum-original-datagram)

Full Context

configure test-oam icmp ipv6 maximum-original-datagram

Description

This command enables the original datagram field of the ICMPv6 error message to be a maximum of 1232 bytes.

The **no** form of this command may result in an original datagram field of the ICMPv6 error message smaller than 1232 bytes be built smaller.

Default

no maximum-original-datagram

Platforms

7705 SAR Gen 2

17.80 maximum-paths

maximum-paths

Syntax

maximum-paths *max-paths* [**ebgp** *ebgp-max-paths*] [**ibgp** *ibgp-max-paths*] [**restrict** {**same-neighbor-as** | **exact-as-path**}] [**unequal-cost**]

no maximum-paths

Context

[Tree] (config>service>vprn>bgp>multi-path maximum-paths)

Full Context

configure service vprn bgp multi-path maximum-paths

Description

This command sets ECMP multi-path parameters that apply to all address families for that BGP multi-path. For some address families it is possible to override these settings on a per address family basis.

When multi-path is enabled, traffic to the destination is load-shared across a set of paths (BGP routes) that the BGP decision process considers equal to the best path. The actual distribution of traffic over the multiple paths may be equal or unequal (that is, based on weights derived from the Link Bandwidth Extended Community).

To qualify as a multi-path, a non-best route must meet the following criteria (some criteria are controlled by this command):

- The multi-path route must be the same type of route as the best path (same AFI/SAFI and, in some cases, same next-hop resolution method).
- The multi-path route must be tied with the best path for all criteria of greater significance than next-hop cost, except for criteria that are configured to be ignored.
- If the best path selection reaches the next-hop cost comparison, the multi-path route must have the same next-hop cost as the best route unless the **unequal-cost** option is configured.
- The multi-path route must not have the same BGP next-hop as the best path or any other multi-path route.
- The multi-path route must not cause the ECMP limit of the routing instance to be exceeded (configured using the **ecmp** command with a value in the range 1 to 64).
- The multi-path route must not cause the applicable *max-paths* limit to be exceeded. If the best path is an EBGP learned route and the **ebgp** option is used, the *ebgp-max-paths* limit overrides the *max-paths* limit. If the best path is an IBGP-learned route and the **ibgp** option is used, the *ibgp-max-paths* limit overrides the *max-paths* limit. All path limits are configurable up to a maximum of 64. Multi-path is effectively disabled if a value is set to 1.
- The multi-path route must have the same neighbor AS in its AS path as the best path if the **restrict same-neighbor-as** option is configured. By default, any path with the same AS path length as the best path (regardless of neighbor AS) is eligible for multi-path.
- The route must have the same AS path as the best path if the **restrict exact-as-path** option is configured. By default, any path with the same AS path length as the best path (regardless of the actual AS numbers) is eligible for multi-path.

The **no** form of this command disables BGP multi-path.

Default

no maximum-paths

Parameters

max-paths

Specifies the maximum number of multipaths per prefix/NLRI if *ebgp-max-paths* or *ibgp-max-paths* does not apply.

Values 1 to 64

egp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an EBGP learned route.

Values 1 to 64

ibgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an IBGP learned route.

Values 1 to 64

restrict same-neighbor-as

Specifies that the non-best path must have the same neighbor AS in its AS path as the best path.

restrict exact-as-path-as

Specifies that the non-best path must have the same AS path as the best path.

unequal-cost

Instructs BGP to ignore differences in the next-hop cost only when determining eligible multipaths.

Platforms

7705 SAR Gen 2

maximum-paths

Syntax

maximum-paths *max-paths* [**ebgp** *ebgp-max-paths*] [**ibgp** *ibgp-max-paths*] [**restrict** {**same-neighbor-as** | **exact-as-path**}] [**unequal-cost**]

no maximum-paths

Context

[Tree] (config>router>bgp>multi-path maximum-paths)

Full Context

configure router bgp multi-path maximum-paths

Description

This command sets ECMP multipath parameters that apply to all address families for that BGP multipath. For some address families it is possible to override these settings on a per address family basis.

When multipath is enabled, traffic to the destination is load-shared across a set of paths (BGP routes) that the BGP decision process considers equal to the best path. The actual distribution of traffic over the multiple paths may be equal or unequal (that is, based on weights derived from the Link Bandwidth Extended Community).

The **no** form of this command disables BGP multipath.

Default

no maximum-paths

Parameters

max-paths

Specifies the maximum number of multipaths per prefix/NLRI if *ebgp-max-paths* or *ibgp-max-paths* does not apply.

Values 1 to 64

ebgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an EBGP learned route.

Values 1 to 64

ibgp-max-paths

Specifies the maximum number of multipaths per prefix or NLRI when the best path is an IBGP learned route.

Values 1 to 64

restrict same-neighbor-as

Specifies that the non-best path must have the same neighbor AS in its AS path as the best path.

restrict exact-as-path

Specifies that the non-best path must have the same AS path as the best path.

unequal-cost

Instructs BGP to ignore differences in the next-hop cost only when determining eligible multipaths.

Platforms

7705 SAR Gen 2

17.81 maximum-recovery-time

maximum-recovery-time

Syntax

maximum-recovery-time *interval*

no maximum-recovery-time

Context

[Tree] (config>router>ldp>graceful-restart maximum-recovery-time)

Full Context

configure router ldp graceful-restart maximum-recovery-time

Description

This command configures the local maximum recovery time.

The **no** form of this command returns the default value.

Default

no maximum-recovery-time (which equals a value of 120 seconds)

Parameters

interval

Specifies the length of time in seconds.

Values 15 to 1800

Platforms

7705 SAR Gen 2

17.82 maximum-routes

maximum-routes

Syntax

maximum-routes *number* [**log-only**] [**threshold** *percentage*]

no maximum-routes

Context

[Tree] (config>service>vprn maximum-routes)

Full Context

configure service vprn maximum-routes

Description

This command specifies the maximum number of remote routes that can be held within a VPN routing/forwarding (VRF) context. The **local**, **host**, **static** and **aggregate** routes are not counted.

The VPRN service ID must be in a shutdown state in order to modify maximum-routes command parameters.

If the **log-only** parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then the offending RIP peer (if applicable) is brought down (but the VPRN instance remains up). BGP peering will remain up but the exceeding BGP routes will not be added to the VRF.

The maximum route threshold can dynamically change to increase the number of supported routes even when the maximum has already been reached. Protocols will resubmit their routes which were initially rejected.

The **no** form of this command disables any limit on the number of routes within a VRF context. Issue the **no** form of this command only when the VPRN instance is shutdown.

Default

0 or disabled — The threshold will not be raised.

Parameters

number

An integer that specifies the maximum number of routes to be held in a VRF context.

Values 1 to 2147483647

log-only

Specifies that if the maximum limit is reached, only log the event. **log-only** does not disable the learning of new routes.

threshold percentage

The percentage at which a warning log message and SNMP trap should be set. There are two warnings, the first is a mid-level warning at the threshold value set and the second is a high-level warning at level between the maximum number of routes and the mid-level rate ([mid+max] / 2).

Values 0 to 100

Platforms

7705 SAR Gen 2

17.83 maximum-sid-depth

maximum-sid-depth

Syntax

maximum-sid-depth

Context

[Tree] (config>router>isis>segm-rtng maximum-sid-depth)

Full Context

configure router isis segment-routing maximum-sid-depth

Description

Commands in this context configure a manual override of the Maximum Segment Depths (MSD) that is announced by the router.

Platforms

7705 SAR Gen 2

17.84 mbb

mbb

Syntax

mbb [**detail**]

no mbb

Context

[Tree] (debug>router>mpls>event mbb)

Full Context

debug router mpls event mbb

Description

This command debugs the state of the most recent invocation of the make-before-break (MBB) functionality.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about MBB events.

Platforms

7705 SAR Gen 2

17.85 mbb-prefer-current-hops

mbb-prefer-current-hops

Syntax

[**no**] **mbb-prefer-current-hops**

Context

[Tree] (config>router>mpls mbb-prefer-current-hops)

Full Context

configure router mpls mbb-prefer-current-hops

Description

This command implements a new option in the CSPF path computation during a Make-Before-Break (MBB) procedure of an RSVP LSP.

When MPLS performs an MBB for the primary or secondary path of a P2P LSP, or the S2L path of a P2MP LSP, and the new **mbb-prefer-current-hops** option is enabled in MPLS context, CSPF will select a path, among equal-cost candidate paths, with the most overlapping links with the current path. Normally, CSPF selects the path randomly.

The procedures of the new MBB CSPF path selection apply to LSP without the least-fill option enabled. If the least-fill rule results in a different path, the LSP path will be moved though. Users can still favor stability over least-fill condition by applying a larger value to the parameter **least-fill-min-thd** under the MPLS context such that a path will only be moved when the difference of the least-available bandwidth becomes significant enough between the most used links in the equal cost paths. If that difference is not significant enough, CSPF will select the path with the most overlapping links instead of selecting a path randomly.

The procedures when the new **mbb-prefer-current-hops** option is enabled apply to all MBB types. Thus, it applies to the auto-bandwidth MBB, the configuration change MBB, the soft preemption MBB, the TE graceful shutdown MBB, the delayed retry MBB (for SRLG secondary LSP path), the path change MBB, the timer resignal MBB, and the manual resignal MBB.

During the FRR global revertive MBB, CSPF selects a random link among the ones available between the PLR node and the Merge Point node, including the failed link if it has restored in the meantime. These links cannot be checked for overlap with the current path.

The TE graceful shutdown MBB will still avoid the link or node that is in maintenance and the soft preemption MBB will still avoid the link that is overbooked.

For an inter-area LSP, this feature applies to the subset of the path from the ingress LER to the exit ABR.

The procedures of this feature are not applied to a zero bandwidth CSPP LSP, including an auto-bandwidth CSPF LSP while its operational bandwidth is zero, and to a non-CSPF LSP.

Platforms

7705 SAR Gen 2

17.86 mbs

mbs

Syntax

mbs *size* [**bytes** | **kilobytes**]

no mbs

Context

[Tree] (config>service>ies>if>sap>ingress>queue-override>queue mbs)

[Tree] (config>service>vprn>if>sap>ingress>queue-override>queue mbs)

[Tree] (config>service>vprn>if>sap>egress>queue-override>queue mbs)

[Tree] (config>service>ies>if>sap>egress>queue-override>queue mbs)

Full Context

configure service ies interface sap ingress queue-override queue mbs

configure service vprn interface sap ingress queue-override queue mbs

configure service vprn interface sap egress queue-override queue mbs

configure service ies interface sap egress queue-override queue mbs

Description

This command overrides specific attributes of the specified queue's MBS parameters. The MBS is a mechanism to override the default maximum size for the queue.

The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer is available when needed or that the packets RED slope is not forced the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The **no** form of this command returns the MBS size assigned to the queue.

Default

mbs default

Parameters

size

This required parameter specifies that the MBS is expressed as an integer representing the required size in either bytes or kilobytes. The default is **kilobytes**. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly to define the size. By specifying the keyword **default** sets the MBS to its default value.

Values 0 to 1073741824, default

bytes

Specifies that the value given for *size* is interpreted as the queue's MBS value in bytes.

kilobytes

Specifies that the value given for *size* is interpreted as the queue's MBS value in kb/s.

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs {*size* [**bytes** | **kilobytes**] | **default**}

no mbs

Context

[Tree] (config>service>vprn>if>sap>ingress>queue-override>queue mbs)

[Tree] (config>service>vprn>if>sap>egress>queue-override>queue mbs)

Full Context

configure service vprn interface sap ingress queue-override queue mbs

configure service vprn interface sap egress queue-override queue mbs

Description

This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The **no** form of this command returns the MBS size assigned to the queue to the value.

Default

mbs default

Parameters

size

The size parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **bytes** and **kilobytes** keywords are mutually exclusive and are used to explicitly define whether the size represents bytes or kilobytes.

Values 0 to 1073741824

default

bytes

When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobytes

When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

default

Keyword that reverts the MBS to its default value.

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs *size* [**bytes** | **kilobytes**]

no mbs

Context

[Tree] (config>service>vpls>sap>egress>queue-override>queue mbs)

[Tree] (config>service>vpls>sap>ingress>queue-override>queue mbs)

Full Context

configure service vpls sap egress queue-override queue mbs

configure service vpls sap ingress queue-override queue mbs

Description

This command overrides specific attributes of the specified queue’s MBS parameters. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS is not guaranteed that a buffer will be available when needed or that the packets RED slope will not force the discard of the packet. Setting correct CBS parameters and controlling CBS over-subscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS assigned to the queue to the default value.

Default

mbs default

Parameters

size

The size parameter is required when specifying mbs and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **bytes** and **kilobytes** keywords are mutually exclusive and are used to explicitly define whether the size represents bytes or kilobytes.

Values 0 to 1073741824

default

bytes

When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobytes

When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs {*size* [**bytes** | **kilobyte**] | **default**}

no mbs

Context

[Tree] (config>card>fp>ingress>network>qgrp>policer-over>plcr mbs)

[Tree] (config>card>fp>ingress>access>qgrp>policer-over>plcr mbs)

Full Context

configure card fp ingress network queue-group policer-override policer mbs

configure card fp ingress access queue-group policer-override policer mbs

Description

This command configures the policer’s PIR leaky bucket’s violate threshold. The **high-prio-only** command is applied to the MBS value to derive the bucket’s low priority violate threshold. For ingress, trusted in-profile packets and untrusted high priority packets use the policer’s high priority violate threshold while trusted out-of-profile and untrusted low priority packets use the policer’s low priority violate threshold.

The PIR bucket’s violate threshold represent the maximum burst tolerance allowed by the policer. If the policer’s offered rate is equal to or less than the policer’s defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low priority violate threshold provides a smaller burst size for the lower priority traffic associated with the

policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by **high-prio-only** is available for the higher priority traffic.

The policer's mbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied.

The **no** form of this command reverts the policer to its default MBS size. By default, the MBS is 16 Mbytes when PIR equals max or is greater than or equal to the FP capacity (this overrides an explicit configured MBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max PIR capped to 3968 kbytes, with a minimum of 256 bytes.

Parameters

size

The *size* parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **bytes** and **kilobytes** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 2683435456, **default**

bytes

When **bytes** is defined, the value given for size is interpreted as the policer's MBS value given in bytes.

kilobytes

When **kilobytes** is defined, the value is interpreted as the policer's MBS value given in kilobytes.

default

Keyword that reverts the MBS to its default value.

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs {*size* [**bytes** | **kilobytes**] | **default**}

no mbs

Context

[Tree] (config>port>ethernet>network>egr>qgrp>qover>q mbs)

[Tree] (config>port>ethernet>access>egr>qgrp>qover>q mbs)

[Tree] (config>port>ethernet>access>ing>qgrp>qover>q mbs)

Full Context

configure port ethernet network egress queue-group queue-overrides queue mbs

configure port ethernet access egress queue-group queue-overrides queue mbs

configure port ethernet access ingress queue-group queue-overrides queue mbs

Description

The Maximum Burst Size (MBS) command specifies the default maximum buffer size for the template queue. The value is given in kilobytes.

The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The **queue-group** or network egress QoS context for mbs provides a mechanism for overriding the default maximum size for the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packets RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

This command applies to egress queue group queues as the queue-delay is only supported on egress queues. This command **the queue-delay** command are mutually exclusive.

The **no** form of this command returns the MBS size assigned to the queue to the value.

Default

mbs default

Parameters

size

The *size* parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 1073741824

bytes

When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobytes

When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

default

Keyword that reverts the MBS to its default value.

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs *size* [**bytes** | **kilobytes**]

no mbs

Context

[Tree] (config>service>epipe>sap>egress>policer-over>plcr mbs)

[Tree] (config>service>epipe>sap>ingress>policer-over>plcr mbs)

Full Context

configure service epipe sap egress policer-override policer mbs

configure service epipe sap ingress policer-override policer mbs

Description

This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured mbs parameter for the specified policer-id.

The **no** form of this command is used to restore the MBS to the default value. By default, the MBS is 16 Mbytes when PIR equals max or is greater than or equal to the FP capacity (this overrides an explicit configured MBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max PIR capped to 3968 kbytes, with a minimum of 256 bytes.

Default

no mbs

Parameters

size

The size parameter is required when specifying mbs override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 2683435456, **default**

bytes

When **bytes** is defined, the value given for *size* is interpreted as the policer MBS value in bytes.

kilobytes

When **kilobytes** is defined, the value given for *size* is interpreted as the policer MBS value in kilobytes.

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs {*size* [**bytes** | **kilobytes**] | **default**}

no mbs

Context

[Tree] (config>service>epipe>sap>ingress>queue-override>queue mbs)

[Tree] (config>service>epipe>sap>egress>queue-override>queue mbs)

Full Context

configure service epipe sap ingress queue-override queue mbs

configure service epipe sap egress queue-override queue mbs

Description

This command overrides specific attributes of the specified queue's MBS parameters. A queue uses its MBS value to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the number of buffers allowed by the MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope associated with a packet. A queue that has not exceeded its MBS is not guaranteed to have buffer available when needed or that the packet's RED slope will not force the discard of the packet. Setting correct CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS assigned to the queue to the default value.

Default

mbs default

Parameters

size

The size parameter is an integer expression of the maximum number of kilobytes or bytes of buffering allowed for the queue. A value of 0 causes the queue to discard all packets.

Values 0 to 1073741824, default

bytes

Indicates that the *size* parameter value is expressed in bytes.

kilobytes

Indicates that the *size* parameter is expressed in kilobytes.

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs *size* [{**bytes** | **kilobytes**}]

no mbs

Context

[Tree] (config>service>vpls>sap>ingress>policer-override>plcr mbs)

[Tree] (config>service>vpls>sap>egress>policer-override>plcr mbs)

Full Context

configure service vpls sap ingress policer-override policer mbs

configure service vpls sap egress policer-override policer mbs

Description

This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured mbs parameter for the specified policer-id.

The **no** form of this command restores the MBS to the default value. By default, the MBS is 16 Mbytes when PIR equals max or is greater than or equal to the FP capacity (this overrides an explicit configured MBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max PIR capped to 3968 kbytes, with a minimum of 256 bytes.

Default

no mbs

Parameters

size

This parameter is required when specifying MBS override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 2683435456, **default**

Default kilobytes

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs *size* [{**bytes** | **kilobytes**}]

no mbs

Context

[Tree] (config>service>ies>if>sap>egress>policer-override>plcr mbs)

[Tree] (config>service>ies>if>sap>ingress>policer-override>plcr mbs)

Full Context

configure service ies interface sap egress policer-override policer mbs

configure service ies interface sap ingress policer-override policer mbs

Description

This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured mbs parameter for the specified policer-id.

The **no** form of this command restores the MBS setting to the default value. By default, the MBS is 16 Mbytes when PIR equals max or is greater than or equal to the FP capacity (this overrides an explicit configured MBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max PIR capped to 3968 kbytes, with a minimum of 256 bytes.

Default

no mbs

Parameters

size

This parameter is required when specifying MBS override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 2683435456, **default**

Default kilobytes

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs *size* [{**bytes** | **kilobytes**}]

no mbs

Context

[Tree] (config>service>vprn>if>sap>egress>policer-override>plcr mbs)

[Tree] (config>service>vprn>if>sap>ingress>policer-override>plcr mbs)

Full Context

configure service vprn interface sap egress policer-override policer mbs

configure service vprn interface sap ingress policer-override policer mbs

Description

This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured mbs parameter for the specified policer-id.

The **no** form of this command restores the MBS to the default value. By default, the MBS is 16 Mbytes when PIR equals max or is greater than or equal to the FP capacity (this overrides an explicit configured MBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max PIR capped to 3968 kbytes, with a minimum of 256 bytes.

Default

no mbs

Parameters

size

This parameter is required when specifying MBS override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 2683435456, **default**

Default kilobytes

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs {*size* [**bytes** | **kilobytes**] | **default**}

no mbs

Context

[Tree] (config>qos>sap-egress>policer mbs)

[Tree] (config>qos>sap-ingress>policer mbs)

Full Context

configure qos sap-egress policer mbs

configure qos sap-ingress policer mbs

Description

This command is used to configure the policer's PIR leaky bucket's high-priority violate threshold. The **high-prio-only** command is applied to the MBS value to derive the bucket's low-priority violate threshold. For ingress, trusted in-profile packets and untrusted high-priority packets use the policer's high-priority violate threshold while trusted out-of-profile and untrusted low-priority packets use the policer's low-priority violate threshold. At egress, inplus-profile, and in-profile packets use the policer's high-priority violate threshold and out-of-profile packets use the policer's low-priority violate threshold. Exceed-profile packets are discarded unless **enable-exceed-pir** is configured, in which case they are forwarded.

The PIR bucket's violate threshold represents the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low-priority violate threshold provides a smaller burst size for the lower priority traffic associated with the policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by **high-prio-only** is available for the higher priority traffic.

The policer's MBS size defined in the QoS policy may be overridden on an SLA profile or SAP where the policy is applied.

The **no** form of this command returns the queue to its default MBS size. By default, the MBS is 16 Mbytes when PIR equals max or is greater than or equal to the FP capacity (this overrides an explicit configured MBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max PIR capped to 3968 kbytes, with a minimum of 256 bytes.

Parameters

size [bytes | kilobytes]

The *size* parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **bytes** and **kilobytes** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 2683435456, **default**

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs {*size* [**bytes** | **kilobytes**] | **default**}

mbs delay-time *microseconds*

mbs delay-percent *percent*

no mbs

Context

[Tree] (config>qos>sap-egress>queue mbs)

Full Context

configure qos sap-egress queue mbs

Description

This command configures the maximum number of buffers, in bytes or kilobytes, allowed for a specific queue. The value overrides the default value for the context.

The **delay-time** command option configures the MBS as a function of the expected delay. The system automatically translates this configuration into kilobytes based on the administrative rate of the queue parent (for example, the scheduler or aggregate-shaper).

The **delay-percent** command option configures the MBS as percentage of the SAP delay budget of the queue configured using the **latency-budget** command.

The **no** form of this command returns the queue to its default MBS.

Default

mbs default

Parameters

size

This parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes.

Default **kilobyte**

default

Keyword to set the MBS to its default value.

bytes

Keyword to interpret the configured value is in bytes.

Values 0 to 1073741824

kilobytes

Keyword to interpret the configured value is in kilobytes.

Values 0 to 1048576

Default kilobytes

microseconds

Specifies the MBS as a function of delay time.

Values 0 to 1000000

percent

Specifies the MBS as a percentage of the SAP latency budget.

Values 0.00 to 100.00

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs {*size* [**bytes** | **kilobytes**] | **default**}

no mbs

Context

[Tree] (config>qos>sap-ingress>queue mbs)

Full Context

configure qos sap-ingress queue mbs

Description

This command configures the maximum number of buffers allowed for a specific queue. The value is given in bytes or kilobytes and overrides the default value for the context.

The **no** form of this command returns the policer to its default MBS.

Default

no mbs

Parameters

size

The *size* parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes.

Default **kilobyte**

default

Sets the MBS to its default value.

bytes

Specifies that the value given for size is interpreted as the queue's MBS value given in bytes.

Values 0 to 2688000

kilobytes

Specifies the value is interpreted as the queue's MBS value given in kilobytes.

Values 0 to 2625

Default kilobytes

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs *percent*

no mbs

Context

[Tree] (config>qos>network-queue>queue mbs)

Full Context

configure qos network-queue queue mbs

Description

This command specifies the relative amount of buffer pool space for the maximum buffers for a specific ingress network FP forwarding class queue or egress network port forwarding class queue. The value is entered as a percentage.

The MBS value is used by a queue to determine whether it has exhausted its total allowed buffers while enqueuing packets. When the queue has exceeded its maximum amount of buffers, all packets are discarded until the queue transmits a packet. A queue that has not exceeded its MBS is not guaranteed to have a buffer available when needed or that the packet's RED slope will not force the discard of the packet. In order to safeguard against queue starvation (when a queue does not receive its fair share of buffers), set proper CBS parameters and control CBS oversubscription. Another safeguard is to properly set the RED slope parameters for the needs of the network queues.

The MBS can sometimes be smaller than the CBS. This will result in a portion of the CBS for the queue to be unused and should be avoided.

The **no** form of this command returns the MBS for the queue to the default for the forwarding class.

Parameters

percent

The percent of buffers from the total buffer pool space for the maximum number of buffers, expressed as a decimal integer. If 10 Mbytes is the total buffer space in the buffer pool, a value of 10 would limit the maximum queue size to 1 Mbyte (10%) of buffer space for the forwarding class queue. If the total size is increased to 20 Mbytes, the existing value of 10 would automatically increase the maximum size of the queue to 2 Mbytes.

Values 0 to 100

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs {*size* [**bytes** | **kilobytes**] | **default**}

no mbs

Context

[Tree] (config>qos>qgrps>egr>qgrp>policer mbs)

[Tree] (config>qos>qgrps>ing>qgrp>policer mbs)

Full Context

configure qos queue-group-templates egress queue-group policer mbs

configure qos queue-group-templates ingress queue-group policer mbs

Description

This command specifies the default maximum buffer size for the template queue in bytes or kilobytes.

The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. When the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The **port>ethernet>access>ingress>queue-group** and **port>ethernet>access>egress>queue-group** contexts for **mbs** provides a mechanism for overriding the default maximum size for the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope that a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard against queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

When configured on an egress queue group queue, this command and the **queue-delay** command are mutually exclusive. In order to change between the **mbs** and **queue-delay** parameters, the current parameter must be removed before adding the new parameter; that is, changing from **mbs** to **queue-delay** requires a **no mbs** before the **queue-delay** is configured and changing from **queue-delay** to **mbs** requires a **no queue-delay** before the **mbs** is configured. If **queue-delay** is configured for an egress queue group queue, it is not possible to override the MBS for that queue.

For policers, this command is used to configure the policer's PIR leaky bucket's high-priority violate threshold. The **high-prio-only** command is applied to the MBS value to derive the bucket's low-priority violate threshold.

At ingress, trusted in-profile packets and untrusted high-priority packets use the policer's high-priority violate threshold while trusted out-of-profile and untrusted low-priority packets use the policer's low-priority violate threshold.

At egress, inplus-profile and in-profile packets use the policer's high-priority violate threshold and out-of-profile packets use the policer's low-priority violate threshold. Exceed-profile packets are discarded unless **enable-exceed-pir** is configured, in which case they are forwarded.

The PIR bucket's violate threshold represents the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low-priority violate threshold provides a smaller burst size for the lower priority traffic associated with the policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by high-prio-only is available for the higher priority traffic.

The policer's mbs size defined in the QoS policy may be overridden on an SLA profile or SAP where the policy is applied.

The **no** form of this command returns the MBS size to its default value. By default, the MBS is 16 Mbytes when PIR equals max or is greater than or equal to the FP capacity (this overrides an explicit configured MBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max PIR capped to 3968 kbytes, with a minimum of 256 bytes.

Default

default

Parameters

size

For queues, the size parameter is an integer expression of the maximum number of bytes or kilobytes of buffering allowed for the queue. For a value of 100 kbytes, enter the value 100. A value of 0 causes the queue to discard all packets. For policers, the size parameter is an integer expression of the maximum number of bytes for the policer's MBS. The queue MBS maximum value used is constrained by the pool size in which the queue exists and by the shared pool space in the corresponding megapool.

Values 0 to 2683435456

Default value: 16 Mbytes when PIR equals max or is greater than or equal to the FP capacity (this overrides an explicitly configured MBS value); otherwise, 10 ms volume of traffic for a configured non-zero/non-max PIR capped to 3968 kbytes, with a minimum of 256 bytes.

[bytes | kilobytes]

Specifies bytes or kilobytes.

Default kilobytes

default

Sets the MBS to its default value.

Platforms

7705 SAR Gen 2

mbs

Syntax

mbs {*size* [**bytes** | **kilobytes**] | **default**}

no mbs

Context

[Tree] (config>qos>qgrps>ing>qgrp>queue mbs)

[Tree] (config>qos>qgrps>egr>qgrp>queue mbs)

Full Context

configure qos queue-group-templates ingress queue-group queue mbs

configure qos queue-group-templates egress queue-group queue mbs

Description

This command specifies the default maximum buffer size for the template queue in bytes or kilobytes.

The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. When the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The **port>ethernet>access>ingress>queue-group** and **port>ethernet>access>egress>queue-group** contexts for **mbs** provides a mechanism for overriding the default maximum size for the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope that a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard against queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

When configured on an egress queue group queue, this command and the **queue-delay** command are mutually exclusive. In order to change between the **mbs** and **queue-delay** parameters, the current parameter must be removed before adding the new parameter; that is, changing from **mbs** to **queue-delay** requires a **no mbs** before the **queue-delay** is configured and changing from **queue-delay** to **mbs** requires

a **no queue-delay** before the **mbs** is configured. If **queue-delay** is configured for an egress queue group queue, it is not possible to override the MBS for that queue.

For policers, this command is used to configure the policer's PIR leaky bucket's high-priority violate threshold. The **high-prio-only** command is applied to the MBS value to derive the bucket's low-priority violate threshold.

At ingress, trusted in-profile packets and untrusted high-priority packets use the policer's high-priority violate threshold while trusted out-of-profile and untrusted low-priority packets use the policer's low-priority violate threshold.

At egress, inplus-profile and in-profile packets use the policer's high-priority violate threshold and out-of-profile packets use the policer's low-priority violate threshold. Exceed-profile packets are discarded unless **enable-exceed-pir** is configured, in which case they are forwarded.

The PIR bucket's violate threshold represents the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low-priority violate threshold provides a smaller burst size for the lower priority traffic associated with the policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by high-prio-only is available for the higher priority traffic.

The policer's mbs size defined in the QoS policy may be overridden on an SLA profile or SAP where the policy is applied.

The **no** form of this command returns the MBS size assigned by the queue.

Default

default

Parameters

size

For queues, the size parameter is an integer expression of the maximum number of bytes or kilobytes of buffering allowed for the queue. For a value of 100 kbytes, enter the value 100. A value of 0 causes the queue to discard all packets. For policers, the size parameter is an integer expression of the maximum number of bytes for the policer's MBS. The queue MBS maximum value used is constrained by the pool size in which the queue exists and by the shared pool space in the corresponding megapool.

Values 0 to 1048576 or **default**

Minimum configurable non-zero value: 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4

Minimum non-zero default value: maximum of 10 ms of CIR, or 6 kbytes on an FP2, 7680 bytes on an FP3, and 16 kbytes on an FP4

[bytes | kilobytes]

Specifies bytes or kilobytes.

Default kilobytes

default

Sets the MBS to its default value.

Platforms

7705 SAR Gen 2

17.87 mbs-contribution

mbs-contribution

Syntax

mbs-contribution *size* [**bytes** | **kilobytes**]

no mbs-contribution

Context

[Tree] (config>card>fp>ing>network>qgrp>policer-ctrl-over>mbs-thrshlds>prio mbs-contribution)

[Tree] (config>card>fp>ingress>access>qgrp>policer-ctrl-over>mbs-thrshlds>prio mbs-contribution)

Full Context

configure card fp ingress network queue-group policer-control-override priority-mbs-thresholds priority mbs-contribution

configure card fp ingress access queue-group policer-control-override priority-mbs-thresholds priority mbs-contribution

Description

This command configures the policy-based burst tolerance for a parent policer instance created when the policy is applied to a SAP or subscriber context. The system uses the parent policer's **min-thresh-separation** value, the priority level's **mbs-contribution** value and the number of child policers currently attached to the priority level to derive the priority level's shared-portion and fair-portion of burst tolerance within the local priority level. The shared-portion and fair-portions for each priority level are then used by the system to calculate each priority level's discard-unfair threshold and discard-all threshold.

The value for a priority level's **mbs-contribution** within the policer-control-policy may be overridden on the SAP or subscriber sub-profile where the policy is applied in order to allow fine tuning of the discard-unfair and discard-all thresholds relevant to the needs of the local child policers on the object.

Accumulative Nature of Burst Tolerance for a Parent Policer Priority Level

When defining **mbs-contribution**, the specified size may only be a portion of the burst tolerance associated with the priority level. The packets associated with the priority level share the burst tolerance of lower within the parent policer. As the parent policer PIR bucket depth increases during congestion, the lower priority packets eventually experience discard based on each priority's discard-unfair and discard-all thresholds. Assuming congestion continues once all the lower priority packets have been prevented from consuming bucket depth, the burst tolerance for the priority level will be consumed by its own packets and any packets associated with higher priorities.

The Effect of Fair and Unfair Child Policer Traffic at a Parent Policer Priority Level

The system continually monitors the offered rate of each child policer on each parent policer priority level and detects when the policer is in a congested state (the aggregate offered load is greater than

the decrement rate defined on the parent policer). As previously stated, the result of congestion is that the parent policer's bucket depth will increase until it eventually hovers around either a discard-unfair or discard-all threshold belonging to one of the priority levels. This threshold is the point where enough packets are being discarded that the increment rate and decrement rate begin to even out. If only a single child policer is associated to the priority level, the discard-unfair threshold is not used since fairness is only applicable when multiple child policers are competing at the same priority level.

When multiple child policers are sharing the congested priority level, the system uses the offered rates and the parenting parameters of each child to determine the fair rate per child when the parent policer is unable to meet the bandwidth needs of each child. The fair rate represents the amount of bandwidth that each child at the priority level should receive relative to the other children at the same level according to the policer control policy instance managing the child policers. This fair rate is applied as the decrement rate for each child's FIR bucket. Changing a child's FIR rate does not modify the amount of packets forwarded by the parent policer for the child's priority level. It simply modifies the forwarded ratio between the children on that priority level. Since each child FIR bucket has some level of burst tolerance before marking its packets as unfair, the current parent policer bucket depth may at times rise above the discard-unfair threshold. The mbs-contribution value provides a means to define how much separation is provided between the priority level's discard-unfair and discard-all threshold to allow the parent policer to absorb some amount of FIR burst before reaching the priority's discard-all threshold.

This level of fair aggregate burst tolerance is based on the decrement rate of the parent policer's PIR bucket while the individual fair bursts making up the aggregate are based on each child's FIR decrement rate. The aggregate fair rate of the priority level is managed by the system with consideration of the current rate of traffic in higher priority levels. In essence, the system ensures that for each iteration of the child FIR rate calculation, the sum of the child FIR decrement rates plus the sum of the higher priority traffic increment rates equals the parent policers decrement rate. This means that dynamic amounts of higher priority traffic can be ignored when sizing a lower priority's fair aggregate burst tolerance. Consider the following:

- The parent policer decrement rate is set to 20 Mb/s (max-rate 20,000).
- A priority level's fair burst size is set to 30 kbytes (mbs-contribution 30 kilobytes).
- Higher priority traffic is currently taking 12 Mb/s.
- The priority level has three child policers attached.
- Each child's PIR MBS is set to 10 kbytes, which makes each child's FIR MBS 10 kbytes.
- The children want 10 Mb/s, but only 8 Mb/s is available.
- The following table describes the FIR rates of the children based on weights.

Table 68: FIR Rates of the Children Based on Weights

	FIR Rate	FIR MBS
Child 1	4 Mb/s	10 kbytes
Child 2	3 Mb/s	10 kbytes
Child 3	1 Mb/s	10 kbytes

The 12 Mb/s of the higher priority traffic and the 8 Mb/s of fair traffic equal the 20 Mb/s decrement rate of the parent policer.

It is clear that the higher priority traffic is consuming 12 Mb/s of the parent policer's decrement rate, leaving 8 Mb/s of decrement rate for the lower priority's fair traffic.

- The burst tolerance of child 1 is based on 10 kbytes above 4 Mb/s,
- The burst tolerance of child 2 is based on 10 kbytes above 3 Mb/s,
- The burst tolerance of child 3 is based on 10 kbytes above 1 Mb/s.

If all three children burst simultaneously (unlikely), they will consume 30 kbytes above 8 Mb/s. This is the same as the remaining decrement rate after the higher priority traffic.

Parent Policer Total Burst Tolerance and Downstream Buffering

The highest in-use priority level's discard-all threshold is the total burst tolerance of the parent policer. In some cases the parent policer represents downstream bandwidth capacity and the max-rate of the parent policer is set to prevent overrunning the downstream bandwidth. The burst tolerance of the parent policer defines how much more traffic may be sent beyond the downstream scheduling capacity. In the worst case scenario, when the downstream buffering is insufficient to handle the total possible burst from the parent policer, downstream discards based on lack of buffering may occur. However, in all likelihood, this is not the case.

In most cases, lower priority traffic in the policer will be responsible for the greater part of congestion above the parent policer rate. Since this traffic is discarded with a lower threshold, this lowers the effective burst tolerance even while the highest priority traffic is present.

Configuring a Priority Level's MBS Contribution Value

In the most conservative case, a priority level's **mbs-contribution** value may be set to be greater than the sum of child policer's MBS and one max-size-frame per child policer. This ensures that even in the absolute worst case where all the lower priority levels are simultaneously bursting to the maximum capacity of each child, enough burst tolerance for the priority's children will exist if they also burst to their maximum capacity.

Since simply adding up all the child policer's PIR MBS values may result in large overall burst tolerances that are not ever likely to be needed, you should consider some level of burst oversubscription when configuring the **mbs-contribution** value for each priority level. The amount of oversubscription should be determined based on the needs of each priority level.

Using the Fixed Keyword to Create Deterministic Parent Policer Discard Thresholds

In the default behavior, the system ignores the **mbs-contribution** values for a priority level on a subscriber or SAP parent policer when a child policer is not currently associated with the level. This prevents additional burst tolerance from being added to higher priority traffic within the parent policer.

This does cause fluctuations in the defined threshold values when child policers are added or removed from a parent policer instance. If this behavior is undesirable, the fixed keyword may be used which causes the **mbs-contribution** value to always be included in the calculation of parent policer's discard thresholds. The defined **mbs-contribution** value may be overridden on a subscriber SLA profile or on a SAP instance, but the fixed nature of the contribution cannot be overridden.

If the defined **mbs-contribution** value for the priority level is zero, the priority level will have no effect on the parent policer's defined discard thresholds. A packet associated with the priority level will use the next lower priority level's discard-unfair and discard-all thresholds.

The **no** form of this command reverts to the policy's priority level's MBS contribution to the default value. When changed, the thresholds for the priority level and all higher priority levels for all instances of the parent policer are recalculated.

Default

no mbs-contribution

Parameters

size

Specifies that the size parameter is required when executing the **mbs-contribution** command. It is expressed as an integer and specifies the priority's specific portion amount of accumulative MBS for the priority level in bytes or kilobytes which is selected by the trailing **bytes** or **kilobytes** keywords. If both **bytes** and **kilobytes** are missing, **kilobytes** is assumed. Setting this value has no effect on parent policer instances where the priority level's **mbs-contribution** value has been overridden. Clearing an override on parent policer instance causes this value to be enforced.

Values 0 to 16777216

bytes, kilobytes

Specifies that the **bytes** keyword is optional and is mutually exclusive with the **kilobytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in bytes.

The **kilobytes** keyword is optional and is mutually exclusive with the **bytes** keyword. When specified, size is interpreted as specifying the size of min-thresh-separation in kilobytes.

Default **kilobytes**

Platforms

7705 SAR Gen 2

mbs-contribution

Syntax

mbs-contribution *size* [**bytes** | **kilobytes**]

Context

[Tree] (config>service>epipe>sap>ingress>policer-ctrl-over>mbs-thrshlds>priority mbs-contribution)

[Tree] (config>service>epipe>sap>egress>policer-ctrl-over>mbs-thrshlds>priority mbs-contribution)

Full Context

configure service epipe sap ingress policer-control-override priority-mbs-thresholds priority mbs-contribution

configure service epipe sap egress policer-control-override priority-mbs-thresholds priority mbs-contribution

Description

The mbs-contribution override command within the SAP ingress and egress contexts is used to override a parent policer's priority level's mbs-contribution parameter that is defined within the policer-control-policy applied to the SAP. This override allow the priority level's burst tolerance to be tuned based on the needs of the SAP's child policers attached to the priority level.

When the override is defined, modifications to the policer-control-policy priority level's mbs-contribution parameter have no effect on the SAP's parent policer priority level until the override is removed using the no mbs-contribution command within the SAP.

The **no** form of this command removes the override and allows the mbs-contribution setting from the policer-control-policy to control the parent policer's priority level's burst tolerance.

Default

no mbs-contribution

Parameters

size

Specifies the mbs-contribution override value.

Values 1 to 16777216 | default

bytes

Specifies that *size* is expressed in bytes. The bytes and kilobytes keywords are mutually exclusive and optional. The default is kilobytes.

kilobytes

Specifies that *size* is expressed in kilobytes. The bytes and kilobytes keywords are mutually exclusive and optional. The default is kilobytes.

Platforms

7705 SAR Gen 2

mbs-contribution

Syntax

mbs-contribution *size* [{**bytes** | **kilobytes**}]

Context

[Tree] (config>service>vpls>sap>ingress>policer-ctrl-over>mbs-thrshlds>priority mbs-contribution)

[Tree] (config>service>vpls>sap>egress>policer-ctrl-over>mbs-thrshlds>priority mbs-contribution)

Full Context

configure service vpls sap ingress policer-control-override priority-mbs-thresholds priority mbs-contribution

configure service vpls sap egress policer-control-override priority-mbs-thresholds priority mbs-contribution

Description

The mbs-contribution override command within the SAP ingress and egress contexts is used to override a parent policer's priority level's mbs-contribution parameter that is defined within the policer-control-policy applied to the SAP. This override allow the priority level's burst tolerance to be tuned based on the needs of the SAP's child policers attached to the priority level.

When the override is defined, modifications to the policer-control-policy priority level's mbs-contribution parameter have no effect on the SAP's parent policer priority level until the override is removed using the no mbs-contribution command within the SAP.

The **no** form of this command removes the override and allows the mbs-contribution setting from the policer-control-policy to control the parent policer's priority level's burst tolerance.

Default

no mbs-contribution

Parameters

size

This parameter is required when specifying MBS contribution override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 16777216 or default

Default kilobytes

Platforms

7705 SAR Gen 2

mbs-contribution

Syntax

mbs-contribution *size* [{**bytes** | **kilobytes**}]

Context

[Tree] (config>service>ies>if>sap>ingress>policer-ctrl-over>mbs-thrshlds>priority mbs-contribution)

[Tree] (config>service>ies>if>sap>egress>policer-ctrl-over>mbs-thrshlds>priority mbs-contribution)

Full Context

configure service ies interface sap ingress policer-control-override priority-mbs-thresholds priority mbs-contribution

configure service ies interface sap egress policer-control-override priority-mbs-thresholds priority mbs-contribution

Description

The mbs-contribution override command within the SAP ingress and egress contexts is used to override a parent policer's priority level's mbs-contribution parameter that is defined within the policer-control-policy applied to the SAP. This override allow the priority level's burst tolerance to be tuned based on the needs of the SAP's child policers attached to the priority level.

When the override is defined, modifications to the policer-control-policy priority level's mbs-contribution parameter have no effect on the SAP's parent policer priority level until the override is removed using the no mbs-contribution command within the SAP.

The **no** form of this command removes the override and allows the mbs-contribution setting from the policer-control-policy to control the parent policer's priority level's burst tolerance.

Default

no mbs-contribution

Parameters

size

This parameter is required when specifying MBS contribution override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 16777216 or default

Default kilobytes

Platforms

7705 SAR Gen 2

mbs-contribution

Syntax

mbs-contribution *size* [**bytes** | **kilobytes**]

Context

[Tree] (config>service>vprn>if>sap>egress>policer-ctrl-over>mbs-thrshlds>priority mbs-contribution)

[Tree] (config>service>vprn>if>sap>ingress>policer-ctrl-over>mbs-thrshlds>priority mbs-contribution)

Full Context

configure service vprn interface sap egress policer-control-override priority-mbs-thresholds priority mbs-contribution

configure service vprn interface sap ingress policer-control-override priority-mbs-thresholds priority mbs-contribution

Description

The mbs-contribution override command within the SAP ingress and egress contexts is used to override a parent policer's priority level's mbs-contribution parameter that is defined within the policer-control-policy applied to the SAP. This override allow the priority level's burst tolerance to be tuned based on the needs of the SAP's child policers attached to the priority level.

When the override is defined, modifications to the policer-control-policy priority level's mbs-contribution parameter have no effect on the SAP's parent policer priority level until the override is removed using the no mbs-contribution command within the SAP.

The **no** form of this command removes the override and allows the mbs-contribution setting from the policer-control-policy to control the parent policer's priority level's burst tolerance.

Default

no mbs-contribution

Parameters

size

This parameter is required when specifying MBS contribution override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 16777216 or default

Default kilobytes

Platforms

7705 SAR Gen 2

mbs-contribution

Syntax

mbs-contribution *size* [**bytes** | **kilobytes**] [**fixed**]

no mbs-contribution

Context

[Tree] (config>qos>plcr-ctrl-plcy>root>priority-mbs-thresholds>priority mbs-contribution)

Full Context

configure qos policer-control-policy root priority-mbs-thresholds priority mbs-contribution

Description

The **mbs-contribution** command is used to configure the policy-based burst tolerance for a parent policer instance created when the policy is applied to a SAP or a subscriber context. The system uses the parent policer's **min-thresh-separation** value, the priority level's **mbs-contribution** value, and the number of child policers currently attached to the priority level to derive the priority level's shared-portion and fair-portion of burst tolerance within the local priority level. The shared-portion and fair-portions for each priority level are then used by the system to calculate each priority level's discard-unfair threshold and discard-all threshold. The mbs-contribution is the minimum separation between two adjacent active discard-all thresholds.

The value for a priority level's **mbs-contribution** within the policer-control-policy may be overridden on the SAP or subscriber sub-profile where the policy is applied in order to allow fine tuning of the discard-unfair and discard-all thresholds relevant to the needs of the local child policers on the object.

Accumulative Nature of Burst Tolerance for a Parent Policer Priority Level

When defining **mbs-contribution**, the specified size may only be a portion of the burst tolerance associated with the priority level. The packets associated with the priority level share the burst tolerance of lower within the parent policer. As the parent policer PIR bucket depth increases during congestion, the lower priority packets eventually experience discard based on each priority's discard-unfair and discard-all thresholds. Assuming congestion continues when all the lower priority packets have been prevented from consuming bucket depth, the burst tolerance for the priority level will be consumed by its own packets and any packets associated with higher priorities.

The Effect of Fair and Unfair Child Policer Traffic at a Parent Policer Priority Level

The system continually monitors the offered rate of each child policer on each parent policer priority level and detects when the policer is in a congested state (the aggregate offered load is greater than the decrement rate defined on the parent policer). As previously stated, the result of congestion is that the parent policer's bucket depth will increase until it eventually hovers around either a discard-unfair or discard-all threshold belonging to one of the priority levels. This threshold is the point where enough packets are being discarded that the increment rate and decrement rate begin to even out. If only a single child policer is associated with the priority level, the discard-unfair threshold is not used since fairness is only applicable when multiple child policers are competing at the same priority level.

When multiple child policers are sharing the congested priority level, the system uses the offered rates and the parenting parameters of each child to determine the fair rate per child when the parent policer is unable to meet the bandwidth needs of each child. The fair rate represents the amount of bandwidth that each child at the priority level should receive relative to the other children at the same level according to the policer control policy instance managing the child policers. This fair rate is applied as the decrement rate for each child's FIR bucket. Changing a child's FIR rate does not modify the number of packets forwarded by the parent policer for the child's priority level. It just modifies the forwarded ratio between the children on that priority level. Since each child FIR bucket has some level of burst tolerance before marking its packets as unfair, the current parent policer bucket depth may at times rise above the discard-unfair threshold. The mbs-contribution value provides a means to define how much separation is provided between the priority level's discard-unfair and discard-all threshold to allow the parent policer to absorb some amount of FIR burst before reaching the priority's discard-all threshold.

This level of fair aggregate burst tolerance is based on the decrement rate of the parent policer's PIR bucket while the individual fair bursts making up the aggregate are based on each child's FIR decrement rate. The aggregate fair rate of the priority level is managed by the system with consideration of the current rate of traffic in higher priority levels. In essence, the system ensures that for each iteration of the child FIR rate calculation, the sum of the child FIR decrement rates plus the sum of the higher priority traffic increment rates equals the parent policers decrement rate. This means that dynamic amounts of higher priority traffic can be ignored when sizing a lower priority's fair aggregate burst tolerance. Consider the following:

- The parent policer decrement rate is set to 20 Mb/s (max-rate 20,000).
- A priority level's fair burst size is set to 30 kbytes (mbs-contribution 30 kbytes).
- Higher priority traffic is currently taking 12 Mb/s.
- The priority level has three child policers attached.
- Each child's PIR MBS is set to 10 kbytes, which makes each child's FIR MBS 10 kbytes.
- The children want 10 Mb/s, but only 8 Mb/s is available

- Based on weights, the children's FIR rates are set as follows.

Table 69: FIR Rates of the Children Based on Weights

	FIR Rate	FIR MBS
Child 1	4 Mb/s	10 kbytes
Child 2	3 Mb/s	10 kbytes
Child 3	1 Mb/s	10 kbytes

The 12 Mb/s of the higher priority traffic and the 8 Mb/s of fair traffic equal the 20 Mb/s decrement rate of the parent policer.

It is clear that the higher priority traffic is consuming 12 Mb/s of the parent policer’s decrement rate, leaving 8 Mb/s of decrement rate for the lower priority’s fair traffic.

- The burst tolerance of child 1 is based on 10 kbytes above 4 Mb/s.
- The burst tolerance of child 2 is based on 10 kbytes above 3 Mb/s.
- The burst tolerance of child 3 is based on 10 kbytes above 1 Mb/s.

If all three children burst simultaneously (unlikely), they will consume 30 kbytes above 8 Mb/s. This is the same as the remaining decrement rate after the higher priority traffic.

Parent Policer Total Burst Tolerance and Downstream Buffering

The highest in-use priority level’s discard-all threshold is the total burst tolerance of the parent policer. In some cases, the parent policer represents downstream bandwidth capacity and the max-rate of the parent policer is set to prevent overrunning the downstream bandwidth. The burst tolerance of the parent policer defines how much more traffic may be sent beyond the downstream scheduling capacity. In the worst-case scenario, when the downstream buffering is insufficient to handle the total possible burst from the parent policer, downstream discards based on lack of buffering may occur. However, in all likelihood, this is not the case.

In most cases, lower priority traffic in the policer will be responsible for the greater part of congestion above the parent policer rate. Since this traffic is discarded with a lower threshold, this lowers the effective burst tolerance even while the highest priority traffic is present.

Configuring a Priority Level's MBS Contribution Value

In the most conservative case, a priority level’s **mbs-contribution** value may be set to be greater than the sum of child policer’s mbs and one max-size-frame per child policer. This ensures that even in the absolute worst case where all the lower priority levels are simultaneously bursting to the maximum capacity of each child, enough burst tolerance for the priority’s children will exist if they also burst to their maximum capacity.

Since simply adding up all the child policer’s PIR MBS values may result in large overall burst tolerances that are not ever likely to be needed, consider some level of burst oversubscription when configuring the **mbs-contribution** value for each priority level. The amount of oversubscription should be determined based on the needs of each priority level.

Using the Fixed Keyword to Create Deterministic Parent Policer Discard Thresholds

In the default behavior, the system ignores the **mbs-contribution** values for a priority level on a subscriber or SAP parent policer when a child policer is not currently associated with the level. This prevents additional burst tolerance from being added to higher priority traffic within the parent policer.

This does cause fluctuations in the defined threshold values when child policers are added or removed from a parent policer instance. If this behavior is undesirable, the fixed keyword may be used that causes the **mbs-contribution** value to always be included in the calculation of parent policer's discard thresholds. The defined **mbs-contribution** value may be overridden on a subscriber sla-profile or on a SAP instance, but the fixed nature of the contribution cannot be overridden.

If the defined **mbs-contribution** value for the priority level is zero, the priority level will have no effect on the parent policer's defined discard thresholds. A packet associated with the priority level will use the next lower priority level's discard-unfair and discard-all thresholds.

The **no** form of this command returns the policy's priority level's MBS contribution to the default value. When changed, the thresholds for the priority level and all higher priority levels for all instances of the parent policer will be recalculated.

Parameters

size

The size parameter is required when executing the **mbs-contribution** command. It is expressed as an integer and specifies the priority's specific portion amount of accumulative MBS for the priority level. Setting this value has no effect on parent policer instances where the priority level's **mbs-contribution** value has been overridden.

Values 0 to 16777216 or **default**

Default 8

bytes | kilobytes:

This parameter indicates whether the size is expressed in bytes or kilobytes.

Default **kilobytes**

fixed

The optional fixed keyword is used to force the inclusion of the defined **mbs-contribution** value in the parent policer's discard threshold calculations. If the **mbs-contribution** command is executed without the **fixed** keyword, the fixed calculation behavior for the priority level is removed.

Platforms

7705 SAR Gen 2

17.88 mbytes

mbytes

Syntax

mbytes {*mbytes* | **disable**}

no mbytes

Context

[Tree] (config>system>security>ssh>key-re-exchange>server mbytes)

[Tree] (config>system>security>ssh>key-re-exchange>client mbytes)

Full Context

configure system security ssh key-re-exchange server mbytes

configure system security ssh key-re-exchange client mbytes

Description

This command configures the maximum bytes to be transmitted before a key re-exchange is initiated by the server.

The **no** form of this command reverts to the default value.

Default

mbytes 1024

Parameters

mbytes

Specifies the number of megabytes, on a SSH session, after which the SSH client initiates the key-re-exchange.

Values 1 to 64000

Default 1024

disable

Specifies that a session will never timeout. To re-enable **mbytes**, enter the command without the **disable** option.

Platforms

7705 SAR Gen 2

17.89 mc-endpoint

mc-endpoint

Syntax

[**no**] **mc-endpoint**

Context

[Tree] (config>redundancy>multi-chassis>peer mc-endpoint)

Full Context

configure redundancy multi-chassis peer mc-endpoint

Description

This command specifies that the endpoint is multi-chassis. This value should be the same on both MC-EP peers for the pseudowires that must be part of the same group.

The **no** form of this command removes the endpoint from the MC-EP. Single chassis behavior applies.

Default

no mc-endpoint

Platforms

7705 SAR Gen 2

mc-endpoint

Syntax

mc-endpoint *mc-ep-id*

no mc-endpoint

Context

[Tree] (config>service>vpls>endpoint mc-endpoint)

Full Context

configure service vpls endpoint mc-endpoint

Description

This command specifies the identifier associated with the multi-chassis endpoint. This value should be the same on both MC-EP peers for the pseudowires that must be part of the same group.

The **no** form of this command removes the endpoint from the MC-EP. Single chassis behavior applies.

Default

no mc-endpoint

Parameters

mc-ep-id

Specifies a multi-chassis endpoint ID

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

17.90 mc-ep-peer

mc-ep-peer

Syntax

mc-ep-peer *name*

mc-ep-peer *ip-address*

no mc-ep-peer

Context

[Tree] (config>service>vpls>endpoint>mc-ep mc-ep-peer)

Full Context

configure service vpls endpoint mc-endpoint mc-ep-peer

Description

This command adds multi-chassis endpoint object.

The **no** form of this command removes the multi-chassis endpoint object.

Default

no mc-ep-peer

Parameters

name

Specifies the name of the multi-chassis endpoint peer

ip-address

Specifies the IP address of multi-chassis endpoint peer

Platforms

7705 SAR Gen 2

17.91 mc-ipsec

mc-ipsec

Syntax

[**no**] **mc-ipsec**

Context

[Tree] (config>redundancy>multi-chassis>peer mc-ipsec)

Full Context

configure redundancy multi-chassis peer mc-ipsec

Description

Commands in this context configure multi-chassis peer parameters.

Platforms

7705 SAR Gen 2

17.92 mc-ipsec-non-forwarding

mc-ipsec-non-forwarding

Syntax

[**no**] **mc-ipsec-non-forwarding tunnel-grp-id**

Context

[Tree] (config>vrrp>policy>priority-event mc-ipsec-non-forwarding)

Full Context

configure vrrp policy priority-event mc-ipsec-non-forwarding

Description

This command configures an instance of a multi-chassis IPsec tunnel-group Priority Event used to override the base priority value of a VRRP virtual router instance depending on the operational state of the event.

Parameters

tunnel-grp-id

Identifies the multi-chassis IPsec tunnel group whose non-forwarding state is monitored by this priority control event.

Platforms

7705 SAR Gen 2

17.93 mc-lag

mc-lag

Syntax

[**no**] **mc-lag**

Context

[Tree] (config>redundancy>multi-chassis>peer mc-lag)

Full Context

configure redundancy multi-chassis peer mc-lag

Description

Commands in this context configure multi-chassis LAG operations and related parameters.

The **no** form of this command administratively disables multi-chassis LAG. MC-LAG can be issued only when mc-lag is shutdown.

Default

no mc-lag

Platforms

7705 SAR Gen 2

17.94 mc-maximum-routes

mc-maximum-routes

Syntax

mc-maximum-routes *number* [**log-only**] [**threshold** *threshold*]

Context

[Tree] (config>service>vprn mc-maximum-routes)

Full Context

configure service vprn mc-maximum-routes

Description

This command specifies the maximum number of multicast routes that can be held in the form of this command in a VPN routing or forwarding (VRF) context. When this limit is reached, a log and SNMP trap are sent. If the **log-only** parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then no new joins are processed.

The **no** form of this command disables the limit of multicast routes within a VRF context. Issue the **no** form of this command only when the VPRN instance is shutdown.

Default

no mc-maximum-routes

Parameters

number

Specifies the maximum number of routes to be held in a VRF context.

Values 1 to 2147483647

log-only

Specifies that if the maximum limit is reached, only log the event. **log-only** does not disable the learning of new routes.

threshold

Specifies the percentage at which a warning log message and SNMP trap should be sent.

Values 0 to 100

Default 10

Platforms

7705 SAR Gen 2

mc-maximum-routes

Syntax

mc-maximum-routes *number* [**log-only**] [**threshold** *threshold*]

no mc-maximum-routes

Context

[Tree] (config>router mc-maximum-routes)

Full Context

configure router mc-maximum-routes

Description

This command specifies the maximum number of multicast routes that can be held within a VPN routing/forwarding (VRF) context. When this limit is reached, a log and SNMP trap are sent. If the **log-only** parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then no new joins will be processed.

The **no** form of this command disables the limit of multicast routes within a VRF context. Issue the **no** form of this command only when the VPRN instance is shutdown.

Default

no mc-maximum-routes

Parameters

number

Specifies the maximum number of routes to be held in a VRF context.

Values 1 to 2147483647

log-only

Specifies that if the maximum limit is reached, only log the event. **log-only** does not disable the learning of new routes.

threshold

Specifies the percentage at which a warning log message and SNMP trap should be sent.

Values 0 to 100

Default 10

Platforms

7705 SAR Gen 2

17.95 mc-ring

mc-ring

Syntax

mc-ring

Context

[Tree] (config>redundancy>multi-chassis>peer mc-ring)

Full Context

configure redundancy multi-chassis peer mc-ring

Description

Commands in this context configure the multi-chassis ring parameters.

The **no** form of this command reverts to the default.

Default

mc-ring

Platforms

7705 SAR Gen 2

mc-ring

Syntax

[**no**] **mc-ring**

Context

[Tree] (config>redundancy>multi-chassis>peer>sync mc-ring)

Full Context

configure redundancy multi-chassis peer sync mc-ring

Description

This command specifies whether multi-chassis ring information should be synchronized with the multi-chassis peer.

Default

no mc-ring

Platforms

7705 SAR Gen 2

17.96 mcast-ipv6-snooping-scope

mcast-ipv6-snooping-scope

Syntax

mcast-ipv6-snooping-scope {**mac-based** | **sg-based**}

no mcast-ipv6-snooping-scope

Context

[Tree] (config>service>vpls mcast-ipv6-snooping-scope)

Full Context

configure service vpls mcast-ipv6-snooping-scope

Description

This command specifies the forwarding scope used for IPv6 multicast traffic when PIM snooping for IPv6 is enabled.

By default, the scope is **mac-based**; IPv6 snooped multicast traffic is forwarded is based on the low-order 32 bits of the destination IPv6 address.

When the scope is configured as **sg-based**, the IPv6 snooped multicast traffic is forwarded based on both its full source (if specified in the join) and destination IPv6 address. SG-based forwarding is only supported on FP3- (or higher) based line cards.

PIM snooping for IPv6 must be disabled to change the forwarding mode within a VPLS service between **mac-based** and **sg-based**.

The **no** form of this command configures the router to use the default value.

Default

mcast-ipv6-snooping-scope mac-based

Parameters

mac-based

Sets forwarding for PIM-snooped IPv6 multicast traffic based on the low-order 32 bits of its destination IPv6 address.

sg-based

Sets forwarding for PIM-snooped IPv6 multicast traffic based on its full source (if specified in the join) and destination IPv6 address.

Platforms

7705 SAR Gen 2

17.97 mcr-default-gtw

mcr-default-gtw

Syntax

mcr-default-gtw

Context

[Tree] (config>service>vpls mcr-default-gtw)

Full Context

configure service vpls mcr-default-gtw

Description

Commands in this context configure the default gateway information when using Dual Homing in L2-TPSDA. The IP and MAC address of the default gateway used for subscribers on an L2 MC-Ring are configured in this context. After a ring heals or fails, the system sends out a gratuitous ARP on an active ring SAP in order to attract traffic from subscribers on the ring with connectivity to that SAP.

Platforms

7705 SAR Gen 2

17.98 mcs

mcs

Syntax

mcs [*ip-int-name*]

no mcs

Context

[Tree] (debug>router>igmp mcs)

Full Context

debug router igmp mcs

Description

This command enables debugging for IGMP multicast servers (MCS).

The **no** form of the command disables the IGMP interface debugging for the specifies interface name.

Parameters

ip-int-name

Debugs the information associated with the specified IP interface name.

Values IP interface address

Platforms

7705 SAR Gen 2

17.99 md

md

Syntax

md *file-url*

Context

[Tree] (file md)

Full Context

file md

Description

This command creates a new directory in a file system.

Directories can only be created one level at a time.

Parameters

file-url

Specifies the directory name to be created.

Values		
	local-url	[*cflash-id/*][*file-path*] up to 200 characters, including cflash-id directory length up to 99 each
	remote-url	[{ftp:// | tftp://}*login*:*pswd*@*remote-locn*/][*file-path*]
		up to 247 characters
		directory length up to 99 characters each
	remote-locn	[hostname | ipv4-address | [ipv6-address]]
	ipv4-address	*a.b.c.d*
	ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
		x:x:x:x:x:x:d.d.d.d[-*interface*]
		x - [0 to FFFF]H
		d - [0 to 255]D
		interface - up to 32 characters, for link local addresses 255
	cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

18 m Commands – Part II

18.1 md-auto-id

md-auto-id

Syntax

md-auto-id

Context

[Tree] (config>qos md-auto-id)

Full Context

configure qos md-auto-id

Description

This command automatically assigns numerical ID values for QoS policies in model-driven (MD) management interfaces.

Classic management interfaces use a numerical policy ID as the primary key for sap-ingress, sap-egress, and network QoS policies. In model-driven interfaces, SAP and network policies use string names as keys. The SAP and network policies can optionally be created in MD interfaces without having to explicitly select and specify a numerical policy ID. In this case, SR OS assigns an ID using the configured ID range.

Platforms

7705 SAR Gen 2

md-auto-id

Syntax

md-auto-id

Context

[Tree] (config>filter md-auto-id)

Full Context

configure filter md-auto-id

Description

This command automatically assigns numerical ID values for filter policies in model-driven management interfaces.

Classic management interfaces use a numerical filter ID as the primary key for IP filters, IPv6 filters, and MAC filters. In model-driven interfaces, IP, IPv6, and MAC filters use string names as keys. The filters can optionally be created in MD interfaces without having to explicitly select and specify a numerical filter ID. In this case, SR OS assigns an ID using the configured ID range.

Platforms

7705 SAR Gen 2

md-auto-id

Syntax

md-auto-id

Context

[Tree] (config>service md-auto-id)

Full Context

configure service md-auto-id

Description

This command automatically assigns numerical ID values for services, customers, and PW templates in model-driven (MD) management interfaces.

Classic management interfaces use a numerical service ID, customer ID, and PW template ID as the primary key for services, customers, and PW templates. In model-driven interfaces, services, customers, and PW templates use string names as keys. The services, customers, and PW templates can optionally be created in MD interfaces without having to explicitly select and specify a numerical ID. In this case, SR OS assigns an ID using the configured ID range.

Platforms

7705 SAR Gen 2

18.2 md-cli

md-cli

Syntax

md-cli

Context

[Tree] (config>system>management-interface>cli md-cli)

Full Context

configure system management-interface cli md-cli

Description

Commands in this context configure the MD-CLI management interface.

Platforms

7705 SAR Gen 2

md-cli

Syntax

md-cli

Context

[Tree] (config>system>security>management-interface md-cli)

Full Context

configure system security management-interface md-cli

Description

Commands in this context configure hash-control for the MD-CLI interface.

Platforms

7705 SAR Gen 2

18.3 md-cli-session

md-cli-session

Syntax

md-cli-session {**permit** | **deny**}

Context

[Tree] (config>system>security>profile>grpc>rpc-authorization md-cli-session)

Full Context

configure system security profile grpc rpc-authorization md-cli-session

Description

This command configures the use of the MdCli Session RPC for the user profile.

The **no** form of this command reverts to the default value.

Default

md-cli-session permit

Parameters

deny

Specifies that the use of MdCli Session RPC is denied.

permit

Specifies that the use of MdCli Session RPC is permitted.

Platforms

7705 SAR Gen 2

18.4 md-interfaces

md-interfaces

Syntax

[**no**] **md-interfaces**

Context

[Tree] (config>system>security>management-interface>output-authorization md-interfaces)

Full Context

configure system security management-interface output-authorization md-interfaces

Description

This command controls output authorization of commands or RPCs for model-driven interfaces that display configuration or state.

When enabled, output authorization is performed for the following commands:

- MD-CLI **info** and **compare** commands
- NETCONF <get> and <get-config> RPCs
- gNMI Get RPC

When disabled, output authorization is not performed, which may significantly decrease the system response time by reducing command authorization requests, especially when remote AAA servers are used. Input to edit configuration is always authorized based on the AAA configuration.

By default, authorization checks are performed for configuration and state output.

The **no** form of this command disables authorization checks, allowing the output to be displayed immediately.

Default

md-interfaces

Platforms

7705 SAR Gen 2

18.5 mda

mda

Syntax

[**no**] **mda** *mda-slot*

Context

[Tree] (config>card mda)

Full Context

configure card mda

Description

This mandatory command enables access to a card's MDA context. In SR OS, MDAs cover MDA and XMA.

Parameters

mda-slot

Specifies the MDA slot number to be configured. Slots are numbered 1 and 2. On vertically oriented slots, the top MDA slot is number 1, and the bottom MDA slot is number 2. On horizontally oriented slots, the left MDA is number 1, and the right MDA slot is number 2.

Values 1, 2

Platforms

7705 SAR Gen 2

mda

Syntax

[**no**] **mda** *mda-id*

Context

[Tree] (config>service>vpls>mesh-sdp>egress>mfib-allowed-mda-destinations mda)

[Tree] (config>service>vpls>spoke-sdp>egress>mfib-allowed-mda-destinations mda)

Full Context

configure service vpls mesh-sdp egress mfib-allowed-mda-destinations mda

configure service vpls spoke-sdp egress mfib-allowed-mda-destinations mda

Description

This command specifies an MFIB-allowed MDA destination for an SDP binding configured in the system.

Parameters

mda-id

Specifies an MFIB-allowed MDA destination

Values

slot/mda	slot: 1 to 10
	mda: 1, 2
slot/xiom/mda	slot: 1 to 10
	xiom: "x1" or "x2"
	mda: 1, 2

Platforms

7705 SAR Gen 2

mda

Syntax

[**no**] **mda** *mda-id*

Context

[Tree] (config>isa>tunnel-grp mda)

Full Context

configure isa tunnel-group mda

Description

This command specifies the MDA ID of the MS-ISA as the member of tunnel-group with multi-active enabled. Up to 16 MDA could be configured under the same tunnel-group.

Parameters

mda-id

Specifies the id of MS-ISA.

Values iom-slot-id/mda-slot-id

Platforms

7705 SAR Gen 2

mda

Syntax

[**no**] **mda** *mda-id*

Context

[Tree] (config>isa>nat-group mda)

Full Context

configure isa nat-group mda

Description

This command configures an ISA NAT group MDA.

Parameters

mda-id

Specifies the MDA ID in the *slot/mda* format.

Values slot: 1 to 10

mda: 1 to 16

Note:
Available parameter values may differ by platform.

Platforms

7705 SAR Gen 2

mda

Syntax

[**no**] **mda** *mda-id*

Context

[Tree] (config>service>pw-template>egress>mfib-mda mda)

Full Context

configure service pw-template egress mfib-allowed-mda-destinations mda

Description

This command specifies an MFIB-allowed media adapter destination for an SDP binding configured in the system.

Parameters

mda-id

Specifies an MFIB-allowed media adapters destination.

Values 1, 2

Platforms

7705 SAR Gen 2

mda

Syntax

mda *mda*

no mda

Context

[Tree] (config>isa>tunnel-mem-pool mda)

Full Context

configure isa tunnel-member-pool mda

Description

This command configures an association between an MDA and the tunnel member pool.

The **no** form of this command removes the association between the MDA and the tunnel member pool.

Parameters

name

Specifies the name of the MDA, up to 32 characters.

Platforms

7705 SAR Gen 2

18.6 mda-type

mda-type

Syntax

mda-type *mda-type* [**level** *mda-level*]

no mda-type

Context

[Tree] (config>card>mda mda-type)

Full Context

configure card mda mda-type

Description

This mandatory command provisions a specific MDA type to the device configuration for the slot. The MDA can be preprovisioned but an MDA must be provisioned before ports can be configured. Ports can be configured once the MDA is properly provisioned.

A maximum of two MDAs can be provisioned on an IOM or XCM. To modify an MDA slot, shut down all port associations.

XMAs are provisioned using MDA commands. A medium severity alarm is generated if an MDA is inserted that does not match the MDA type configured for the slot. This alarm is cleared when the correct MDA is inserted or the configuration is modified. A high severity alarm is raised when an administratively enabled MDA is removed from the chassis. This alarm is cleared if the either the correct MDA type is inserted or the configuration is modified. A low severity trap is issued if an MDA is removed that is administratively disabled.

An MDA can only be provisioned in a slot if the MDA type is allowed in the MDA slot. An error message is generated when an MDA is provisioned in a slot where it is not allowed.

Some MDA hardware can support two different firmware loads. One load includes the base Ethernet functionality, including 10G WAN mode, but does not include 1588 port-based timestamping. The second load includes the base Ethernet functionality and 1588 port-based timestamping, but does not include 10G WAN mode. These are identified as two MDA types that are the same, except for a "-ptp” suffix to indicate the second loadset; for example, *x40-10gb-sfp* and *x40-10gb-sfp-ptp*. A hard reset of the MDA occurs when switching between the two provisioned types.

A medium severity alarm is generated if an MDA is inserted that does not match the MDA type configured for the slot. This alarm is cleared when the correct MDA is inserted or the configuration is modified.

A high severity alarm is raised when an administratively enabled MDA is removed from the chassis. This alarm is cleared if the either the correct MDA type is inserted or the configuration is modified. A low severity trap is issued if an MDA is removed that is administratively disabled.

An alarm is raised if partial or complete MDA failure is detected. The alarm is cleared when the error condition ceases.

All parameters in the MDA context remain and if non-default values are required then their configuration remains as it is on all existing MDAs.

New generations of XMAs include variants controlled through hardware and software licensing. For these XMAs, the license level must be provisioned in addition to the MDA type. An XMA cannot become operational unless the provisioned license level matches the license level of the XMA installed into the slot. The set of license levels varies by MDA type.

The provisioned level controls aspects related to connector provisioning and the consumption of hardware egress queues and egress policers. Changes to the provisioned license level may be blocked if configuration that would not be permitted with the new target license level exists.

If the license level is not specified, the level is set to the highest license level for that XMA.

The **no** form of this command deletes the MDA from the configuration. The MDA must be administratively shut down before it can be deleted from the configuration.

Parameters

mda-type

Specifies the type of MDA selected for the slot position. Values for this attribute vary by platform and release. The release notes include a listing of all supported mda-types and their CLI strings. In addition, the command can be queried to check which mda-types are relevant for the active platform type. Some examples include me6-10gb-spf+ and x4-100g-cfp2.

mda-level

Specifies the MDA level. Possible values vary by MDA type.

Platforms

7705 SAR Gen 2

18.7 meas-interval

meas-interval

Syntax

meas-interval {**5-mins** | **15-mins** | **1-hour** | **1-day**} [**create**]

no meas-interval {**5-mins** | **15-mins** | **1-hour** | **1-day**}

Context

[Tree] (config>oam-pm>session meas-interval)

Full Context

configure oam-pm session meas-interval

Description

This command establishes the parameters of the individual measurement intervals utilized by the session. Multiple measurement intervals may be specified within the session. A maximum of three different measurement intervals may be configured under each session.

The **no** form of this command deletes the specified measurement interval.

Parameters

meas-interval

Specifies the duration of the measurement interval.

Values 1-min, 5-mins, 15-mins, 1-hour, 1-day

create

Creates the measurement interval.

Platforms

7705 SAR Gen 2

18.8 med-out

med-out

Syntax

med-out {*number* | **igp-cost**}

no med-out

Context

[Tree] (config>service>vprn>bgp>group med-out)

[Tree] (config>service>vprn>bgp med-out)

[Tree] (config>service>vprn>bgp>group>neighbor med-out)

Full Context

configure service vprn bgp group med-out

configure service vprn bgp med-out

configure service vprn bgp group neighbor med-out

Description

This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set.

The specified value can be overridden by any value set via a route policy.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command used at the global level reverts to default where the MED is not advertised.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no med-out

Parameters

number

Specifies the MED path attribute value, expressed as a decimal integer.

Values 0 to 4294967295

igp-cost

Specifies the MED is set to the IGP cost of the given IP prefix.

Platforms

7705 SAR Gen 2

med-out

Syntax

med-out {*number* | **igp-cost**}

no med-out

Context

[Tree] (config>router>bgp>group>neighbor med-out)

[Tree] (config>router>bgp med-out)

[Tree] (config>router>bgp>group med-out)

Full Context

configure router bgp group neighbor med-out

configure router bgp med-out

configure router bgp group med-out

Description

This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set.

The specified value can be overridden by any value set via a route policy.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command used at the global level reverts to default where the MED is not advertised.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no med-out

Parameters

number

Specifies the MED path attribute value, expressed as a decimal integer.

Values 0 to 4294967295

igp-cost

Sets MED to the IGP cost of the given IP prefix.

Platforms

7705 SAR Gen 2

18.9 member

member

Syntax

[**no**] **member** *interface-name*

Context

[Tree] (config>service>vprn>isis>link-group>level member)

Full Context

configure service vprn isis link-group level member

Description

This command adds or removes a links to the associated link-group. The interface name should already exist before it is added to a link-group.

The **no** form of this command removes the specified interface from the associated link-group.

Parameters

interface-name

Specifies the name of the interface to be added or removed from the associated link-group.

Platforms

7705 SAR Gen 2

member

Syntax

member *user-profile-name* [*user-profile-name*]

no member *user-profile-name*

Context

[Tree] (config>system>security>user>console member)

Full Context

configure system security user console member

Description

This command is used to associate the user with a local command authorization profile.

A user can be associated with up to eight profiles.

When a user is a member of multiple profiles, profiles are evaluated in the order that they are configured. Evaluation stops if there is a match, or when the default action of the a profile is **deny-all**, **permit-all** or **read-only-all**. When the profile default action is **none** and if no match conditions are met in the profile, the next profile is evaluated. When the default action of the last profile is **none** and no explicit match is found, the command is denied.

The **no** form of this command removes the association between the user and the profile.

Default

member default

Parameters

user-profile-name

Specifies up to eight user profile names, up to 32 characters.

Platforms

7705 SAR Gen 2

member

Syntax

[**no**] **member** *interface-name*

Context

[Tree] (config>router>isis>link-group>level member)

Full Context

configure router isis link-group level member

Description

This command adds or removes a link to the associated link-group. The interface name should already exist before it is added to a link-group.

The **no** form of this command removes the specified interface from the associated link-group.

Parameters

interface-name

Specifies the name of the interface to be added or removed from the associated link-group.

Platforms

7705 SAR Gen 2

18.10 member-pool

member-pool

Syntax

member-pool *name*

no member-pool

Context

[Tree] (config>isa>tunnel-grp member-pool)

Full Context

configure isa tunnel-group member-pool

Description

This command associates the tunnel group with a tunnel member pool. This tunnel group is used as the designated standby in an N:M IPsec redundancy configuration.

The **no** form of this command removes the tunnel member pool from the configuration.

Parameters

name

Specifies the name of the member pool, up to 32 characters.

Platforms

7705 SAR Gen 2

18.11 members

members

Syntax

[**no**] **members** *comm-id* [*comm-id*]

Context

[Tree] (config>router>policy-options>community members)

Full Context

configure router policy-options community members

Description

This command adds members to a route policy community list to use in route policy entries.

Each member of a community list is a standard, extended, or large community value or a regular expression that potentially matches many community values. A regular expression incorporates terms and operators that use the terms. An individual numerical digit is an elementary term in the community regular expression. More complex terms can be built from elementary terms. The following are key operators supported by SR OS:

- .
- *
- ?
- {n}
- {m,n}
- {m, }

To reverse the match criteria when specifying a list of ranges or single values using square brackets, use the non-match operator (^) before the elements within the square brackets.

The **no** version of this command deletes route policy community members.

Parameters

comm-id

Specifies a BGP community value, up to 72 characters. A community ID can be specified in different forms.

Values [*as-num*:*comm-val* | *reg-ex* | *ext-comm* | *well-known-comm* | *large-comm*]

where:

- *as-num* — 0 to 65535
- *comm-val* — 0 to 65535
- *reg-ex* — A regular expression string. Allowed values are any string up to 72 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (such as "#", "$", or spaces), the entire string must be enclosed within double quotes.
- *ext-comm* — The extended community, defined as one of the following:
  - {*target* | *origin*}:*ip-address*:*comm-val*
  - {*target* | *origin*}:*reg-ex1*&*reg-ex2*
  - {*target* | *origin*}:*ip-address*:*reg-ex2*
  - {*target* | *origin*}:*asnum*:*ext-comm-val*
  - {*target* | *origin*}:*ext-asnum*:*comm-val*
  - **bandwidth**:*asnum*:*val-in-mbps*
  - **ext:4300**:*ovstate*
  - **ext**:*value1*:*value2*
  - **flowspec-set**:*ext-asnum*:*group-id*
  - **flowspec-set-trans**:*ext-asnum*:*group-id*
  - **ipv6-redirect**: *ipv6-addr*
  - **color**:*co-bits*:*color-value*

  where:

  - *target* — route target
  - *origin* — route origin
  - *ip-address* — a.b.c.d
  - *ext-comm-val* — 0 to 4294967295
  - *ext-asnum* — 0 to 4294967295
  - *reg-ex1*, *reg-ex2* — A regular expression string. Allowed values are any string up to 63 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.
  - **bandwidth** — bandwidth
  - *val-in-mbps* — 0 to 16777215

- **ext:4300** — origin verification
- *ovstate* — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid)
- **ext** — extended
- *value1* — 0000 to FFFF
- *value2* — 0 to FFFFFFFFFFFF
- **flowspec-set —** FlowSpec set
- **flowspec-set-trans —** FlowSpec set transitive
- *ipv6-addr* — x:x:x:x:x:x:x:x (eight 16-bit pieces)
- *group-id* — 0 to 16383
- *co-bits* — 00, 01, 10 or 11
- *color-value* — 0 to 4294967295

- *well-known-comm* — **null**, **no-export**, **no-export-subconfed**, **no-advertise**, **llgr-stale**, **no-llgr**, **blackhole**
- *large-comm* — large community, defined as one of the following:
  - *ext-asnum*:*ext-comm-val*:*ext-comm-val*
  - *reg-ex3*&*reg-ex3*&*reg-ex3*

  where:
  - *reg-ex3* — A regular expression string. Allowed values are any string up to 68 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

18.12 memory-use-alarm

memory-use-alarm

Syntax

memory-use-alarm rising-threshold *threshold* [**falling-threshold** *threshold*] **interval** *seconds* [*rmon-event-type*] [**startup-alarm** *alarm-type*]

no memory-use-alarm

Context

[Tree] (config>system>thresholds memory-use-alarm)

Full Context

configure system thresholds memory-use-alarm

Description

The memory thresholds are based on monitoring the TIMETRA-SYSTEM-MIB sgiMemoryUsed object. This object contains the amount of memory currently used by the system. The severity level is Alarm. The absolute sample type method is used.

The **no** form of this command removes the configured memory threshold warning.

Parameters

rising-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is greater than or equal to this threshold and the associated startup-alarm is equal to rising or either.

After a rising threshold crossing event is generated, another such event will not be generated until the sampled value falls below this threshold and reaches less than or equal the falling-threshold value.

The threshold value represents units in bytes.

Values -2147483648 to 2147483647

Default 0

falling-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is less than or equal to this threshold and the associated startup-alarm is equal to falling or either.

After a falling threshold crossing event is generated, another such event will not be generated until the sampled value rises above this threshold and reaches greater than or equal the rising-threshold value.

The threshold value represents units in bytes.

Values -2147483648 to 2147483647

Default 0

seconds

Specifies the polling period over which the data is sampled and compared with the rising and falling thresholds.

The threshold value represents units in bytes.

Values 1 to 2147483647

rmon-event-type

Specifies the type of notification action to be taken when this event occurs.

Values log — An entry is made in the RMON-MIB log table for each event occurrence. This does not create an OS logger entry. The RMON-MIB log table entries can be viewed using the CLI command.

trap — An SR OS logger event is generated. The SR OS logger utility then distributes the notification of this event to its configured log destinations which may be CONSOLE, telnet session, memory log, cflash file, syslog, or SNMP trap destinations logs.

both — Both an entry in the RMON-MIB logTable and an SR OS logger event are generated.

none — No action is taken.

Default both

alarm-type

Specifies the alarm that may be sent when this alarm is first created. If the first sample is greater than or equal to the rising threshold value and startup-alarm is equal to rising or either, then a single rising threshold crossing event is generated. If the first sample is less than or equal to the falling threshold value and startup-alarm is equal to falling or either, a single falling threshold crossing event is generated.

Values rising, falling, either

Default either

Configuration example

```
memory-use-alarm rising-threshold 50000000 falling-threshold 45999999
interval 500 rmon-event-type both start-alarm either
```

Platforms

7705 SAR Gen 2

18.13 memory-use-warn

memory-use-warn

Syntax

memory-use-warn rising-threshold *threshold* [**falling-threshold** *threshold*] **interval** *seconds* [*rmon-event-type*] [**startup-alarm** *alarm-type*]

no memory-use-warn

Context

[Tree] (config>system>thresholds memory-use-warn)

Full Context

configure system thresholds memory-use-warn

Description

The memory thresholds are based on monitoring MemoryUsed object. This object contains the amount of memory currently used by the system. The severity level is Alarm.

The absolute sample type method is used.

The **no** form of this command removes the configured compact flash threshold warning.

Parameters

rising-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is greater than or equal to this threshold and the associated startup-alarm is equal to rising or either.

After a rising threshold crossing event is generated, another such event will not be generated until the sampled value falls below this threshold and reaches less than or equal the falling-threshold value.

The threshold value represents units in bytes.

Values -2147483648 to 2147483647

Default 0

falling-threshold *threshold*

Specifies a threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single threshold crossing event will be generated. A single threshold crossing event will also be generated if the first sample taken is less than or equal to this threshold and the associated startup-alarm is equal to falling or either.

After a falling threshold crossing event is generated, another such event will not be generated until the sampled value rises above this threshold and reaches greater than or equal the rising-threshold value.

The threshold value represents units in bytes.

Values -2147483648 to 2147483647

Default 0

seconds

Specifies the polling period over which the data is sampled and compared with the rising and falling thresholds.

Values 1 to 2147483647

rmon-event-type

Specifies the type of notification action to be taken when this event occurs.

Values log — An entry is made in the RMON-MIB log table for each event occurrence.

This does not create an SR OS logger entry. The RMON-MIB log table entries can be viewed using the **show>system>thresholds** CLI command.

trap — An SR OS logger event is generated. The SR OS logger utility then distributes the notification of this event to its configured log destinations which may be CONSOLE, telnet session, memory log, cflash file, syslog, or SNMP trap destinations logs.

both — Both an entry in the RMON-MIB logTable and an SR OS logger event are generated.

none — No action is taken.

Default both

startup-alarm *alarm-type*

Specifies the alarm that may be sent when this alarm is first created. If the first sample is greater than or equal to the rising threshold value and startup-alarm is equal to rising or either, then a single rising threshold crossing event is generated. If the first sample is less than or equal to the falling threshold value and startup-alarm is equal to falling or either, a single falling threshold crossing event is generated.

Default either

Values rising, falling, either

Configuration example

```
memory-use-warn rising-threshold 500000 falling-threshold 400000 interval 800
 rmon-
event-type log start-alarm falling
```

Platforms

7705 SAR Gen 2

18.14 mesh-group

mesh-group

Syntax

mesh-group {**value** | **blocked**}

no mesh-group

Context

[Tree] (config>service>vprn>isis>if mesh-group)

Full Context

configure service vprn isis interface mesh-group

Description

This command assigns an interface to a mesh group. Mesh groups limit the amount of flooding that occurs when a new or changed LSP is advertised throughout an area.

All routers in a mesh group should be fully meshed. When LSPs need to be flooded, only a single copy is received rather than a copy per neighbor.

To create a mesh group, configure the same mesh group value for each interface that is part of the mesh group. All routers must have the same mesh group value configured for all interfaces that are part of the mesh group.

To prevent an interface from flooding LSPs, the optional **blocked** parameter can be specified. Configure mesh groups carefully. It is easy to create isolated islands that do not receive updates as (other) links fail.

The **no** form of this command removes the interface from the mesh group. The interface does not belong to a mesh group.

Default

no mesh-group

Parameters

value

Specifies a unique decimal integer value distinguishes this mesh group from other mesh groups on this or any other router that is part of this mesh group.

Values 1 to 2000000000

blocked

Prevents an interface from flooding LSPs.

Platforms

7705 SAR Gen 2

mesh-group

Syntax

mesh-group {**value** | **blocked**}

no mesh-group

Context

[Tree] (config>router>isis>interface mesh-group)

Full Context

configure router isis interface mesh-group

Description

This command assigns an interface to a mesh group. Mesh groups limit the amount of flooding that occurs when a new or changed LSP is advertised throughout an area.

All routers in a mesh group should be fully meshed. When LSPs need to be flooded, only a single copy is received rather than a copy per neighbor.

To create a mesh group, configure the same mesh group value for each interface that is part of the mesh group. All routers must have the same mesh group value configured for all interfaces that are part of the mesh group.

To prevent an interface from flooding LSPs, the optional **blocked** parameter can be specified. Configure mesh groups carefully to avoid creating isolated islands that do not receive updates as (other) links fail.

The **no** form of this command removes the interface from the mesh group.

Parameters

value

Specifies the unique decimal integer value that distinguishes this mesh group from other mesh groups on this or any other router that is part of this mesh group.

Values 1 to 2000000000

blocked

Prevents an interface from flooding LSPs.

Platforms

7705 SAR Gen 2

18.15 mesh-sdp

mesh-sdp

Syntax

mesh-sdp *sdp-id*[*:vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**]

mesh-sdp *sdp-id*[*:vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **leaf-ac**

mesh-sdp *sdp-id*[*:vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **root-leaf-tag**

no mesh-sdp *sdp-id*[*:vc-id*]

Context

[Tree] (config>service>vpls mesh-sdp)

Full Context

configure service vpls mesh-sdp

Description

This command binds a VPLS service to an existing service destination point (SDP).

Mesh SDPs bound to a service are logically treated like a single bridge "port" for flooded traffic where flooded traffic received on any mesh SDP on the service is replicated to other "ports" (spoke SDPs and SAPs) and not transmitted on any mesh SDPs.

Note:
This command creates a binding between a service and an SDP. The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service is down.

The SDP must already be defined in the **config>service>sdp** context in order to associate the SDP with an Epipe or VPLS service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end router devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Default

No *sdp-id* is bound to a service.

Parameters

sdp-id

Specifies an SDP identifier.

Values 1 to 17407

vc-id

Specifies a virtual circuit identifier. This value is used to validate the VC ID portion of each mesh SDP binding defined in the service. The default value of this object is equal to the service ID. Any value may be used for the *vc-id* when there is no existing mesh SDP within the service; if a mesh SDP exists then all other mesh SDPs in the service must be configured with the same *vc-id*.

Values 1 to 4294967295

vc-type

Specifies to override the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the binding's VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

The VC type value for Ethernet is 0x0005.

The VC type value for an Ethernet VLAN is 0x0004.

ether

Defines the VC type as Ethernet. The **vlan** keyword is mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke SDP binding (hex 5).

vlan

Defines the VC type as VLAN. The top VLAN tag, if a VLAN tag is present, is stripped from traffic received on the pseudowire, and a vlan-tag is inserted when forwarding into the pseudowire. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for mesh SDP bindings.

Note:
The system expects a symmetrical configuration with its peer, specifically it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. As some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. Consequently, protocol extractions will not necessarily function for asymmetrical configurations as they would with a symmetrical configurations resulting in an unexpected operation.

root-leaf-tag

Specifies a tagging mesh SDP under an E-Tree VPLS. When a tag SDP binding is required, it is created with a root-leaf-tag flag. Only VLAN tag SDP bindings are supported. The VLAN type must be set to VC VLAN type. The root-leaf-tag parameter indicates this SDP binding is a tag SDP that uses a default VID 1 for root and 2 for leaf. The SDP binding tags egress E-Tree traffic with root and leaf VIDs as appropriate. Root and leaf VIDs are

only significant between peering VPLS but the values must be consistent on each end. On ingress a tag SDP binding removes the VID tag on the interface between VPLS in the same E-Tree service. The tag SDP receives root tagged traffic and marks the traffic with a root indication internally. This option is not available on BGP EVPN-enabled E-Tree services.

leaf-ac

Specifies an access (AC) mesh SDP binding under a E-Tree VPLS as a leaf access (AC) SDP. The default E-Tree SDP type is a root AC if *leaf-ac* or *root-leaf-tag* is not specified at SDP binding creation. This option is only available when the VPLS is designated as an E-Tree VPLS. BGP EVPN-enabled E-Tree VPLS services support the **leaf-ac** option.

create

Keyword used to create the mesh SDP. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

18.16 mesh-sdp-binding

mesh-sdp-binding

Syntax

[**no**] **mesh-sdp-binding**

Context

[Tree] (config>service>vpls>site mesh-sdp-binding)

Full Context

configure service vpls site mesh-sdp-binding

Description

This command enables applications to all mesh SDPs.

The **no** form of reverts the default.

Default

no mesh-sdp-binding

Platforms

7705 SAR Gen 2

18.17 message

message

Syntax

message {**eq** | **neq**} **pattern** *pattern* [**regexp**]

no message

Context

[Tree] (config>service>vprn>log>filter>entry>match message)

Full Context

configure service vprn log filter entry match message

Description

This command adds system messages as a match criterion.

The **no** form of this command removes messages as a match criterion.

Parameters

eq

Specifies if the matching criteria should be equal to the specified value.

neq

Specifies if the matching criteria should not be equal to the specified value.

pattern

Specifies a message, up to 400 characters, to be used in the match criteria.

regexp

Specifies the type of string comparison to use to determine if the log event matches the value of **message** command parameters. When the **regexp** keyword is not specified, the default matching algorithm used is a basic substring match.

Platforms

7705 SAR Gen 2

message

Syntax

message {**eq** | **neq**} **pattern** *pattern* [**regexp**]

no message

Context

[Tree] (config>log>filter>entry>match message)

Full Context

configure log filter entry match message

Description

This command adds system messages as a match criterion.

The **no** form of this command removes messages as a match criterion.

Parameters

eq

Determines if the matching criteria should be equal to the specified value.

neq

Determines if the matching criteria should not be equal to the specified value.

pattern

Specifies a message up to 400 characters in length to be used in the match criteria.

regexp

Specifies the type of string comparison to use to determine if the log event matches the value of **message** command parameters. When the **regexp** keyword is not specified, the default matching algorithm used is a basic substring match.

Platforms

7705 SAR Gen 2

18.18 message-count

message-count

Syntax

message-count *count*

Context

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>icmp6-gen>pkt-too-big message-count)

[Tree] (config>service>vprn>if>sap>ipsec-tun>icmp6-gen>pkt-too-big message-count)

[Tree] (config>ipsec>tnl-temp>icmp6-gen>pkt-too-big message-count)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>icmp6-gen>pkt-too-big message-count)

Full Context

configure service ies interface ipsec ipsec-tunnel icmp6-generation pkt-too-big message-count

configure service vprn interface sap ipsec-tunnel icmp6-generation pkt-too-big message-count

configure ipsec tunnel-template icmp6-generation pkt-too-big message-count

configure service vprn interface ipsec ipsec-tunnel icmp6-generation pkt-too-big message-count

Description

This command configures the maximum number of ICMPv6 messages that can be sent during the configured interval.

Parameters

count

Specifies the maximum number of ICMPv6 messages that can be sent during the configured interval

Values 10 to 1000

Default 100

Platforms

7705 SAR Gen 2

message-count

Syntax

message-count *number*

Context

[Tree] (config>router>if>ipsec>ipsec-tunnel>icmp-gen>frag-required message-count)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>icmp-gen>frag-required message-count)

[Tree] (config>ipsec>tnl-temp>icmp-gen>frag-required message-count)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>icmp-gen>frag-required message-count)

[Tree] (config>service>vprn>if>sap>ip-tunnel>icmp-gen>frag-required message-count)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel>icmp-gen>frag-required message-count)

Full Context

configure router interface ipsec ipsec-tunnel icmp-generation frag-required message-count

configure service vprn interface ipsec ipsec-tunnel icmp-generation frag-required message-count

configure ipsec tunnel-template icmp-generation frag-required message-count

configure service ies interface ipsec ipsec-tunnel icmp-generation frag-required message-count

configure service vprn interface sap ip-tunnel icmp-generation frag-required message-count

configure service vprn interface sap ipsec-tunnel icmp-generation frag-required message-count

Description

This command configures the maximum number of ICMP Destination Unreachable "fragmentation needed and DF set" messages (type 3, code 4) that can be sent during the period specified by the **interval** *seconds* command.

Default

message-count 100

Parameters

number

Specifies the number of ICMP Destination Unreachable "fragmentation needed and DF set" messages that are transmitted within the **interval** *seconds* command time.

Values 10 to 1000

Platforms

7705 SAR Gen 2

18.19 message-digest-key

message-digest-key

Syntax

message-digest-key *keyid* **md5** [*key* | *hash-key*] [**hash** | **hash2** | **custom**]

no message-digest-key *keyid*

Context

[Tree] (config>service>vprn>ospf>area>sham-link message-digest-key)

[Tree] (config>service>vprn>ospf>area>if message-digest-key)

[Tree] (config>service>vprn>ospf>area>virtual-link message-digest-key)

Full Context

configure service vprn ospf area sham-link message-digest-key

configure service vprn ospf area interface message-digest-key

configure service vprn ospf area virtual-link message-digest-key

Description

This command configures a message digest key when MD5 authentication is enabled on the interface, virtual-link or sham-link. Multiple message digest keys can be configured.

This command is not valid in the OSPF3 context.

The **no** form of this command removes the message digest key identified by the *key-id*.

Default

No message digest keys are defined.

Parameters

keyid

Specifies the key ID. The *keyid* is expressed as a decimal integer.

Values 1 to 255

md5 *key*

Specifies the MD5 key. The *key* can be any alphanumeric string up to 16 characters in length.

md5 hash-key

Specifies the MD5 hash key. The key can be any combination of ASCII characters up to 32 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

message-digest-key

Syntax

message-digest-key *key-id* **md5** [*key* | *hash-key* | *hash2-key*] [**hash** | **hash2** | **custom**]

no message-digest-key *key-id*

Context

[Tree] (config>router>ospf>area>interface message-digest-key)

[Tree] (config>router>ospf>area>virtual-link message-digest-key)

Full Context

configure router ospf area interface message-digest-key

configure router ospf area virtual-link message-digest-key

Description

This command configures a message digest key when MD5 authentication is enabled on the interface. Multiple message digest keys can be configured.

The **no** form of this command removes the message digest key identified by the *key-id*.

By default, no message keys are defined.

Parameters

key-id

Specifies the key ID. The *keyid* is expressed as a decimal integer.

Values 1 to 255

key

Specifies the MD5 key. The *key* can be any alphanumeric string up to 16 characters in length.

hash-key* | *hash2-key

Specifies the MD5 hash or hash2 key. the hash key. The key can be any combination of ASCII characters up to 32 (hash1-key) or 55 (hash2-key) characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

18.20 message-fast-tx

message-fast-tx

Syntax

message-fast-tx *time*

no message-fast-tx

Context

[Tree] (config>system>lldp message-fast-tx)

Full Context

configure system lldp message-fast-tx

Description

This command configures the duration of the fast transmission period.

Default

no message-fast-tx

Parameters

time

Specifies the fast transmission period in seconds.

Values 1 to 3600

Default 1

Platforms

7705 SAR Gen 2

18.21 message-fast-tx-init

message-fast-tx-init

Syntax

message-fast-tx-init *count*

no message-fast-tx-init

Context

[Tree] (config>system>lldp message-fast-tx-init)

Full Context

configure system lldp message-fast-tx-init

Description

This command configures the number of LLDPDUs to send during the fast transmission period.

Default

no message-fast-tx-init

Parameters

count

Specifies the number of LLDPDUs to send during the fast transmission period.

Values 1 to 8

Default 4

Platforms

7705 SAR Gen 2

18.22 message-interval

message-interval

Syntax

message-interval {[*seconds*] [**milliseconds** *milliseconds*]}

no message-interval

Context

[Tree] (config>service>ies>if>ipv6>vrrp message-interval)

Full Context

configure service ies interface ipv6 vrrp message-interval

Description

This command sets the advertisement timer and indirectly sets the master down timer on the virtual router instance. The message-interval setting must be the same for all virtual routers participating as a virtual router. Any VRRP advertisement message received with an Advertisement Interval field different than the virtual router instance configured message-interval value will be silently discarded.

The message-interval command is available in both non-owner and owner **vrrp** *virtual-router-id* nodal contexts. If the message-interval command is not executed, the default message interval of 1 second will be used.

The **no** form of this command restores the default message interval value of 1 second to the virtual router instance.

Default

message-interval 1

Parameters

seconds

The number of seconds that will transpire before the advertisement timer expires.

Values 1 to 255

Default 1

milliseconds ***milliseconds***

Specifies the time interval, in milliseconds, between sending advertisement messages.

Values 100 to 900

Platforms

7705 SAR Gen 2

message-interval

Syntax

message-interval {[*seconds*] [**milliseconds** *milliseconds*]}

no message-interval

Context

[Tree] (config>service>ies>if>vrrp message-interval)

Full Context

configure service ies interface vrrp message-interval

Description

This command sets the advertisement timer and indirectly sets the master down timer on the virtual router instance. The message-interval setting must be the same for all virtual routers participating as a virtual router. Any VRRP advertisement message received with an Advertisement Interval field different than the virtual router instance configured message-interval value will be silently discarded.

The message-interval command is available in both non-owner and owner **vrrp** *virtual-router-id* nodal contexts. If the message-interval command is not executed, the default message interval of 1 second will be used.

The **no** form of this command restores the default message interval value of 1 second to the virtual router instance.

Parameters

seconds

The number of seconds that will transpire before the advertisement timer expires.

Values 1 to 255

Default 1

milliseconds ***milliseconds***

Specifies the time interval, in milliseconds, between sending advertisement messages. This parameter is not supported on non-redundant chassis.

Values 100 to 900

Platforms

7705 SAR Gen 2

message-interval

Syntax

message-interval {[*seconds*] [**milliseconds** *milliseconds*]}

no message-interval

Context

[Tree] (config>service>vprn>if>ipv6>vrrp message-interval)

Full Context

configure service vprn interface ipv6 vrrp message-interval

Description

This command sets the advertisement timer and indirectly sets the master down timer on the virtual router instance. The message-interval setting must be the same for all virtual routers participating as a virtual router. Any VRRP advertisement message received with an Advertisement Interval field different than the virtual router instance configured message-interval value will be silently discarded.

The message-interval command is available in both non-owner and owner **vrrp** *virtual-router-id* nodal contexts. If the message-interval command is not executed, the default message interval of 1 second will be used.

The **no** form of this command restores the default message interval value of 1 second to the virtual router instance.

Parameters

seconds

The number of seconds that will transpire before the advertisement timer expires.

Values 1 to 255

Default 1

milliseconds *milliseconds*

Specifies the time interval, in milliseconds, between sending advertisement messages. This parameter is not supported on single-slot chassis.

Values 100 to 900

Platforms

7705 SAR Gen 2

message-interval

Syntax

message-interval {[*seconds*] [**milliseconds** *milliseconds*]}

no message-interval

Context

[Tree] (config>router>if>vrrp message-interval)

[Tree] (config>router>if>ipv6>vrrp message-interval)

Full Context

configure router interface vrrp message-interval

configure router interface ipv6 vrrp message-interval

Description

This command configures the administrative advertisement message timer used by the master virtual router instance to send VRRP advertisement messages and to derive the master down timer as backup.

For an owner virtual router instance, the administrative advertisement timer directly sets the operational advertisement timer and indirectly sets the master down timer for the virtual router instance.

Non-owner virtual router instances usage of the **message-interval** setting is dependent on the state of the virtual router (master or backup) and the state of the **master-int-inherit** parameter.

- When a non-owner is operating as master for the virtual router, the configured **message-interval** is used as the operational advertisement timer similar to an owner virtual router instance. The **master-int-inherit** command has no effect when operating as master.
- When a non-owner is in the backup state with **master-int-inherit** disabled, the configured **message-interval** value is used to match the incoming VRRP advertisement message advertisement interval

field. If the locally configured message interval does not match the advertisement interval field, the VRRP advertisement is discarded.

- When a non-owner is in the backup state with **master-int-inherit** enabled, the configured **message-interval** is ignored. The master down timer is indirectly derived from the incoming VRRP advertisement message advertisement interval field value.

VRRP advertisements messages that are fragmented, or contain IP options (IPv4), or contain extension headers (IPv6) require a longer message interval to be configured.

The in-use value of the message interval is used to derive the master down timer to be used when the virtual router is operating in backup mode based on the following formula:

(3x (in-use message interval) + skew time)

The skew time portion is used to slow down virtual routers with relatively low priority values when competing in the master election process.

The command is available in both non-owner and owner **vrrp** nodal contexts.

By default, a **message-interval** of 1 second is used.

The **no** form of the command reverts to the default value.

Default

message-interval 1 — Advertisement timer set to 1 second.

Parameters

seconds

The number of seconds that will transpire before the advertisement timer expires expressed as a decimal integer.

Values IPv4: 1 to 255 IPv6: 1 to 40

milliseconds

Specifies the time interval, in milliseconds, between sending advertisement messages.

Values 100 to 900 IPv6: 10 to 990

Platforms

7705 SAR Gen 2

18.23 message-length

message-length

Syntax

message-length *message-length*

no message-length

Context

[Tree] (config>service>sdp>keep-alive message-length)

Full Context

configure service sdp keep-alive message-length

Description

This command configures the SDP monitoring keepalive request message length transmitted.

The **no** form of this command reverts the **message-length** *octets* value to the default setting.

Default

no message-length — The message length should be equal to the SDP's operating path MTU as configured in the **config>service>sdp path-mtu** command. If the default size is overridden, the actual size used will be the smaller of the operational SDP ID Path MTU and the size specified.

Parameters

message-length

Specifies the size of the keepalive request messages in octets, expressed as a decimal integer. The **size** keyword overrides the default keepalive message size.

Values 40 to 9198

Platforms

7705 SAR Gen 2

18.24 message-severity-level

message-severity-level

Syntax

message-severity-level

Context

[Tree] (config>system>management-interface>cli>md-cli>environment message-severity-level)

Full Context

configure system management-interface cli md-cli environment message-severity-level

Description

This command configures the message severity level.

Platforms

7705 SAR Gen 2

18.25 message-size

message-size

Syntax

message-size *max-num-of-routes*

no message-size

Context

[Tree] (config>service>vprn>ripng>group>neighbor message-size)

[Tree] (config>service>vprn>ripng>group message-size)

[Tree] (config>service>vprn>rip>group message-size)

[Tree] (config>service>vprn>rip message-size)

[Tree] (config>service>vprn>rip>group>neighbor message-size)

[Tree] (config>service>vprn>ripng message-size)

Full Context

configure service vprn ripng group neighbor message-size

configure service vprn ripng group message-size

configure service vprn rip group message-size

configure service vprn rip message-size

configure service vprn rip group neighbor message-size

configure service vprn ripng message-size

Description

This command sets the maximum number of routes per RIP update message.

The **no** form of this command resets the maximum number of routes back to the default of 25.

Default

no message-size

Parameters

size

An Integer.

Values 25 to 255

Default 25

Platforms

7705 SAR Gen 2

message-size

Syntax

message-size *max-num-of-routes*

no message-size

Context

[Tree] (config>router>rip message-size)

[Tree] (config>router>rip>group message-size)

[Tree] (config>router>ripng>group>neighbor message-size)

[Tree] (config>router>ripng message-size)

[Tree] (config>router>rip>group>neighbor message-size)

[Tree] (config>router>ripng>group message-size)

Full Context

configure router rip message-size

configure router rip group message-size

configure router ripng group neighbor message-size

configure router ripng message-size

configure router rip group neighbor message-size

configure router ripng group message-size

Description

This command configures the maximum number of routes per RIP update message.

The **no** form of the command reverts to the default value.

Default

message-size 25

Parameters

max-num-of-routes

The maximum number of RIP routes per RIP update message expressed as a decimal integer.

Values 25 to 255

Platforms

7705 SAR Gen 2

18.26 messages

messages

Syntax

[**no**] **messages**

Context

[Tree] (debug>router>ldp>if>event messages)

[Tree] (debug>router>ldp>peer>event messages)

Full Context

debug router ldp interface event messages

debug router ldp peer event messages

Description

This command displays specific information (for example, message type, source, and destination) regarding LDP messages sent to and received from LDP peers.

The **no** form of the command disables debugging output for LDP messages.

Platforms

7705 SAR Gen 2

18.27 metric

metric

Syntax

metric *metric-value*

no metric [*metric-value*]

Context

[Tree] (config>service>vprn>static-route-entry>ipsec-tunnel metric)

[Tree] (config>service>vprn>static-route-entry>black-hole metric)

[Tree] (config>service>vprn>static-route-entry>next-hop metric)

[Tree] (config>service>vprn>static-route-entry>grt metric)

[Tree] (config>service>vprn>static-route-entry>indirect metric)

Full Context

configure service vprn static-route-entry ipsec-tunnel metric

configure service vprn static-route-entry black-hole metric

configure service vprn static-route-entry next-hop metric

configure service vprn static-route-entry grt metric

configure service vprn static-route-entry indirect metric

Description

This command specifies the cost metric for the static route, expressed as a decimal integer. This value is used when importing the static route into other protocols such as OSPF. When the metric is configured as 0 then the metric configured in OSPF, default-import-metric, applies. When modifying the metric of an existing static route, the preference will not change unless specified. This value is also used to determine which static route to install in the forwarding table.

If there are multiple static routes with the same preference but different metrics then the lower cost (metric) route will be installed.

The **no** form of this command returns the metric to the default value

Default

metric 1

Parameters

metric-value

Specifies the cost metric value.

Values 0 to 65535

Platforms

7705 SAR Gen 2

metric

Syntax

metric *metric*

no metric

Context

[Tree] (config>service>vprn>isis>if>level metric)

Full Context

configure service vprn isis interface level metric

Description

This command configures the metric used for the level on the interface.

In order to calculate the lowest cost to reach a given destination, each configured level on each interface must have a cost. The costs for each level on an interface may be different.

If the metric is not configured, the default of 10 is used unless reference bandwidth is configured.

The **no** form of this command reverts to the default value.

Default

no metric

Parameters

metric

The metric assigned for this level on this interface.

Values 1 to 16777215

Default 10

Platforms

7705 SAR Gen 2

metric

Syntax

metric *metric*

no metric

Context

[Tree] (config>service>vprn>ospf3>area>if metric)

[Tree] (config>service>vprn>ospf>area>sham-link metric)

[Tree] (config>service>vprn>ospf>area>if metric)

Full Context

configure service vprn ospf3 area interface metric

configure service vprn ospf area sham-link metric

configure service vprn ospf area interface metric

Description

This command configures an explicit route cost metric for the OSPF interface that overrides the metrics calculated based on the speed of the underlying link.

The **no** form of this command deletes the manually configured interface metric, so the interface uses the computed metric based on the **reference-bandwidth** command setting and the speed of the underlying link.

Default

no metric

Parameters

metric

The metric to be applied to the interface expressed as a decimal integer.

Values 1 to 65535

Platforms

7705 SAR Gen 2

metric

Syntax

metric *metric*

no metric

Context

[Tree] (config>router>mpls>static-lsp metric)

[Tree] (config>router>mpls>lsp-template metric)

[Tree] (config>router>mpls>lsp metric)

Full Context

configure router mpls static-lsp metric

configure router mpls lsp-template metric

configure router mpls lsp metric

Description

This command allows the user to override the LSP operational metric with a constant administrative value that will not change regardless of the actual path the LSP is using over its lifetime.

The LSP operational metric will match the metric the active path of this LSP is using at any given time. For a CSPF LSP, this metric represents the cumulative IGP metric of all the links the active path is using. If CSPF for this LSP is configured to use the TE metric, the LSP operational metric is set to the maximum value. For a non-CSPF LSP, the operational metric is the shortest IGP cost to the destination of the LSP.

The LSP operational metric is used by some applications to select an LSP among a set of LSPs that are destined to the same egress router. The LSP with the lowest operational metric will be selected. If more than one LSP with the same lowest LSP metric exists, the LSP with the lowest tunnel index will be selected. The configuration of a constant metric by the user will make sure the LSP always maintains its preference in this selection regardless of the path it is using at any given time. Applications that use the LSP operational metric include LDP-over-RSVP, VPRN auto-bind, and IGP, BGP and static route shortcuts.

The **no** form of this command disables the administrative LSP metric and reverts to the default setting in which the metric value will represent the LSP metric returned by MPLS. The same behavior is obtained if the user entered a metric of value zero (0).

Default

no metric. The LSP operational metric defaults to the metric returned by MPLS.

Parameters

metric

Specifies the integer value which specifies the value of the LSP administrative metric. A value of zero command reverts to the default setting and disables the administrative LSP metric.

Values 0 to 16777215

Platforms

7705 SAR Gen 2

metric

Syntax

metric *metric*

no metric [*metric*]

Context

[Tree] (config>router>static-route-entry>black-hole metric)

[Tree] (config>router>static-route-entry>indirect metric)

[Tree] (config>router>static-route-entry>next-hop metric)

Full Context

configure router static-route-entry black-hole metric

configure router static-route-entry indirect metric

configure router static-route-entry next-hop metric

Description

This command specifies the cost metric for the static route, expressed as a decimal integer. This value is used when importing the static route into other protocols such as OSPF. When the metric is configured

as 0 then the metric configured in OSPF, default-import-metric, applies. When modifying the metric of an existing static route, the preference will not change unless specified. This value is also used to determine which static route to install in the forwarding table.

If there are multiple static routes with the same preference but different metrics then the lower cost (metric) route will be installed.

The **no** form of this command returns the metric to the default value

Default

metric 1

Parameters

metric

Specifies the cost metric value.

Values 0 to 65535

Platforms

7705 SAR Gen 2

metric

Syntax

metric *metric*

no metric

Context

[Tree] (config>service>sdp metric)

Full Context

configure service sdp metric

Description

This command specifies the metric to be used within the tunnel table manager for decision making purposes. When multiple SDPs going to the same destination exist, this value is used as a tie-breaker by tunnel table manager users such as MP-BGP to select the route with the lower value.

Parameters

metric

Specifies the SDP metric.

Values 0 to 65535

Platforms

7705 SAR Gen 2

metric

Syntax

metric *metric*

no metric

Context

[Tree] (config>router>isis>if>level metric)

Full Context

configure router isis interface level metric

Description

This command configures the metric used for the level on the interface.

In order to calculate the lowest cost to reach a given destination, each configured level on each interface must have a cost. The costs for each level on an interface may be different.

If the metric is not configured, the default of 10 is used unless reference bandwidth is configured.

The **no** form of this command reverts to the default value.

Default

metric 10

Parameters

metric

Specifies the metric assigned for this level on this interface.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

metric

Syntax

metric *metric*

no metric

Context

[Tree] (config>router>ospf>area>interface metric)

[Tree] (config>router>ospf3>area>interface metric)

Full Context

configure router ospf area interface metric

configure router ospf3 area interface metric

Description

This command configures an explicit route cost metric for the OSPF interface that overrides the metrics calculated based on the speed of the underlying link.

The **no** form of this command deletes the manually configured interface metric, so the interface uses the computed metric based on the **reference-bandwidth** command setting and the speed of the underlying link.

Default

no metric

Parameters

metric

Specifies the metric to be applied to the interface expressed as a decimal integer.

Values 1 to 65535

Platforms

7705 SAR Gen 2

metric

Syntax

metric *metric* [**equal** | **or-higher** | **or-lower**]

no metric

Context

[Tree] (config>router>policy-options>policy-statement>entry>from metric)

Full Context

configure router policy-options policy-statement entry from metric

Description

This command matches BGP routes based on local preference (the value in the MULTI_EXIT_DISC attribute).

If no comparison qualifiers are present (**equal**, **or-higher**, **or-lower**), then **equal** is the implied default.

A non-BGP route does not match a policy entry if it contains the **metric** command. In addition, a BGP route without a MED attribute also does not match a policy entry if it contains a **metric** command.

Default

no metric

Parameters

metric

Specifies the MED value.

Values 0 to 4294967295, or a parameter name delimited by starting and ending at-sign (@) characters

equal

Specifies that matched routes should have the same MED as the value specified.

or-higher

Specifies that matched routes should have the same or a greater MED as the value specified.

or-lower

Specifies that matched routes should have the same or a lower MED as the value specified.

Platforms

7705 SAR Gen 2

metric

Syntax

metric {**add** | **subtract**} *metric*

metric set [**igp** | *metric-value*]

no metric

Context

[Tree] (config>router>policy-options>policy-statement>entry>action metric)

[Tree] (config>router>policy-options>policy-statement>default-action metric)

Full Context

configure router policy-options policy-statement entry action metric

configure router policy-options policy-statement default-action metric

Description

In a BGP import or export policy, this command assigns a MED value to routes matched by the policy statement entry. The MED value may be set to a fixed value (overriding the received value), set to the routing table cost of the route used to resolve the next hop of the BGP route, or modified by adding or subtracting a fixed value offset.

The **no** form of this command removes the MED attribute from the matched routes.

Default

no metric

Parameters

add

Specifies that an integer is added to any existing metric. If the result of the addition results in a number greater than 4294967295, the value 4294967295 is used.

subtract

Specified *integer* is subtracted from any existing metric. If the result of the subtraction results in a number less than 0, the value of 0 is used.

set

Specifies that *integer* replaces any existing metric.

igp

Sets the MED value to the routing table cost of the route used to resolve the next hop of the BGP route.

metric

Specifies the metric modifier expressed as a decimal integer.

Values 0 to 4294967295

param-name —Specifies the metric parameter variable name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@”

Platforms

7705 SAR Gen 2

metric

Syntax

metric *metric*

Context

[Tree] (config>ipsec>tnl-temp>rev-route metric)

Full Context

configure ipsec tunnel-template reverse-route metric

Description

This command configures the metric for reverse routes. The system uses the metric when selecting a route to install in the route table.

Default

metric 0

Parameters

metric

Specifies the metric value for reverse routes.

Values 0 to 65535

Platforms

7705 SAR Gen 2

18.28 metric-in

metric-in

Syntax

metric-in *metric*

no metric-in

Context

[Tree] (config>service>vprn>ripng metric-in)

[Tree] (config>service>vprn>rip metric-in)

[Tree] (config>service>vprn>ripng>group>neighbor metric-in)

[Tree] (config>service>vprn>rip>group>neighbor metric-in)

[Tree] (config>service>vprn>rip>group metric-in)

[Tree] (config>service>vprn>ripng>group metric-in)

Full Context

configure service vprn ripng metric-in

configure service vprn rip metric-in

configure service vprn ripng group neighbor metric-in

configure service vprn rip group neighbor metric-in

configure service vprn rip group metric-in

configure service vprn ripng group metric-in

Description

This command sets the metric added to routes received from a RIP neighbor.

The **no** form of this command reverts the *metric* value back to the default.

Default

no metric-in

Parameters

metric

The value added to the metric of routes received from a RIP neighbor, expressed as a decimal integer.

Values 1 to 16

Platforms

7705 SAR Gen 2

metric-in

Syntax

metric-in *metric*

no metric-in

Context

[Tree] (config>router>ripng>group metric-in)

[Tree] (config>router>ripng metric-in)

[Tree] (config>router>rip metric-in)

[Tree] (config>router>rip>group metric-in)

[Tree] (config>router>rip>group>neighbor metric-in)

[Tree] (config>router>ripng>group>neighbor metric-in)

Full Context

configure router ripng group metric-in

configure router ripng metric-in

configure router rip metric-in

configure router rip group metric-in

configure router rip group neighbor metric-in

configure router ripng group neighbor metric-in

Description

This command configures the metric added to routes received from a RIP neighbor.

When applying an export policy to a RIP configuration, the policy overrides the metric values determined through calculations involving the **metric-in** and **metric-out** values.

The **no** form of the command reverts to the default value.

Default

metric-in 1

Parameters

metric

Specifies the value added to the metric of routes received from a RIP neighbor expressed as a decimal integer.

Values 1 to 16

Platforms

7705 SAR Gen 2

18.29 metric-out

metric-out

Syntax

metric-out *metric*

no metric-out

Context

[Tree] (config>service>vprn>rip>group metric-out)

[Tree] (config>service>vprn>rip>group>neighbor metric-out)

[Tree] (config>service>vprn>ripng>group metric-out)

[Tree] (config>service>vprn>ripng>group>neighbor metric-out)

[Tree] (config>service>vprn>rip metric-out)

[Tree] (config>service>vprn>ripng metric-out)

Full Context

configure service vprn rip group metric-out

configure service vprn rip group neighbor metric-out

configure service vprn ripng group metric-out

configure service vprn ripng group neighbor metric-out

configure service vprn rip metric-out

configure service vprn ripng metric-out

Description

This command sets the metric added to routes exported into RIP and advertised to RIP neighbors.

The **no** form of this command removes the command from the config and resets the metric-in value back to the default.

Default

no metric-out

Parameters

metric

The value added to the metric for routes exported into RIP and advertised to RIP neighbors, expressed as a decimal integer.

Values 1 to 16

Platforms

7705 SAR Gen 2

metric-out

Syntax

metric-out *metric*

no metric-out

Context

[Tree] (config>router>ripng>group>neighbor metric-out)

[Tree] (config>router>ripng metric-out)

[Tree] (config>router>rip>group>neighbor metric-out)

[Tree] (config>router>rip>group metric-out)

[Tree] (config>router>ripng>group metric-out)

[Tree] (config>router>rip metric-out)

Full Context

configure router ripng group neighbor metric-out

configure router ripng metric-out

configure router rip group neighbor metric-out

configure router rip group metric-out

configure router ripng group metric-out

configure router rip metric-out

Description

This command configures the metric assigned to routes exported into RIP and advertised to RIP neighbors.

When applying an export policy to a RIP configuration, the policy overrides the metric values determined through calculations involving the **metric-in** and **metric-out** values.

The **no** form of the command reverts to the default value.

Default

metric-out 1

Parameters

metric

Specifies the value added to the metric for routes exported into RIP and advertised to RIP neighbors expressed as a decimal integer.

Values 1 to 16

Platforms

7705 SAR Gen 2

18.30 metric-type

metric-type

Syntax

metric-type *metric-type*

no metric-type

Context

[Tree] (config>router>mpls>lsp-template metric-type)

[Tree] (config>router>mpls>lsp metric-type)

Full Context

configure router mpls lsp-template metric-type

configure router mpls lsp metric-type

Description

This command configures the link metric type used by the local CSPF or the PCE controller in the SR-TE LSP path computation.

The **no** form of this command returns the link metric type to its default value.

Default

metric-type igp

Parameters

metric-type

Specifies the metric type for the LSP.

Values **igp** — Specifies use of the IGP metric.

te — Specifies use of the traffic-engineering metric. This is the default metric type.

delay — Specifies computation delay metrics.

Platforms

7705 SAR Gen 2

metric-type

Syntax

metric-type {**igp** | **te-metric** | **delay**}

no metric-type

Context

[Tree] (config>router>fad>flex-algo metric-type)

Full Context

configure router flexible-algorithm-definitions flex-algo metric-type

Description

This command configures the type of metric for the flexible algorithm. The topology graph assumes that all links are configured with the correct metric type.

For example, if the flexible algorithm definition instructs the use of **te-metric** keyword, it is assumed that all links have *te-metric* configured. Links without the *te-metric* configuration are excluded from the flexible algorithm topology graph.

The **no** form of this command removes the configured metric type and sets it to its default value.

Default

metric-type igp

Parameters

igp

Keyword to use the IGP metric for the flexible algorithm topology graph.

te-metric

Keyword to use the TE metric for the flexible algorithm topology graph.

delay

Keyword to use the delay metric for the flexible algorithm topology graph.

Platforms

7705 SAR Gen 2

18.31 mfib-allowed-mda-destinations

mfib-allowed-mda-destinations

Syntax

mfib-allowed-mda-destinations

Context

[Tree] (config>service>vpls>mesh-sdp>egress mfib-allowed-mda-destinations)

[Tree] (config>service>vpls>spoke-sdp>egress mfib-allowed-mda-destinations)

Full Context

configure service vpls mesh-sdp egress mfib-allowed-mda-destinations

configure service vpls spoke-sdp egress mfib-allowed-mda-destinations

Description

Commands in this context configure MFIB-allowed MDA destinations.

The allowed-mda-destinations node and the corresponding **mda** command are used on spoke and mesh SDP bindings to provide a list of MDA destinations in the chassis that are allowed as destinations for multicast streams represented by [*,g] and [s,g] multicast flooding records on the VPLS service. The MDA list only applies to IP multicast forwarding when IGMP snooping is enabled on the VPLS service. The MDA list has no effect on normal VPLS flooding such as broadcast, L2 multicast, unknown destinations or non-snooped IP multicast.

At the IGMP snooping level, a spoke or mesh SDP binding is included in the flooding domain for an IP multicast stream when it has either been defined as a multicast router port, received a IGMP query through the binding or has been associated with the multicast stream through an IGMP request by a host over the binding. Due to the dynamic nature of the way that a spoke or mesh SDP binding is associated with one or more egress network IP interfaces, the system treats the binding as appearing on all network ports. This causes all possible network destinations in the switch fabric to be included in the multicast streams flooding domain. The MDA destination list provides a simple mechanism that narrows the IP multicast switch fabric destinations for the spoke or mesh SDP binding.

If no MDAs are defined within the allowed-mda-destinations node, the system operates normally and will forward IP multicast flooded packets associated with the spoke or mesh SDP binding to all switch fabric taps containing network IP interfaces.

The MDA inclusion list should include all MDAs that the SDP binding may attempt to forward through. A simple way to ensure that an MDA that is not included in the list is not being used by the binding is to define the SDP the binding is associated with as MPLS and use an RSVP-TE LSP with a strict egress hop. The MDA associated with the IP interface defined as the strict egress hop should be present in the inclusion list. If the inclusion list does not currently contain the MDA that the binding is forwarding through, the multicast packets will not reach the destination represented by the binding.

By default, the MDA inclusion list is empty.

If an MDA is removed from the list, the MDA is automatically removed from the flooding domain of any snooped IP multicast streams associated with a destination on the MDA unless the MDA was the last MDA on the inclusion list. Once the inclusion list is empty, all MDAs are eligible for snooped IP multicast flooding for streams associated with the SDP binding.

Platforms

7705 SAR Gen 2

mfib-allowed-mda-destinations

Syntax

mfib-allowed-mda-destinations

Context

[Tree] (config>service>pw-template>egress mfib-allowed-mda-destinations)

Full Context

configure service pw-template egress mfib-allowed-mda-destinations

Description

Commands in this context configure MFIB-allowed XMA or MDA destinations.

The allowed-mda-destinations node and the corresponding **mda** command are used on spoke and mesh SDP bindings to provide a list of XMA or MDA destinations in the chassis that are allowed as destinations for multicast streams represented by [*,g] and [s,g] multicast flooding records on the VPLS service. The XMA or MDA list only applies to IP multicast forwarding when IGMP snooping is enabled on the VPLS service. The XMA or MDA list has no effect on normal VPLS flooding such as broadcast, Layer 2 multicast, unknown destinations or non-snooped IP multicast.

At the IGMP snooping level, a spoke or mesh SDP binding is included in the flooding domain for an IP multicast stream when it has either been defined as a multicast router port, received a IGMP query through the binding or has been associated with the multicast stream through an IGMP request by a host over the binding. Due to the dynamic nature of the way that a spoke or mesh SDP binding is associated with one or more egress network IP interfaces, the system treats the binding as appearing on all network ports. This causes all possible network destinations in the switch fabric to be included in the multicast streams flooding domain. The XMA or MDA destination list provides a simple mechanism that narrows the IP multicast switch fabric destinations for the spoke or mesh SDP binding.

If no XMAs or MDAs are defined within the allowed-mda-destinations node, the system operates normally and will forward IP multicast flooded packets associated with the spoke or mesh SDP binding to all switch fabric taps containing network IP interfaces.

The XMA or MDA inclusion list should include all XMAs or MDAs that the SDP binding may attempt to forward through. A simple way to ensure that an XMA or MDA that is not included in the list is not being used by the binding is to define the SDP the binding is associated with as MPLS and use an RSVP-TE LSP with a strict egress hop. The XMA or MDA associated with the IP interface defined as the strict egress hop should be present in the inclusion list.

If the inclusion list does not currently contain the XMA or MDA that the binding is forwarding through, the multicast packets will not reach the destination represented by the binding. By default, the XMA or MDA inclusion list is empty.

If an XMA or MDA is removed from the list, the XMA or MDA is automatically removed from the flooding domain of any snooped IP multicast streams associated with a destination on the XMA or MDA unless the XMA or MDA was the last XMA or MDA on the inclusion list. Once the inclusion list is empty, all XMAs or MDAs are eligible for snooped IP multicast flooding for streams associated with the SDP binding.

Platforms

7705 SAR Gen 2

18.32 mfib-table-high-wmark

mfib-table-high-wmark

Syntax

[**no**] **mfib-table-high-wmark** *high-water-mark*

Context

[Tree] (config>service>vpls mfib-table-high-wmark)

Full Context

configure service vpls mfib-table-high-wmark

Description

This command specifies the multicast FIB high watermark. When the percentage filling level of the multicast FIB exceeds the configured value, a trap is generated and a log entry is added.

The **no** form of this command reverts to the default.

Default

mfib-table-high-wmark 95

Parameters

high-water-mark

Specifies the multicast FIB high watermark as a percentage.

Values 1 to 100

Platforms

7705 SAR Gen 2

18.33 mfib-table-low-wmark

mfib-table-low-wmark

Syntax

[**no**] **mfib-table-low-wmark** *low-water-mark*

Context

[Tree] (config>service>vpls mfib-table-low-wmark)

Full Context

configure service vpls mfib-table-low-wmark

Description

This command specifies the multicast FIB low watermark. When the percentage filling level of the multicast FIB drops below the configured value, the corresponding trap is cleared and a log entry is added.

The **no** form of this command reverts to the default.

Default

mfib-table-low-wmark

Parameters

low-water-mark

Specifies the multicast FIB low watermark as a percentage.

Values 1 to 100

Default 90

Platforms

7705 SAR Gen 2

18.34 mfib-table-size

mfib-table-size

Syntax

mfib-table-size *size*

no mfib-table-size

Context

[Tree] (config>service>vpls mfib-table-size)

Full Context

configure service vpls mfib-table-size

Description

This command specifies the maximum number of (s,g) entries in the multicast forwarding database (MFIB) for this VPLS instance.

The *mfib-table-size* parameter specifies the maximum number of multicast database entries for both learned and static multicast addresses for the VPLS instance. When a table-size limit is set on the mfib of a service which is lower than the current number of dynamic entries present in the mfib then the number of entries remains above the limit.

The **no** form of this command removes the configured maximum MFIB table size.

Default

no mfib-table-size

Parameters

size

Specifies the maximum number of (s,g) entries allowed in the Multicast FIB

Values 1 to 40959

Platforms

7705 SAR Gen 2

18.35 mh-mode

mh-mode

Syntax

mh-mode {**access** | **network**}

Context

[Tree] (config>service>vpls>bgp-evpn>mpls mh-mode)

Full Context

configure service vpls bgp-evpn mpls mh-mode

Description

This command configures each BGP-EVPN instance in a multi-instance EVPN VPLS service to behave as network or access.

You can only configure one network instance for the service. If the service has a provider tunnel enabled, it requires a network instance.

Default

mh-mode network

Parameters

access

Specifies that the BGP-EVPN instance does not participate in multihoming procedures, such as DF election processing or local bias forwarding.

network

Specifies that the BGP-EVPN instance participates in multihoming procedures, such as DF election processing or local bias forwarding.

In services with two instances, only one of the two instances can be configured as **network**.

Platforms

7705 SAR Gen 2

mh-mode

Syntax

mh-mode {**access** | **network**}

Context

[Tree] (config>service>epipe>bgp-evpn>mpls mh-mode)

Full Context

configure service epipe bgp-evpn mpls mh-mode

Description

This command configures each BGP-EVPN instance in a multi-instance Epipe service to behave as network or access.

You can only configure one network instance for the service. If the service has a provider tunnel enabled, it requires a network instance.

Default

mh-mode network

Parameters

access

Specifies that the BGP-EVPN instance does not participate in multihoming procedures, such as DF election processing or local bias forwarding.

network

Specifies that the BGP-EVPN instance participates in multihoming procedures, such as DF election processing or local bias forwarding.

In services with two instances, only one of the two instances can be configured as **network**.

Platforms

7705 SAR Gen 2

18.36 micro-loop-avoidance

micro-loop-avoidance

Syntax

micro-loop-avoidance [**fib-delay** *fib-delay*]

no micro-loop-avoidance

Context

[Tree] (config>router>isis>segment-routing micro-loop-avoidance)

Full Context

configure router isis segment-routing micro-loop-avoidance

Description

This command enables, in the IGP instance, the microloop avoidance feature to prevent microloops from using segment routing (SR) loop-free tunnels for packets that are forwarded over SR IS-IS node SIDs.

This command enables microloop avoidance for MT0. Microloop avoidance for MT2 is enabled when this command is enabled along with SR-MPLS MT2 using the **configure router isis segment-routing multi-topology mt2** command.

When enabled, the behavior of the feature is triggered by the receipt of a single event on a P2P link or broadcast link with two neighbors:

- link addition or restoration
- link removal or failure
- link metric change

IGP then performs the following procedures:

1. IGP runs the main SPF and LFA SPFs.
2. For a node or a prefix in which the SPF resulted in no change to its next hops and metrics, IGP takes no action.
3. For a node or a prefix in which SPF resulted in a change to its next hops or metrics, IGP marks the route as eligible for microloop avoidance.
   a. Activate, for each node SID that uses a microloop avoidance eligible route with ECMP next hops, the common set of next hops between the previous and new SPF.
   b. Compute and activate, for each node SID which uses a microloop avoidance eligible route, with a single next hop loop-free SR tunnel that is applicable to the specific link event.

      This tunnel acts the microloop avoidance primary path for the route and uses the same outgoing interface as the new computed primary next hop.
   c. Program the TI-LFA, base LFA, or remote LFA backup path that protects the new primary next hop for the node SID.
4. Start the **fib-delay** timer to delay programming of new main and LFA SPF results into the FIB.
5. After the expiry of the **fib-delay** timer, program the new primary next hops for node SIDs routes that were marked eligible for microloop avoidance procedures.

The **no** form of this command disables the microloop avoidance feature.

Default

no micro-loop-avoidance

Parameters

fib-delay

Specifies the delay, in 100s of milliseconds, before the system programs the new next hops for the SR tunnel.

Values 1 to 300

Default 15

Platforms

7705 SAR Gen 2

micro-loop-avoidance

Syntax

[no] micro-loop-avoidance

Context

[Tree] (config>router>isis>flex-algos>flex-algo micro-loop-avoidance)

Full Context

configure router isis flexible-algorithms flex-algo micro-loop-avoidance

Description

This command enables flexible algorithms-aware microloop avoidance. When enabled, the microloop configuration parameters are inherited from the base SPF.

This command enables microloop avoidance with flexible algorithms for MT0. Microloop avoidance with flexible algorithms for MT2 is enabled when this command is enabled along with SR-MPLS MT2 using the **configure router isis segment-routing multi-topology mt2** command.

The **no** form of this command disables the microloop avoidance for flexible algorithms.

Default

no micro-loop-avoidance

Platforms

7705 SAR Gen 2

micro-loop-avoidance

Syntax

[**no**] **micro-loop-avoidance**

Context

[Tree] (config>router>ospf>flex-algos>flex-algo micro-loop-avoidance)

Full Context

configure router ospf flexible-algorithms flex-algo micro-loop-avoidance

Description

This command enables microloop avoidance for an SR-OSPF flexible algorithm, and consequently, inherits the FIB delay timer from the SR-OSPF **configure router ospf segment-routing** context.

When enabled, FIB updates are delayed before programming new primary next hops to avoid microloops.

When enabled, the feature applies to the following contexts:

- OSPFv2 SR-OSPF IPv4 tunnel (node SID)
- IPv4 and IPv6 SR-TE LSPs that use a node SID in their segment list
- IPv4 and IPv6 SR policies that use a node SID in their segment list

The **no** form of this command disables microloop avoidance. When microloop avoidance is disabled, the system forces any running FIB delay to expire immediately and programs the new next hops for all impacted node SIDs. When disabled, microloop avoidance is disabled instantaneously and will be disabled for the next SPF runs. Microloop avoidance remains disabled until it is re-enabled.

Default

no micro-loop-avoidance

Platforms

7705 SAR Gen 2

micro-loop-avoidance

Syntax

micro-loop-avoidance [**fib-delay** *fib-delay*]

no micro-loop-avoidance

Context

[Tree] (config>router>ospf>segm-rtng micro-loop-avoidance)

Full Context

configure router ospf segment-routing micro-loop-avoidance

Description

This command enables microloop avoidance for SR-OSPF.

When enabled, the feature applies to the following contexts:

- SPFv2 SR-OSPF IPv4 tunnel (node SID)
- IPv4 and IPv6 SR-TE LSPs that use a node SID in their segment list
- IPv4 and IPv6 SR policies that use a node SID in their segment list

The **no** form of this command disables microloop avoidance. When microloop avoidance is disabled, the system forces any running FIB delay to expire immediately and programs the new next hops for all impacted node SIDs. When disabled, microloop avoidance is disabled instantaneously and will be disabled for the next SPF runs. Microloop avoidance remains disabled until it is re-enabled.

Default

no micro-loop-avoidance

Parameters

fib-delay

Specifies the FIB delay before programming new primary next hops. Configure the FIB delay timer to a value that corresponds to the worst-case IGP convergence in a network domain. The default FIB delay timer value corresponds to a network with a nominal convergence time.

Values 1 to 300

Default 15 (equal to 1.5 seconds or 15 units of 100 ms)

Platforms

7705 SAR Gen 2

18.37 min-advertisement-interval

min-advertisement-interval

Syntax

[**no**] **min-advertisement-interval** *seconds*

Context

[Tree] (config>router>router-advert>if min-advertisement-interval)

Full Context

configure router router-advertisement interface min-advertisement-interval

Description

This command configures the minimum interval between sending ICMPv6 neighbor discovery router advertisement messages.

Default

min-advertisement-interval 200

Parameters

seconds

Specifies the minimum interval in seconds between sending ICMPv6 neighbor discovery router advertisement messages.

Values 3 to 1350

Platforms

7705 SAR Gen 2

18.38 min-delay

min-delay

Syntax

min-delay [*delay*]

no min-delay

Context

[Tree] (config>log>event-handling>handler>action-list>entry min-delay)

Full Context

configure log event-handling handler action-list entry min-delay

Description

This command specifies the minimum delay in seconds between subsequent executions of the action specified in this entry. This is useful, for example, to ensure that a script does not get triggered to execute too often.

Default

no min-delay

Parameters

delay

Specifies the unit in seconds.

Values 1 to 604800

Platforms

7705 SAR Gen 2

18.39 min-frame-length

min-frame-length

Syntax

min-frame-length *byte-length*

Context

[Tree] (config>port>ethernet min-frame-length)

Full Context

configure port ethernet min-frame-length

Description

This command configures the minimum transmitted frame length.

Note: The *byte-length* value of 72 is only supported on FP4 and newer generations of XMA, MDA-e-XP, and MDA-s.

Parameters

byte-length

Specifies the number of bytes for the minimum frame length.

Values 64, 68, 72

Default 64

Platforms

7705 SAR Gen 2

18.40 min-lease-time

min-lease-time

Syntax

min-lease-time [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no min-lease-time

Context

[Tree] (config>service>vprn>dhcp>server>pool min-lease-time)

[Tree] (config>router>dhcp>server>pool min-lease-time)

Full Context

configure service vprn dhcp local-dhcp-server pool min-lease-time

configure router dhcp local-dhcp-server pool min-lease-time

Description

This command configures the minimum lease time.

The **no** form of this command reverts to the default.

Default

min-lease-time min 10

Parameters

min-lease-time

Specifies the minimum lease time.

Values		
	days	0 to 3650
	hours	0 to 23
	minutes	0 to 59
	seconds	0 to 59

Platforms

7705 SAR Gen 2

18.41 min-route-advertisement

min-route-advertisement

Syntax

min-route-advertisement *seconds*

no min-route-advertisement

Context

[Tree] (config>service>vprn>bgp>group>neighbor min-route-advertisement)

[Tree] (config>service>vprn>bgp>group min-route-advertisement)

[Tree] (config>service>vprn>bgp min-route-advertisement)

Full Context

configure service vprn bgp group neighbor min-route-advertisement

configure service vprn bgp group min-route-advertisement

configure service vprn bgp min-route-advertisement

Description

This command configures the minimum interval, in seconds, at which a prefix can be advertised to a peer.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command reverts to default values.

Default

min-route-advertisement 30

Parameters

seconds

Specifies the minimum route advertising interval, in seconds, expressed as a decimal integer.

Values 1 to 255

Platforms

7705 SAR Gen 2

min-route-advertisement

Syntax

min-route-advertisement *seconds*

no min-route-advertisement

Context

[Tree] (config>router>bgp>group>neighbor min-route-advertisement)

[Tree] (config>router>bgp min-route-advertisement)

[Tree] (config>router>bgp>group min-route-advertisement)

Full Context

configure router bgp group neighbor min-route-advertisement

configure router bgp min-route-advertisement

configure router bgp group min-route-advertisement

Description

This command configures the minimum interval, in seconds, between successive updates of a prefix towards a peer.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group), or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command used at the global level reverts to default.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

The **rapid-update** command can be used to override the peer-level **min-route-advertisement** time and applies the minimum setting (0 seconds) to routes belonging to address families specified by the **rapid-update** command; routes of other address families continue to be advertised according to the session-level MRAI setting.

The **rapid-update** and **rapid-withdrawal** commands may result in the routes being sent before the peer-level MRAI timer expires.

Default

min-route-advertisement 30

Parameters

seconds

Specifies the minimum route advertising interval, in seconds, expressed as a decimal integer.

Values 1 to 255

Platforms

7705 SAR Gen 2

18.42 min-threshold-separation

min-threshold-separation

Syntax

min-threshold-separation *size* [**bytes** | **kilobytes**]

no min-threshold-separation

Context

[Tree] (config>card>fp>ingress>network>queue-group>policer-control-override>priority-mbs-thresholds min-threshold-separation)

[Tree] (config>card>fp>ingress>access>queue-group>policer-control-override>priority-mbs-thresholds min-threshold-separation)

Full Context

configure card fp ingress network queue-group policer-control-override priority-mbs-thresholds min-threshold-separation

configure card fp ingress access queue-group policer-control-override priority-mbs-thresholds min-threshold-separation

Description

This command defines the minimum required separation between each in-use discard threshold maintained for each parent policer context associated with the policer-control-policy. The min-thresh-separation value may be overridden on each SAP or sub-profile to which the policy is applied.

The system uses the default or specified min-thresh-separation value in order to determine the minimum separation required between each of the of the parent policer discard thresholds. The system enforces the minimum separation based on the following behavior in two ways. The first is determining the size of the shared-portion for each priority level (when the **mbs-contribution** command's optional fixed keyword is not specified):

- When a parent policer instance's priority level has less than two child policers associated, the shared-portion for the level will be zero.
- When a parent policer instance's priority level has two or more child policers associated, the shared-portion for the level will be equal to the current value of **min-thresh-separation**.

The second function the system uses the **min-thresh-separation** value for is determining the value per priority level for the fair-portion:

- When a parent policer instance's priority level has no child policers associated, the fair-portion for the level will be zero.
- When a parent policer instance's priority level has one child policer associated, the fair-portion will be equal to the maximum of the min-thresh-separation value and the priority level's mbs-contribution value.
- When a parent policer instance's priority level has two or more child policers associated, the fair-portion will be equal to the maximum of the following:
  - **min-thresh-separation** value
  - The priority level's **mbs-contribution** value less **min-thresh-separation** value

When the **mbs-contribution** command's optional fixed keyword is defined for a priority level within the policy, the system will treat the defined **mbs-contribution** value as an explicit definition of the priority level's MBS. While the system will continue to track child policer associations with the parent policer priority levels, the association counters will have no effect. Instead the following rules will be used to determine a fixed priority level's shared-portion and fair-portion:

- If a fixed priority level's **mbs-contribution** value is set to zero, both the shared-portion and fair-portion will be set to zero
- If the **mbs-contribution** value is not set to zero:
  - The shared-portion will be set to the current **min-thresh-separation** value
  - The fair-portion will be set to the maximum of the following:
    - **min-thresh-separation** value
    - **mbs-contribution** value less **min-thresh-separation value**

Each time the **min-thresh-separation** value is modified, the thresholds for all instances of the parent policer created through association with this **policer-control-policy** are reevaluated except for parent policer instances that currently have a min-thresh-separation override.

Determining the Correct Value for the Minimum Threshold Separation Value

The minimum value for **min-thresh-separation** should be set equal to the maximum size packet that will be handled by the parent policer. This ensures that when a lower priority packet is incrementing the bucket, the size of the increment will not cause the bucket's depth to equal or exceed a higher priority threshold. It

also ensures that an unfair packet within a priority level cannot cause the PIR bucket to increment to the discard-all threshold within the priority.

When evaluating maximum packet size, each child policer's per-packet-offset setting should be taken into consideration. If the maximum size packet is 1518 bytes and a per-packet-offset parameter is configured to add 20 bytes per packet, min-thresh-separation should be set to 1538 due to the fact that the parent policer will increment its PIR bucket using the extra 20 bytes.

In most circumstances, a value larger than the maximum packet size is not necessary. Management of priority level aggregate burst tolerance is intended to be implemented using the priority level **mbs-contribution** command. Setting a value larger than the maximum packet size will not adversely affect the policer performance, but it may increase the aggregate burst tolerance for each priority level.

One thing to note is that a priority level's shared-portion of the parent policer's PIR bucket depth is only necessary to provide some separation between a lower priority's discard-all threshold and this priority's discard-unfair threshold. It is expected that the burst tolerance for the unfair packets is relatively minimal since the child policers feeding the parent policer priority level all have some amount of fair burst before entering into an FIR exceed or unfair state. The fair burst amount for a priority level is defined using the mbs-contribution command.

The **no** form of this command returns the policy's **min-thresh-separation** value to the default value. This has no effect on instances of the parent policer where **min-thresh-separation** is overridden unless the override is removed.

Default

no min-thresh-separation

Parameters

size

Specifies that the size parameter is required when executing the **min-thresh-separation** command. It is expressed as an integer and specifies the shared portion in bytes or kilobytes that is selected by the trailing bytes or kilobytes keywords. If both bytes and kilobytes are missing, kilobytes is the assumed value. Setting this value has no effect on parent policer instances where the **min-thresh-separation** value has been overridden. Clearing an override on parent policer instance causes this value to be enforced.

Values 0 to 16777216

bytes | kilobytes

Specifies that the **bytes** keyword is optional and is mutually exclusive with the **kilobytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in bytes.

The **kilobytes** keyword is optional and is mutually exclusive with the **bytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in kilobytes.

Values bytes or kilobytes

Default kilobytes

Platforms

7705 SAR Gen 2

min-thresh-separation

Syntax

min-thresh-separation *size* [**bytes** | **kilobytes**]

Context

[Tree] (config>service>epipe>sap>egress>policy-ctrl-over>mbs-thrshlds min-thresh-separation)

[Tree] (config>service>epipe>sap>ingress>policy-ctrl-over>mbs-thrshlds min-thresh-separation)

Full Context

configure service epipe sap egress policer-control-override priority-mbs-thresholds min-thresh-separation

configure service epipe sap ingress policer-control-override priority-mbs-thresholds min-thresh-separation

Description

This command, within the SAP ingress and egress contexts, is used to override the root arbiter's parent policer min-thresh-separation parameter that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy min-thresh-separation parameter have no effect on the SAP's parent policer until the override is removed using the no min-thresh-separation command within the SAP.

The no form of this command removes the override and allows the min-thresh-separation setting from the policer-control-policy to control the root arbiter's parent policer's minimum discard threshold separation size.

Default

no min-thresh-separation

Parameters

size

The minimum discard threshold separation override value.

Values 1 to 16777216 | default

bytes

Signifies that *size* is expressed in bytes. The bytes and kilobytes keywords are mutually exclusive and optional. The default is kilobytes.

kilobytes

Signifies that *size* is expressed in kilobytes. The bytes and kilobytes keywords are mutually exclusive and optional. The default is kilobytes.

Platforms

7705 SAR Gen 2

min-thresh-separation

Syntax

min-thresh-separation *size* [{**bytes** | **kilobytes**}]

Context

[Tree] (config>service>vpls>sap>egress>policer-ctrl-over>mbs-thrshlds min-thresh-separation)

[Tree] (config>service>vpls>sap>ingress>policer-ctrl-over>mbs-thrshlds min-thresh-separation)

Full Context

configure service vpls sap egress policer-control-override priority-mbs-thresholds min-thresh-separation

configure service vpls sap ingress policer-control-override priority-mbs-thresholds min-thresh-separation

Description

This command within the SAP ingress and egress contexts is used to override the root arbiter's parent policer min-thresh-separation parameter that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy min-thresh-separation parameter have no effect on the SAP's parent policer until the override is removed using the **no min-thresh-separation** command within the SAP.

The **no** form of this command removes the override and allows the min-thresh-separation setting from the policer-control-policy to control the root arbiter's parent policer's minimum discard threshold separation size.

Default

no min-thresh-separation

Parameters

size

This parameter is required when specifying min-thresh-separation override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 16777216 or default

Default kilobytes

Platforms

7705 SAR Gen 2

min-thresh-separation

Syntax

min-thresh-separation *size* [{**bytes** | **kilobytes**}]

Context

[Tree] (config>service>ies>if>sap>ingress>policer-ctrl-over>mbs-thrshlds min-thresh-separation)

[Tree] (config>service>ies>if>sap>egress>policer-ctrl-over>mbs-thrshlds min-thresh-separation)

Full Context

configure service ies interface sap ingress policer-control-override priority-mbs-thresholds min-thresh-separation

configure service ies interface sap egress policer-control-override priority-mbs-thresholds min-thresh-separation

Description

This command within the SAP ingress and egress contexts is used to override the root arbiter's parent policer min-thresh-separation parameter that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy min-thresh-separation parameter have no effect on the SAP's parent policer until the override is removed using the no min-thresh-separation command within the SAP.

The **no** form of this command removes the override and allows the min-thresh-separation setting from the policer-control-policy to control the root arbiter's parent policer's minimum discard threshold separation size.

Default

no min-thresh-separation

Parameters

size

This parameter is required when specifying min-thresh-separation override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 16777216 or default

Default kilobytes

Platforms

7705 SAR Gen 2

min-thresh-separation

Syntax

min-thresh-separation *size* [{**bytes** | **kilobytes**}]

Context

[Tree] (config>service>vprn>if>sap>egress>policy-ctrl-over>mbs-thrshlds min-thresh-separation)

[Tree] (config>service>vprn>if>sap>ingress>policy-ctrl-over>mbs-thrshlds min-thresh-separation)

Full Context

configure service vprn interface sap egress policer-control-override priority-mbs-thresholds min-thresh-separation

configure service vprn interface sap ingress policer-control-override priority-mbs-thresholds min-thresh-separation

Description

This command within the SAP ingress and egress contexts is used to override the root arbiter's parent policer min-thresh-separation parameter that is defined within the policer-control-policy applied to the SAP.

When the override is defined, modifications to the policer-control-policy min-thresh-separation parameter have no effect on the SAP's parent policer until the override is removed using the no min-thresh-separation command within the SAP.

The **no** form of this command removes the override and allows the min-thresh-separation setting from the policer-control-policy to control the root arbiter's parent policer's minimum discard threshold separation size.

Default

no min-thresh-separation

Parameters

size

This parameter is required when specifying min-thresh-separation override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

Values 0 to 16777216 or default

Default kilobytes

Platforms

7705 SAR Gen 2

min-thresh-separation

Syntax

min-thresh-separation *size* [**bytes** | **kilobytes**]

no min-thresh-separation

Context

[Tree] (config>qos>plcr-ctrl-plcy>root>priority-mbs-thresholds min-thresh-separation)

Full Context

configure qos policer-control-policy root priority-mbs-thresholds min-thresh-separation

Description

The **min-thresh-separation** command defines the minimum required separation between each in-use discard threshold maintained for each parent policer context associated with the policer-control-policy. The min-thresh-separation value may be overridden on each SAP or sub-profile to which the policy is applied.

The system uses the default or specified **min-thresh-separation** value in order to determine the minimum separation required between each of the of the parent policer discard thresholds. The system enforces the minimum separation based on the following behavior in two ways. The first is determining the size of the shared-portion for each priority level (when the **mbs-contribution** command's optional fixed keyword is not specified):

- When a parent policer instance's priority level has less than two child policers associated, the shared-portion for the level will be zero.
- When a parent policer instance's priority level has two or more child policers associated, the shared-portion for the level will be equal to the current value of **min-thresh-separation**.

The second function the system uses the **min-thresh-separation** value for is determining the value per priority level for the fair-portion:

- When a parent policer instance's priority level has no child policers associated, the fair-portion for the level will be zero.
- When a parent policer instance's priority level has one child policer associated, the fair-portion will be equal to the maximum of the min-thresh-separation value and the priority level's mbs-contribution value.
- When a parent policer instance's priority level has two or more child policers associated, the fair-portion will be equal to the maximum of the following:
  - **min-thresh-separation** value
  - **mbs-contribution** value less **min-thresh-separation** value

When the **mbs-contribution** command's optional fixed keyword is defined for a priority level within the policy, the system will treat the defined **mbs-contribution** value as an explicit definition of the priority level's MBS. While the system will continue to track child policer associations with the parent policer priority levels, the association counters will have no effect. Instead, the following rules will be used to determine a fixed priority level's shared-portion and fair-portion:

- If a fixed priority level's **mbs-contribution** value is set to zero, both the shared-portion and fair-portion will be set to zero

- If the **mbs-contribution** value is not set to zero:
  - The shared-portion will be set to the current **min-thresh-separation** value
  - The fair-portion will be set to the maximum of the following:

    min-thresh-separation value

    mbs-contribution value less **min-thresh-separation value**

Each time the **min-thresh-separation** value is modified, the thresholds for all instances of the parent policer created through association with this **policer-control-policy** are reevaluated.

Determining the Correct Value for the Minimum Threshold Separation Value

The minimum value for **min-thresh-separation** should be set equal to the maximum size packet that will be handled by the parent policer. This ensures that when a lower priority packet is incrementing the bucket, the size of the increment will not cause the bucket's depth to equal or exceed a higher priority threshold. It also ensures that an unfair packet within a priority level cannot cause the PIR bucket to increment to the discard-all threshold within the priority.

When evaluating maximum packet size, each child policer's per-packet-offset setting should be taken into consideration. If the maximum size packet is 1518 bytes and a per-packet-offset parameter is configured to add 20 bytes per packet, min-thresh-separation should be set to 1538 due to the fact that the parent policer will increment its PIR bucket using the extra 20 bytes.

In most circumstances, a value larger than the maximum packet size is not necessary. Management of priority level aggregate burst tolerance is intended to be implemented using the priority level **mbs-contribution** command. Setting a value larger than the maximum packet size will not adversely affect the policer performance, but it may increase the aggregate burst tolerance for each priority level.

Note:
A priority level's shared-portion of the parent policer's PIR bucket depth is only necessary to provide some separation between a lower priority's discard-all threshold and this priority's discard-unfair threshold. It is expected that the burst tolerance for the unfair packets is relatively minimal since the child policers feeding the parent policer priority level all have some amount of fair burst before entering into an FIR exceed or unfair state. The fair burst amount for a priority level is defined using the mbs-contribution command.

The **no** form of this command returns the policy's **min-thresh-separation** value to the default value.

Parameters

size

The *size* parameter is required when executing the **min-thresh-separation** command. It is expressed as an integer. Setting this value has no effect on parent policer instances where the **min-thresh-separation** value has been overridden.

Values 0 to 16777216 or **default**

Default 1536

bytes | kilobytes

This parameter indicates whether the size is expressed in bytes or kilobytes.

Default **kilobytes**

Platforms

7705 SAR Gen 2

18.43 min-wait-to-advertise

min-wait-to-advertise

Syntax

min-wait-to-advertise *seconds*

no min-wait-to-advertise

Context

[Tree] (config>service>vprn>bgp>convergence min-wait-to-advertise)

Full Context

configure service vprn bgp convergence min-wait-to-advertise

Description

This command configures the minimum amount of time that BGP waits, after the first session establishment following a restart of the BGP instance, until it can start advertising IPv4-unicast and IPv6-unicast routes to its BGP peers, to allow time for re-convergence.

The time limit configured by this command should allow sufficient time for all important peers to re-establish their sessions with the restarting router.

The **no** form of this command implements the default time limit of 0 seconds, which disables all forms of delayed route advertisement. In other words, it causes IPv4-unicast and IPv6-unicast routes to be re-advertised as soon as possible after BGP instance restart.

Default

no min-wait-to-advertise

Parameters

seconds

Specifies the minimum amount of time, in seconds, that BGP waits until IPv4-unicast and IPv6-unicast routes can be advertised to peers.

Values 0 to 3600

Platforms

7705 SAR Gen 2

min-wait-to-advertise

Syntax

min-wait-to-advertise *seconds*

no min-wait-to-advertise

Context

[Tree] (config>router>bgp>convergence min-wait-to-advertise)

Full Context

configure router bgp convergence min-wait-to-advertise

Description

This command configures the minimum amount of time that BGP waits, after the first session establishment following a restart of the BGP instance, until it can start advertising IPv4-unicast and IPv6-unicast routes to its BGP peers, to allow time for re-convergence.

The time limit configured by this command should allow sufficient time for all important peers to re-establish their sessions with the restarting router.

The **no** form of this command implements the default time limit of 0 seconds, which disables all forms of delayed route advertisement. In other words, it causes IPv4-unicast and IPv6-unicast routes to be re-advertised as soon as possible after BGP instance restart.

Default

no min-wait-to-advertise

Parameters

seconds

Specifies the minimum amount of time, in seconds, that BGP waits until IPv4-unicast and IPv6-unicast routes can be advertised to peers.

Values 0 to 3600

Platforms

7705 SAR Gen 2

18.44 minimum

minimum

Syntax

minimum [**percent** [*percent*]] [**number** [*number*]]

no minimum

Context

[Tree] (config>service>vprn>dhcp6>server>pool>prefix>thresholds>minimum-free minimum)

[Tree] (config>router>dhcp6>server>pool>prefix>thresholds>minimum-free minimum)

Full Context

configure service vprn dhcp6 local-dhcp-server pool prefix thresholds minimum-free minimum

configure router dhcp6 local-dhcp-server pool prefix thresholds minimum-free minimum

Description

This command configures a percentage-based or number-based threshold which represents the minimal available percentage or number of the prefix with a configured length in the provisioned prefix. The system sends out a warning if the actual percentage or number is lower than the configured threshold.

For example:

```
prefix 2001:0:0:ffe0::/50 pd wan-host create
   thresholds
       minimum-free prefix-length 64
           minimum number 3
```

With the above configuration, the system sends a warning when the number of available /64 in prefix 2001:0:0:ffe0::/50 is less than 3.

The **no** form of this command removes the command parameters from the configuration.

Parameters

percent

Specifies the percentage of used prefixes with the minimum free threshold length in the pool compared to the number of provisioned prefixes.

Values 0 to 100

number

Specifies the number of prefixes.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

18.45 minimum-age

minimum-age

Syntax

minimum-age [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no minimum-age

Context

[Tree] (config>system>security>password minimum-age)

Full Context

configure system security password minimum-age

Description

Configure the minimum required age of a password before it can be changed again.

Default

minimum-age min 10

Parameters

days

Specifies the minimum required days of a password before it can be changed again.

Values 0 to 1

hours

Specifies the minimum required hours of a password before it can be changed again.

Values 0 to 23

minutes

Specifies the minimum required minutes of a password before it can be changed again.

Values 0 to 59

seconds

Specifies the minimum required seconds of a password before it can be changed again.

Values 0 to 59

Note:
This command applies to local users.

Platforms

7705 SAR Gen 2

18.46 minimum-change

minimum-change

Syntax

minimum-change *distance*

no minimum-change

Context

[Tree] (config>system>security>password minimum-change)

Full Context

configure system security password minimum-change

Description

This command configures the minimum number of characters required to be different in the new password from a previous password.

The **no** form of this command reverts to default value.

Default

minimum-change 5

Parameters

distance

Specifies how many characters must be different in the new password from the old password.

Values 1 to 20

Note:
This command applies to local users.

Platforms

7705 SAR Gen 2

18.47 minimum-classes

minimum-classes

Syntax

minimum-classes *minimum*

no minimum-classes

Context

[Tree] (config>system>security>password>complexity-rules minimum-classes)

Full Context

configure system security password complexity-rules minimum-classes

Description

Force the use of at least this many different character classes

The **no** form of this command resets to default.

Default

no minimum-classes

Parameters

minimum

Specifies the minimum number of classes to be configured.

Values 2 to 4

Platforms

7705 SAR Gen 2

18.48 minimum-delay

minimum-delay

Syntax

minimum-delay [*minimum-delay*]

no minimum-delay

Context

[Tree] (config>router>isis>if>delay-normalization minimum-delay)

Full Context

configure router isis interface delay-normalization minimum-delay

Description

This command configures the minimum delay on the interface. For example, if 5 microseconds is configured, the system advertises this value as the minimum delay, even if TWAMP measures a delay of 2 microseconds.

The **no** form of this command reverts to the default.

Default

1 usec

Parameters

minimum-delay

Specifies the minimum delay in microseconds.

Values 1 to 10000000

Platforms

7705 SAR Gen 2

minimum-delay

Syntax

minimum-delay [*minimum-delay*]

no minimum-delay

Context

[Tree] (config>router>ospf>area>if>delay-normalization minimum-delay)

Full Context

configure router ospf area interface delay-normalization minimum-delay

Description

This command configures the minimum delay on the interface. For example, if 5 microseconds is configured, the system advertises this value as the minimum delay, even if TWAMP measures a delay of 2 microseconds.

The **no** form of this command reverts to the default.

Default

1 usec

Parameters

minimum-delay

Specifies the minimum delay, in microseconds.

Values 1 to 10000000

Platforms

7705 SAR Gen 2

18.49 minimum-free

minimum-free

Syntax

minimum-free *minimum-free* [**percent**] [**event-when-depleted**]

no minimum-free

Context

[Tree] (config>router>dhcp>server>pool minimum-free)

[Tree] (config>service>vprn>dhcp>server>pool minimum-free)

Full Context

configure router dhcp local-dhcp-server pool minimum-free

configure service vprn dhcp local-dhcp-server pool minimum-free

Description

This command specifies the desired minimum number of free addresses in this pool.

The **no** form of this command reverts to the default.

Default

minimum-free 1

Parameters

minimum-free

Specifies the minimum number of free addresses.

Values 0 to 255

percent

Specifies that the value indicates a percentage.

event-when-depleted

This parameter enables a system-generate event when all available addresses in the pool/subnet of local DHCP server are depleted.

Platforms

7705 SAR Gen 2

minimum-free

Syntax

minimum-free *minimum-free* [**percent**] [**event-when-depleted**]

no minimum-free

Context

[Tree] (config>router>dhcp>server>pool>subnet minimum-free)

[Tree] (config>service>vprn>dhcp>server>pool>subnet minimum-free)

Full Context

configure router dhcp local-dhcp-server pool subnet minimum-free

configure service vprn dhcp local-dhcp-server pool subnet minimum-free

Description

This command configures the minimum number of free addresses in this subnet. If the actual number of free addresses in this subnet falls below this configured minimum, a notification is generated.

The **no** form of the reverts to the default.

Default

minimum-free 1

Parameters

minimum-free

Specifies the minimum number of free addresses in this subnet.

Values 0 to 255

percent

Specifies that the value indicates a percentage.

event-when-depleted

Enables a system-generate event when all available addresses in the pool or subnet of local DHCP server are depleted.

Platforms

7705 SAR Gen 2

minimum-free

Syntax

[**no**] **minimum-free prefix-length** [*prefix-length*]

Context

[Tree] (config>service>vprn>dhcp6>server>pool>thresholds minimum-free)

[Tree] (config>router>dhcp6>server>pool>thresholds minimum-free)

Full Context

configure service vprn dhcp6 local-dhcp-server pool thresholds minimum-free

configure router dhcp6 local-dhcp-server pool thresholds minimum-free

Description

This command creates a threshold for a given prefix length on the pool level. Up to 128 thresholds could be created. For example, with **minimum-free prefix-length 64**, then the usage of /64 prefix in the pool is counted.

There are two types of thresholds that could be defined at the pool level:

- Depleted — The system sends out a warning when the prefix with the configured length is no long available in the pool.
- Minimum free — A percentage-based threshold which represents the minimal available percentage of prefix with the configured length in the pool. The system will send out warning if the actual percentage is lower than the configured percentage.

Configuring this command also enables the system stats collection for **configure prefix length**, which could be displayed with the **show router** *router-id* **dhcp6 local-dhcp-server** *dhcp6-server-name* **pool-threshold-stats** command.

The **no** form of this command removes the prefix-length from the configuration.

Parameters

prefix-length

Specifies the IPv6 prefix length.

Values 1 to 128

Platforms

7705 SAR Gen 2

18.50 minimum-length

minimum-length

Syntax

minimum-length *length*

no minimum-length

Context

[Tree] (config>system>security>password>complexity-rules minimum-length)

Full Context

configure system security password complexity-rules minimum-length

Description

This command configures the minimum number of characters required for locally administered passwords, HMAC-MD5-96, HMAC-SHA-96, and des-keys configured in the system security section.

If multiple minimum-length commands are entered each command overwrites the previous entered command.

The **no** form of this command reverts to default value.

Default

minimum-length 6

Parameters

value

Specifies the minimum number of characters required for a password.

Values 6 to 50

Platforms

7705 SAR Gen 2

18.51 minute

minute

Syntax

minute {*minute-number* [..*minute-number*] | **all**}

no minute

Context

[Tree] (config>system>cron>sched minute)

Full Context

configure system cron schedule minute

Description

This command specifies the minute to schedule a command. Multiple minutes of the hour can be specified. When multiple minutes are configured, each of them will cause the schedule to occur. If a minute is configured, but no **hour** or day is configured, the event will not execute. If a minute is configured without configuring the month, weekday, day-of-month, and minute, the event will not execute.

The **no** form of this command removes the specified minute from the configuration.

Default

no minute

Parameters

minute-number

Specifies the minute to schedule a command.

Values 0 to 59 (maximum 60 minute-numbers)

all

Specifies all minutes.

Platforms

7705 SAR Gen 2

18.52 minutes

minutes

Syntax

minutes {*minutes* | **disable**}

no minutes

Context

[Tree] (config>system>security>ssh>key-re-exchange>server minutes)

[Tree] (config>system>security>ssh>key-re-exchange>client minutes)

Full Context

configure system security ssh key-re-exchange server minutes

configure system security ssh key-re-exchange client minutes

Description

This command configures the maximum time, in minutes, before a key re-exchange is initiated by the server.

The **no** form of this command reverts to the default value.

Default

minutes 60

Parameters

minutes

Specifies the time interval, in minutes, after which the SSH client will initiate the key-re-exchange.

Values 1 to 1440

Default 60

disable

Specifies that a session will never timeout. To re-enable **minutes**, enter the command without the **disable** option.

Platforms

7705 SAR Gen 2

18.53 mirror-dest

mirror-dest

Syntax

mirror-dest *service-id* [**create**] [**type** *mirror-type*] [**name** *name*]

no mirror-dest *service-id*

Context

[Tree] (config>mirror mirror-dest)

Full Context

configure mirror mirror-dest

Description

This command creates a context to set up a service that is intended for packet mirroring. It is configured as a service to allow mirrored packets to be directed locally (within the same router) or remotely, over the core of the network and have a far-end decode mirror encapsulation.

The mirror destination service is comprised of destination parameters that define where the mirrored packets are to be sent. It also specifies whether the defined *service-id* will receive mirrored packets from far-end router over the network core.

The mirror destination service IDs are persistent between boots of the router and are included in the configuration saves. The local sources of mirrored packets for the service ID are defined within the **debug mirror mirror-source** command that references the same *service-id*. Up to 255 mirror destination service IDs can be created within a single system.

The **mirror-dest** command creates or edits a service ID for mirroring purposes. If the *service-id* does not exist within the context of all defined services, the mirror destination service is created and the context of the CLI is changed to that service ID. If the *service-id* exists within the context of defined mirror destination services, the CLI context is changed for editing parameters on that service ID. If the *service-id* exists within the context of another service type, an error message is returned and CLI context is not changed from the current context.

The **no** form of this command removes a mirror destination from the system. The mirror source associations with the mirror destination *service-id* do not need to be removed or shutdown first. The mirror destination *service-id* must be shutdown before the service ID can be removed. When the service ID is removed, all **mirror-source** commands that have the service ID defined will also be removed from the system.

Parameters

service-id

The service identification identifies the service in the service domain. This ID is unique to this service and cannot be used by any other service, regardless of service type. The same service ID must be configured on every router that this particular service is defined on.

If particular a service ID already exists for a service, then the same value cannot be used to create a mirror destination service ID with the same value. For example:

If an Epipe service ID **11** exists, then a mirror destination service ID **11** cannot be created. If a VPLS service ID **12** exists, then a mirror destination service ID **12** cannot be created.

If an IES service ID **13** exists, then a mirror destination service ID **13** cannot be created.

Values

service-id:	1 to 2147483647
svc-name:	64 characters maximum

create

Keyword used to create a mirror destination service.

mirror-type

The type describes the encapsulation supported by the mirror service.

Values ether, ip-only

name

Configures an optional service name identifier, up to 64 characters, to a given service. This service name can then be used in configuration references, display, and show commands throughout the system. A defined service name can help the service provider or administrator to identify and manage services within the SR OS platforms.

To create a service, you must assign a service ID; however, after it is created, either the service ID or the service name can be used to identify and reference a service.

If a name is not specified at creation time, then SR OS assigns a string version of the *service-id* as the name.

Service names may not begin with an integer (0 to 9).

Platforms

7705 SAR Gen 2

18.54 mirror-source

mirror-source

Syntax

[**no**] **mirror-source** *service-id*

Context

[Tree] (config>mirror mirror-source)

Full Context

configure mirror mirror-source

Description

This command configures mirror source parameters for a mirrored service.

The **mirror-source** command is used to enable mirroring of packets specified by the association of the **mirror-source** to sources of packets defined within the context of the *mirror-dest-service-id*. The mirror destination service must already exist within the system.

A mirrored packet cannot be mirrored to multiple destinations. If a packet matches multiple mirror source entries (for example, a SAP on one **mirror-source** and a port on another **mirror-source**), then the packet is mirrored to a single *mirror-dest-service-id* based on the following precedence:

1. Filter entry
2. SAP
3. Physical port

The precedence is structured so the most specific match criteria has precedence over a less specific match. For example, if a **mirror-source** defines a port and a SAP on that port, then a packet arriving on

the SAP will be mirrored using the SAP mirror (and not mirrored using the Port mirror) because the SAP is more specific than the port.

The **no** form of this command deletes all related source commands within the context of the **mirror-source** *service-id*. The command does not remove the service ID from the system.

Parameters

service-id

Specifies the service identification identifies the service in the service domain. This ID is unique to this service and cannot be used by any other service, regardless of service type. The same service ID must be configured on every router that this particular service is defined on.

Values

service-id: 1 to 2147483647

svc-name: 64 characters maximum

Platforms

7705 SAR Gen 2

mirror-source

Syntax

[**no**] **mirror-source** *service-id*

Context

[Tree] (debug mirror-source)

Full Context

debug mirror-source

Description

This command configures mirror source parameters for a mirrored service.

The **mirror-source** command is used to enable mirroring of packets specified by the association of the **mirror-source** to sources of packets defined within the context of the *mirror-dest-service-id*. The mirror destination service must already exist within the system.

A mirrored packet cannot be mirrored to multiple destinations. If a packet matches multiple mirror source entries (for example, a SAP on one **mirror-source** and a port on another **mirror-source**), then the packet is mirrored to a single *mirror-dest-service-id* based on the following precedence:

1. Filter entry
2. SAP
3. Physical port

The precedence is structured so the most specific match criteria has precedence over a less specific match. For example, if a **mirror-source** defines a port and a SAP on that port, then a packet arriving on

the SAP will be mirrored using the SAP mirror (and not mirrored using the Port mirror) because the SAP is more specific than the port.

The **mirror-source** configuration is not saved when a configuration is saved. A **mirror-source** manually configured within an ASCII configuration file will not be preserved if that file is overwritten by a **save** command. Define the **mirror-source** within a file associated with a **config exec** command to make a **mirror-source** persistent between system reboots.

By default, all mirror destination service IDs have a **mirror-source** associated with them. The **mirror-source** is not technically created with this command. Instead the service ID provides a contextual node for storing the current mirroring sources for the associated mirror destination service ID. The **mirror-source** is created for the mirror service when the operator enters the **debug>mirror-source** *svcId* for the first time. The **mirror-source** is also automatically removed when the mirror destination service ID is deleted from the system.

The **no** form of this command deletes all related source commands within the context of the **mirror-source** *service-id*. The command does not remove the service ID from the system.

Parameters

service-id

Specifies the mirror destination service ID for which match criteria will be defined. The *service-id* must already exist within the system.

Values *service-id*: 1 to 2147483647

svc-name: 64 characters maximum

Platforms

7705 SAR Gen 2

18.55 misc

misc

Syntax

[**no**] **misc**

Context

[Tree] (debug>router>igmp misc)

Full Context

debug router igmp misc

Description

This command enables debugging for IGMP miscellaneous information.

The **no** form of the command disables the debugging.

Platforms

7705 SAR Gen 2

Output

The following output is an example of debugged IGMP miscellaneous information.

Output Example

```
A:ALA-CA# debug router 100 igmp misc
*A:ALA-CA# show debug
debug
    router "100"
        igmp
            misc
        exit
    exit
exit
*A:ALA-CA#
```

misc

Syntax

misc [**detail**]

no misc

Context

[Tree] (debug>router>mpls>event misc)

[Tree] (debug>router>rsvp>event misc)

Full Context

debug router mpls event misc

debug router rsvp event misc

Description

This command debugs miscellaneous events.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about miscellaneous events.

Platforms

7705 SAR Gen 2

misc

Syntax

[**no**] **misc**

Context

[Tree] (debug>router>isis misc)

Full Context

debug router isis misc

Description

This command enables debugging for IS-IS misc.

The **no** form of the command disables debugging.

Platforms

7705 SAR Gen 2

misc

Syntax

[**no**] **misc**

Context

[Tree] (debug>router>ospf3 misc)

[Tree] (debug>router>ospf misc)

Full Context

debug router ospf3 misc

debug router ospf misc

Description

This command enables debugging for miscellaneous OSPF events.

Platforms

7705 SAR Gen 2

18.56 mixed-lsp-mode

mixed-lsp-mode

Syntax

[**no**] **mixed-lsp-mode**

Context

[Tree] (config>service>sdp mixed-lsp-mode)

Full Context

configure service sdp mixed-lsp-mode

Description

This command enables the use by an SDP of the mixed-LSP mode of operation. This command indicates to the service manager that it must allow a primary LSP type and a backup LSP type in the same SDP configuration. For example, the **lsp** and **ldp** commands are allowed concurrently in the SDP configuration. The user can configure one or two types of LSPs under the same SDP. Without this command, these commands are mutually exclusive.

The user can configure an RSVP LSP as a primary LSP type with an LDP LSP as a backup type. The user can also configure an RFC 8277 BGP LSP as a backup LSP type.

If the user configures an LDP LSP as a primary LSP type, then the backup LSP type must be an RFC 8277 BGP labeled route.

At any given time, the service manager programs only one type of LSP in the linecard that will activate it to forward service packets according to the following priority order:

- RSVP LSP type. Up to 16 RSVP LSPs can be entered by the user and programmed by the service manager in ingress linecard to load balance service packets. This is the highest priority LSP type.
- LDP LSP type. One LDP FEC programmed by the service manager but the ingress card can use up to 16 LDP ECMP paths for the FEC to load balance service packets when ECMP is enabled on the node.
- BGP LSP type. One RFC 8277-labeled BGP prefix programmed by the service manager. The ingress card can use more than one next-hop for the prefix.

In the case of the RSVP/LDP SDP, the service manager will program the NHLFE(s) for the active LSP type preferring the RSVP LSP type over the LDP LSP type. If no RSVP LSP is configured or all configured RSVP LSPs go down, the service manager will re-program the card with the LDP LSP if available. If not, the SDP goes operationally down.

When a higher priority type LSP becomes available, the service manager reverts back to this LSP at the expiry of the sdp-revert-time timer or the failure of the currently active LSP, whichever comes first. The service manager then re-programs the card accordingly. If the infinite value is configured, then the SDP reverts to the highest priority type LSP only if the currently active LSP failed.

Note:
LDP uses a tunnel down damp timer which is set to three seconds by default. When the LDP LSP fails, the SDP will revert to the RSVP LSP type after the expiry of this timer. For an immediate

switchover, this timer must be set to zero. Use the **config>router>ldp>tunnel-down-damp-time** command.

If the user changes the value of the sdp-revert-time timer, it will take effect only at the next use of the timer. Any timer which is outstanding at the time of the change will be restarted with the new value.

If class based forwarding is enabled for this SDP, the forwarding of the packets over the RSVP LSPs will be based on the FC of the packet as in current implementation. When the SDP activates the LDP LSP type, then packets are forwarded over the LDP ECMP paths using the regular hash routine.

In the case of the LDP/BGP SDP, the service manager will prefer the LDP LSP type over the BGP LSP type. The service manager will re-program the card with the BGP LSP if available otherwise it brings down the SDP operationally.

Also note the following difference in behavior of the LDP/BGP SDP compared to that of an RSVP/LDP SDP. For a given /32 prefix, only a single route will exist in the routing table: the IGP route or the BGP route. Thus, either the LDP FEC or the BGP label route is active at any given time. The impact of this is that the tunnel table needs to be re-programmed each time a route is deactivated and the other is activated. Furthermore, the SDP revert-time cannot be used since there is no situation where both LSP types are active for the same /32 prefix.

The **no** form of this command disables the mixed-LSP mode of operation. The user first has to remove one of the LSP types from the SDP configuration or the command will fail.

Default

no mixed-lsp-mode

Platforms

7705 SAR Gen 2

18.57 mka-hello-interval

mka-hello-interval

Syntax

mka-hello-interval *interval*

no mka-hello-interval

Context

[Tree] (config>macsec>conn-assoc>static-cak mka-hello-interval)

Full Context

configure macsec connectivity-association static-cak mka-hello-interval

Description

This command specifies the MKA hello interval.

The **no** form of this command disables the MKA hello interval.

Default

mka-hello-interval 2

Parameters

interval

Specifies the MKA hello interval, in seconds.

Values 1 to 6 s in 1-s increments, 500ms

Platforms

7705 SAR Gen 2

18.58 mka-key-server-priority

mka-key-server-priority

Syntax

mka-key-server-priority *priority*

no mka-key-server-priority

Context

[Tree] (config>macsec>conn-assoc>static-cak mka-key-server-priority)

Full Context

configure macsec connectivity-association static-cak mka-key-server-priority

Description

This command specifies the key server priority used by the MACsec Key Agreement (MKA) protocol to select the key server when MACsec is enabled using static connectivity association key (CAK) security mode.

The **no** form of this command disables the **mka-key-server-priority**.

Default

mka-key-server-priority 16

Parameters

priority

Specifies the priority of the server.

Values 0 to 255

Platforms
7705 SAR Gen 2

18.59 mld

mld

Syntax
[**no**] **mld**

Context
[Tree] (config>service>vprn mld)

Full Context
configure service vprn mld

Description
Commands in this context configure Multicast Listener Discovery (MLD) parameters.

The **no** form of this command disables MLD.

Default
no mld

Platforms
7705 SAR Gen 2

mld

Syntax
[**no**] **mld**

Context
[Tree] (config>router mld)

Full Context
configure router mld

Description
Commands in this context configure Multicast Listener Discovery (MLD) parameters.

The **no** form of the command disables MLD.

Default
no mld

Platforms
7705 SAR Gen 2

18.60 mld-snooping

mld-snooping

Syntax
mld-snooping

Context
[Tree] (config>service>vpls>allow-ip-int-bind mld-snooping)
[Tree] (config>service>vpls>sap mld-snooping)
[Tree] (config>service>vpls mld-snooping)
[Tree] (config>service>vpls>mesh-sdp mld-snooping)
[Tree] (config>service>vpls>spoke-sdp mld-snooping)

Full Context
configure service vpls allow-ip-int-bind mld-snooping
configure service vpls sap mld-snooping
configure service vpls mld-snooping
configure service vpls mesh-sdp mld-snooping
configure service vpls spoke-sdp mld-snooping

Description
Commands in this context configure MLD snooping parameters.

Platforms
7705 SAR Gen 2

mld-snooping

Syntax
[**no**] **mld-snooping**

Context

[Tree] (debug>service>id mld-snooping)

Full Context

debug service id mld-snooping

Description

This command enables and configures MLD-snooping debugging.

The **no** form of this command disables MLD-snooping debugging.

Platforms

7705 SAR Gen 2

18.61 mode

mode

Syntax

mode {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}

no mode

Context

[Tree] (debug>router>local-dhcp-server mode)

[Tree] (debug>router>ip>dhcp6 mode)

[Tree] (debug>router>ip>dhcp mode)

Full Context

debug router local-dhcp-server mode

debug router ip dhcp6 mode

debug router ip dhcp mode

Description

This command debugs the DHCP tracing detail level.

Parameters

dropped-only

Displays only dropped packets.

ingr-and-dropped

Displays only ingress packet and dropped packets.

egr-ingr-and-dropped

Displays ingress, egress and dropped packets.

Platforms

7705 SAR Gen 2

mode

Syntax

mode {**strict** | **loose** | **strict-no-ecmp**}

no mode

Context

[Tree] (config>service>vprn>if>ipv6>urpf-check mode)

[Tree] (config>service>vprn>nw-if>urpf-check mode)

[Tree] (config>service>ies>if>ipv6>urpf-check mode)

[Tree] (config>service>vprn>if>urpf-check mode)

[Tree] (config>service>ies>if>urpf-check mode)

Full Context

configure service vprn interface ipv6 urpf-check mode

configure service vprn network-interface urpf-check mode

configure service ies interface ipv6 urpf-check mode

configure service vprn interface urpf-check mode

configure service ies interface urpf-check mode

Description

This command specifies the mode of unicast RPF check.

The **no** form of this command reverts to the default (strict) mode.

Default

mode strict

Parameters

strict

When specified, uRPF checks whether incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix.

loose

In **loose** mode, uRPF checks whether incoming packet has source address with a corresponding prefix in the routing table. However, the loose mode does not check whether

the interface expects to receive a packet with a specific source address prefix. This object is valid only when **urpf-check** is enabled.

strict-no-ecmp

When a packet is received on an interface in this mode and the SA matches an ECMP route the packet is dropped by uRPF.

Platforms

7705 SAR Gen 2

mode

Syntax

mode {**automatic** | **manual**}

Context

[Tree] (config>port>dwdm>coherent mode)

Full Context

configure port dwdm coherent mode

Description

This command configures the mode used to compensate for chromatic dispersion.

Parameters

automatic

Sets to automatic mode.

manual

Sets to manual mode.

Platforms

7705 SAR Gen 2

mode

Syntax

mode {**access** | **network** | **hybrid**}

no mode

Context

[Tree] (config>lag mode)

[Tree] (config>port>ethernet mode)

Full Context

configure lag mode

configure port ethernet mode

Description

This command configures an Ethernet port, TDM channel, or SONET/SDH path (sub-port) for **access**, **network**, or **hybrid** mode operation.

An **access** port or channel is used for customer facing traffic on which services are configured. A Service Access Point (SAP) can only be configured on an access port or channel. When a port is configured for access mode, the appropriate **encap-type** must be specified to distinguish the services on the port or SONET path. Once an Ethernet port, a TDM channel or a SONET path has been configured for access mode, multiple services can be configured on the Ethernet port, a TDM channel or SONET path.

A network port or channel participates in the service provider transport or infrastructure network when a network mode is selected. When the network option is configured, the encap-type cannot be configured for the port/channel.

When network mode is selected on a SONET/SDH path, the appropriate control protocols are activated when the need arises. For example, configuring an IP interface on the SONET path activates IPCP while the removal of the IP interface causes the IPCP to be removed. The same applies for MPLS, MPLSCP, and OSICP. When configuring a SONET/SDH port, the mode command must be entered in the channel context or an error message is generated.

A hybrid Ethernet port allows the combination of network and access modes of operation on a per-VLAN basis and must be configured as either dot1q or QinQ encapsulation.

When the hybrid port is configured to the dot1q encapsulation, the user configures a SAP inside a service simply by providing the SAP ID which must include the port-id value of the hybrid mode port and an unused VLAN tag value. The format is <*port-id*>:*qtag1*. A SAP of format <*port-id*>:* also supported.

The user configures a network IP interface under **config>router>if>port** by providing the port name which consists of the port-id of the hybrid mode port and an unused VLAN tag value. The format is <*port-id*>:*qtag1*. The user must explicitly enter a valid value for qtag1. The <*port-id*>:* value is not supported on a network IP interface. The 4096 VLAN tag space on the port is shared among VLAN SAPs and VLAN network IP interfaces.

When the hybrid port is configured to QinQ encapsulation, the user configures a SAP inside a service simply by providing the SAP ID which must include the port-id value of the hybrid mode port and the outer and inner VLAN tag values. The format is <port-id>:qtag1.qtag2. A SAP of format <*port-id*>: *qtag1*.* is also supported. The outer VLAN tag value must not have been used to create an IP network interface on this port. In addition, the qtag1.qtag2 value combination must not have been used by another SAP on this port.

The user configures a network IP interface under **config>router>if>port** by providing the port name which consists of the port-id of the hybrid mode port and a VLAN tag value. The format is <*port-id*>:*qtag1*.*. An outer VLAN tag qtag2 of * creates an IP network interface. In addition, the qtag1.qtag2 value combination must not have been used on another SAP or IP network interface on this port.

The **no** form of this command restores the default.

Default

mode network — For Ethernet ports

mode access — For TDM channel or SONET paths

Parameters

access

Configures the Ethernet port, TDM channel or SONET path as service access.

network

Configures the Ethernet port, TDM channel or SONET path for transport network use.

hybrid

Configures the Ethernet port for hybrid use.

Platforms

7705 SAR Gen 2

mode

Syntax

mode {*rstp* | *comp-dot1w* | *dot1w* | *mstp* | *pmstp*}

no mode

Context

[Tree] (config>service>vpls>stp mode)

[Tree] (config>service>template>vpls-template>stp mode)

Full Context

configure service vpls stp mode

configure service template vpls-template stp mode

Description

This command specifies the version of Spanning Tree Protocol the bridge is currently running.

See section Spanning Tree Operating Modes for more information about these modes.

The **no** form of this command returns the STP variant to the default.

Default

mode rstp

Parameters

rstp

Corresponds to the Rapid Spanning Tree Protocol specified in IEEE 802.1D/D4-2003

dot1w

Corresponds to the mode where the Rapid Spanning Tree is backward compatible with IEEE 802.1w

compdot1w

Corresponds to the Rapid Spanning Tree Protocol in conformance with IEEE 802.1w

mstp

Sets MSTP as the STP mode of operation. Corresponds to the Multiple Spanning Tree Protocol specified in 802.1Q REV/D5.0-09/200

pmstp

The PMSTP mode is only supported in VPLS services where the M-VPLS flag is configured

Platforms

7705 SAR Gen 2

mode

Syntax

mode {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}

no mode

Context

[Tree] (debug>service>id>igmp-snooping mode)

Full Context

debug service id igmp-snooping mode

Description

This command enables and configures the IGMP tracing mode.

The **no** form of this command disables the configures the IGMP tracing mode.

Platforms

7705 SAR Gen 2

mode

Syntax

mode {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}

no mode

Context

[Tree] (debug>service>id>mld mode)

Full Context

debug service id mld-snooping mode

Description

This command enables and configures the MLD tracing mode.

The **no** form of this command disables the configures the MLD tracing mode.

Platforms

7705 SAR Gen 2

mode

Syntax

mode {**auto** | **napt** | **one-to-one**}

no mode

Context

[Tree] (config>service>vprn>nat>outside>pool mode)

Full Context

configure service vprn nat outside pool mode

Description

This command configures the mode of operation of this NAT address pool.

The mode value is only relevant while the value of pool type is equal to largeScale; while the value of pool type is equal to l2Aware, the mode of operation is always NAPT.

Default

mode auto

Parameters

napt

Specifies NAPT (Network Address Port Translation)

auto

The system selects the actual mode based upon other configuration parameters; the actual mode can be NAPT or 1:1 NAT (also known as 'Basic NAT').

one-to-one

Indicates 1:1 NAT (also known as 'Basic NAT')

Platforms

7705 SAR Gen 2

mode

Syntax

mode {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}

no mode

Context

[Tree] (debug>service>id>dhcp mode)

Full Context

debug service id dhcp mode

Description

This command configures the DHCP tracing mode.

The **no** form of the command disables debugging.

Parameters

dropped-only

Only displays dropped packets.

ingr-and-dropped

Only displays ingress packet and dropped packets.

egr-ingr-and-dropped

Displays ingress, egress and dropped packets.

Platforms

7705 SAR Gen 2

mode

Syntax

mode {**auto** | **napt** | **one-to-one**}

no mode

Context

[Tree] (config>router>nat>outside>pool mode)

Full Context

configure router nat outside pool mode

Description

This command specifies the mode of operation of this NAT address pool.

The **no** form of the command reverts to the default.

Default

auto

Parameters

{auto | napt | one-to-one}

Specifies the mode of operation of this NAT pool.

Platforms

7705 SAR Gen 2

mode

Syntax

mode {**strict** | **loose** | **strict-no-ecmp**}

Context

[Tree] (config>router>if>ipv6>urpf-check mode)

[Tree] (config>router>if>urpf-check mode)

Full Context

configure router interface ipv6 urpf-check mode

configure router interface urpf-check mode

Description

This command specifies the mode of unicast RPF check.

The **no** form of this command reverts to the default (strict) mode.

Default

mode strict

Parameters

strict

When specified, uRPF checks whether incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix.

loose

In **loose** mode, uRPF checks whether incoming packet has source address with a corresponding prefix in the routing table. However, the loose mode does not check whether

the interface expects to receive a packet with a specific source address prefix. This object is valid only when **urpf-check** is enabled.

strict-no-ecmp

When a packet is received on an interface in this mode and the SA matches an ECMP route the packet is dropped by uRPF.

Platforms

7705 SAR Gen 2

mode

Syntax

mode {**ecmp-protected** | **linear**}

no mode

Context

[Tree] (config>router>segment-routing>maintenance-policy mode)

Full Context

configure router segment-routing maintenance-policy mode

Description

This command specifies the data path programming and protection mechanism for SR policy candidate paths to which the maintenance policy is applied.

In both the **linear** mode and **ecmp-protected** modes, if two or more candidate paths of the same {headend, color, endpoint} and also have the same mode, then the best preference path is treated as the primary while the next best preference path is the standby. If a third path is present in the linear mode, then this is treated as a tertiary and also programmed in the IOM.

If the currently active path goes unavailable due to S-BFD, the system failovers to the next best preference available candidate path. If S-BFD is down on all segment lists of all programmed candidate paths of an SR Policy, then the SR Policy is marked as down in TTM.

If the default mode is specified, the router only programs the segment lists of the best preference paths in the IOM.

The **no** form of this command removes the configured mode.

Default

no mode

Parameters

ecmp-protected

Specifies only the top two routes (paths) are programmed in the IOM. Up to 32 segment lists can be programmed for each path.

linear

Specifies the top three routes are programmed in the IOM. Only one segment list is allowed per path.

Platforms

7705 SAR Gen 2

mode

Syntax

mode {**target-defined** | **on-change** | **sample**}

Context

[Tree] (config>system>telemetry>persistent-subscriptions>subscription mode)

Full Context

configure system telemetry persistent-subscriptions subscription mode

Description

This command configures the subscription path mode for telemetry notifications that are sent for the persistent subscription.

Default

mode target-defined

Parameters

target-defined

Keyword specifying that target defined mode is used.

on-change

Keyword specifying that on change mode is used.

sample

Keyword specifying that sample mode is used.

Platforms

7705 SAR Gen 2

18.62 monitor-oper-group

monitor-oper-group

Syntax

monitor-oper-group *name*

no monitor-oper-group

Context

[Tree] (config>lag monitor-oper-group)

Full Context

configure lag monitor-oper-group

Description

This command, supported on access LAG only, specifies the operational group to monitor. The state of the operational group affects the state of this LAG. When the operational group is inactive, the state of the LAG goes down and the LAG uses the configured **lag**>**standby-signaling** mechanism (**lacp** or **power-off**) to signal the CE that the LAG is not available.

Default

no monitor-oper-group

Parameters

name

Specifies the name of the **oper-group**, up to 32 characters.

Platforms

7705 SAR Gen 2

monitor-oper-group

Syntax

monitor-oper-group *name*

no monitor-oper-group

Context

[Tree] (config>service>vpls>sap monitor-oper-group)

[Tree] (config>service>vpls>spoke-sdp monitor-oper-group)

[Tree] (config>service>vpls>bgp>pw-template-binding monitor-oper-group)

[Tree] (config>service>vpls>site monitor-oper-group)

Full Context

configure service vpls sap monitor-oper-group

configure service vpls spoke-sdp monitor-oper-group

configure service vpls bgp pw-template-binding monitor-oper-group

configure service vpls site monitor-oper-group

Description

This command specifies the operational group to be monitored by the object under which it is configured. The **oper-group** *name* must be already configured under the **config>service** context before its name is referenced in this command.

The **no** form of this command removes the association.

Default

no monitor-oper-group

Parameters

group-name

Specifies a character string of maximum 32 ASCII characters identifying the group instance.

Platforms

7705 SAR Gen 2

monitor-oper-group

Syntax

monitor-oper-group *group-name*

no monitor-oper-group

Context

[Tree] (config>service>epipe>spoke-sdp monitor-oper-group)

[Tree] (config>service>epipe>sap monitor-oper-group)

Full Context

configure service epipe spoke-sdp monitor-oper-group

configure service epipe sap monitor-oper-group

Description

This command specifies the operational group to be monitored by the object under which it is configured. The **oper-group** *name* must be already configured under the **config>service** context before its name is referenced in this command.

The **no** form of this command removes the association.

Parameters

group-name

Specifies an oper group name.

Platforms

7705 SAR Gen 2

monitor-oper-group

Syntax

monitor-oper-group *name*

no monitor-oper-group

Context

[Tree] (config>service>ies>if monitor-oper-group)

Full Context

configure service ies interface monitor-oper-group

Description

This command specifies the operational group to be monitored by the object under which it is configured. The oper-group name must be already configured under the config>service context before its name is referenced in this command.

The **no** form of this command removes the association from the configuration.

Default

no monitor-oper-group

Parameters

name

Specifies a character string of maximum 32 ASCII characters identifying the group instance.

Platforms

7705 SAR Gen 2

monitor-oper-group

Syntax

monitor-oper-group *name*

no monitor-oper-group

Context

[Tree] (config>service>vprn>if monitor-oper-group)

Full Context

configure service vprn interface monitor-oper-group

Description

This command specifies the operational group to be monitored by the object under which it is configured. The oper-group name must be already configured under the config>service context before its name is referenced in this command.

The **no** form of this command removes the association from the configuration.

Default

no monitor-oper-group

Parameters

name

Specifies a character string, up to 32 ASCII characters, identifying the group instance.

Platforms

7705 SAR Gen 2

monitor-oper-group

Syntax

monitor-oper-group *group-name* **family** {**ipv4** | **ipv6**} **add** [*1..4294967295*]

monitor-oper-group *group-name* **family** {**ipv4** | **ipv6**} **set** [*1..4294967295*]

monitor-oper-group *group-name* **family** {**ipv4** | **ipv6**} **subtract** [*1..4294967295*]

no monitor-oper-group [**family** {**ipv4** | **ipv6**}]

Context

[Tree] (config>service>vprn>pim>if monitor-oper-group)

Full Context

configure service vprn pim interface monitor-oper-group

Description

This command configures PIM to monitor the state of an operational group to provide a redundancy mechanism. PIM monitors the operational group and changes its DR priority to the specified value when the status of the operational group is up. This enables the router to become the PIM DR only when the operational group is up. If the operational group status changes to down, PIM changes its DR priority to the default or the value configured with **priority** under **config**>**service**>**vprn**>**pim**>**if**. The **oper-group** *group-name* must already be configured under the **config**>**service** context before its name is referenced in this command. Two operational groups are supported per PIM interface.

The **no** form of this command removes the operational group from the configuration.

Parameters

group-name

Specifies the operational group identifier up to 32 characters in length.

family

Specifies the address family.

ipv4

Specifies the IPv4 designated router priority.

ipv6

Specifies the IPv6 designated router priority.

add

Specifies that the value is to be added to the existing priority to become the designated router.

subtract

Specifies that the value is to be subtracted from the existing priority to become the designated router.

set

Specifies the priority to become the designated router.

value

Specifies the priority modifier expressed as a decimal integer.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

monitor-oper-group

Syntax

monitor-oper-group *group-name* **family** {**ipv4** | **ipv6**} **add** [*value*]

monitor-oper-group *group-name* **family** {**ipv4** | **ipv6**} **set** [*value*]

monitor-oper-group *group-name* **family** {**ipv4** | **ipv6**} **subtract** [*value*]

no monitor-oper-group [**family** {**ipv4** | **ipv6**}]

Context

[Tree] (config>router>pim>if monitor-oper-group)

Full Context

configure router pim interface monitor-oper-group

Description

This command configures PIM to monitor the state of an operational group to provide a redundancy mechanism. PIM monitors the operational group and changes its DR priority to the specified value when the status of the operational group is up. This enables the router to become the PIM DR only when the operational group is up. If the operational group status changes to down, PIM changes its DR priority to the default or the value configured with **priority** under **config**>**router**>**pim**>**if**. The **oper-group** *group-name* must already be configured under the **config**>**service** context before its name is referenced in this command. Two operational groups are supported per PIM interface.

The **no** form of this command removes the operational group from the configuration.

Parameters

group-name

Specifies the operational group identifier, up to 32 characters.

family

Specifies the address family.

ipv4

Specifies the IPv4 designated router priority.

ipv6

Specifies the IPv6 designated router priority.

add

Specifies that the value is to be added to the existing priority to become the designated router.

subtract

Specifies that the value is to be subtracted from the existing priority to become the designated router.

set

Specifies the priority to become the designated router.

value

Specifies the priority modifier expressed as a decimal integer.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

monitor-oper-group

Syntax

monitor-oper-group *group-name*

no monitor-oper-group

Context

[Tree] (config>port monitor-oper-group)

Full Context

configure port monitor-oper-group

Description

This command configures the operational group to monitor the operational group state. The state of the operational group affects the state of this port. When the operational group is inactive, the state of the port goes down and powers off the port to signal the CE that the connected port is not available.

Default

no monitor-oper-group

Parameters

group-name

Specifies operational group name to monitor, up to 32 characters.

Platforms

7705 SAR Gen 2

18.63 month

month

Syntax

month {*month-number* [*..month-number*] | *month-name* [*..month-name*] | **all**}

no month

Context

[Tree] (config>system>cron>sched month)

Full Context

configure system cron schedule month

Description

This command specifies the month when the event should be executed. Multiple months can be specified. When multiple months are configured, each of them will cause the schedule to trigger. If a month is configured without configuring the month, weekday, day-of-month, and minute, the event will not execute.

The **no** form of this command removes the specified month from the configuration.

Default

no month

Parameters

month-number

Specifies a month number.

Values 1 to 12 (maximum 12 month-numbers)

month-name

Specifies a month by name.

Values january, february, march, april, may, june, july, august, september, october, november, december (maximum 12 month names)

all

Specifies all months.

Platforms

7705 SAR Gen 2

18.64 more

more

Syntax

[**no**] **more**

Context

[Tree] (environment more)

Full Context

environment more

Description

This command enables per-screen CLI output, meaning that the output is displayed on a screen-by- screen basis. The terminal screen length can be modified with the **terminal** command.

The following prompt appears at the end of each screen of paginated output:

```
Press any key to continue (Q to quit)
```

The **no** form of the command displays the output all at once. If the output length is longer than one screen, the entire output will be displayed, which may scroll the screen.

Default

more

Platforms

7705 SAR Gen 2

more

Syntax

[**no**] **more**

Context

[Tree] (config>system>management-interface>cli>md-cli>environment more)

Full Context

configure system management-interface cli md-cli environment more

Description

This command configures pagination of the output text.

The **no** form of this command reverts to the default value.

Default

more

Platforms

7705 SAR Gen 2

18.65 motd

motd

Syntax

motd {**url** *url-prefix*: *source-url* | **text** *motd-text-string*}

no motd

Context

[Tree] (config>system>login-control motd)

Full Context

configure system login-control motd

Description

This command creates the message of the day displayed after a successful console login. Only one message can be configured.

The **no** form of this command removes the message.

Default

no motd

Parameters

url *url-prefix*: *source-url*

When the message of the day is present as a text file, provide both url-prefix and the source-url of the file containing the message of the day. The URL prefix can be local or remote.

text *motd-text-string*

Specifies the text of the message of the day. The *motd-text-string* must be enclosed in double quotes. Multiple text strings are not appended to one another.

Some special characters can be used to format the message text. The \n character can be used to create multi-line messages. A \n in the message moves to the beginning of the next line by sending ASCII/UTF-8 chars 0xA (LF) and 0xD (CR) to the client terminal. An \r in the message sends the ASCII/UTF-8 char 0xD (CR) to the client terminal.

Platforms

7705 SAR Gen 2

18.66 move

move

Syntax

move *old-file-url new-file-url* [**force**] [**no-redirect**] [**client-tls-profile** *profile*] [**proxy** *proxy-url*]

Context

[Tree] (file move)

Full Context

file move

Description

This command moves a local file, system file, or a directory. If the target already exists, the command fails and an error message displays.

The following prompt appears if the destination file already exists:

"Overwrite destination file (y/n)?”

Parameters

old-file-url

Specifies the file or directory to be moved.

Values		
	local-url	[*cflash-id*/][*file-path*] up to 200 characters, including cflash-id directory length up to 99 each
	remote-url	[{ftp:// | tftp:// | http:// | https://}*login*:*pswd*@*remote-locn*/][*file-path*]
		up to 247 characters
		directory length up to 99 characters each
	remote-locn	[hostname | ipv4-address | [ipv6-address]]
	ipv4-address	*a.b.c.d*
	ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
		x:x:x:x:x:x:d.d.d.d[-*interface*]
		x - [0 to FFFF]H
		d - [0 to 255]D
		interface - up to 32 characters, for link local addresses 255
	cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

new-file-url

Specifies the new destination to place the old-file-url.

Values		
	local-url	[*cflash-id*/][*file-path*] up to 200 characters, including cflash-id directory length up to 99 each
	remote-url	[{ftp:// | tftp://}*login*:*pswd*@*remote-locn*/][*file-path*]
		up to 247 characters
		directory length up to 99 characters each

remote-locn	[hostname | ipv4-address | [ipv6-address]]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D
	interface - up to 32 characters, for link local addresses 255
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

force

Forces an immediate move of the specified file(s).

The **file move force** command moves the specified file(s) without displaying a user prompt message. This command also automatically accepts HTTP redirects unless overridden by the **no-redirect** parameter.

profile

Specifies the TLS client profile configured under **config**>**system**>**security**>**tls**>**client-tls-profile** to use.

proxy-url

Specifies the URL of an HTTP proxy. For example, http://proxy.mydomain.com:8000. This URL must be an HTTP URL and not an HTTPS URL.

no-redirect

Specifies to automatically refuse any HTTP redirects without prompting the user.

Platforms

7705 SAR Gen 2

18.67 move-frequency

move-frequency

Syntax

move-frequency *frequency*

no move-frequency

Context

[Tree] (config>service>vpls>mac-move move-frequency)

[Tree] (config>service>template>vpls-template>mac-move move-frequency)

Full Context

configure service vpls mac-move move-frequency

configure service template vpls-template mac-move move-frequency

Description

This command indicates the maximum rate at which MACs can be relearned in the VPLS service before the SAP where the moving MAC was last seen is automatically disabled to protect the system against undetected loops or duplicate MACs. The rate (relearns per second) is measured in a 5-second window.

The **no** form of this command reverts to the default value.

Default

2 (relearns per second, when mac-move is enabled). For example, the value 2 requires 10 MAC relearns in a 5-second period for the MAC to be considered duplicate.

Parameters

frequency

Specifies the rate, in relearns per second.

Values 1 to 10

Platforms

7705 SAR Gen 2

18.68 mp-bgp-keep

mp-bgp-keep

Syntax

[**no**] **mp-bgp-keep**

Context

[Tree] (config>router>bgp mp-bgp-keep)

Full Context

configure router bgp mp-bgp-keep

Description

As a result of enabling this command, route refresh messages are no longer needed, or issued when VPN route policy changes are made; RIB-IN will retain all MP-BGP routes.

The **no** form of this command is used to disable this feature.

Default

no mp-bgp-keep

Platforms

7705 SAR Gen 2

18.69 mpls

mpls

Syntax

mpls [**bgp** *bgp*] [**endpoint** *endpoint-name*]

no mpls [**bgp** *bgp*]

Context

[Tree] (config>service>vprn>bgp-evpn mpls)

[Tree] (config>service>epipe>bgp-evpn mpls)

[Tree] (config>service>vpls>bgp-evpn mpls)

Full Context

configure service vprn bgp-evpn mpls

configure service epipe bgp-evpn mpls

configure service vpls bgp-evpn mpls

Description

Commands in this context configure the BGP EVPN MPLS parameters. In VPLS, either instance BGP 1 or BGP 2 can be configured, but not both simultaneously in the same service. Epipe and VPRN services only support instance 1. If the **bgp** *bgp* parameter is not specified, the instance is set to 1.

The **endpoint** option is only supported for Epipe services. When configured, the same endpoint name can be configured for the **bgp-evpn>mpls** context and an additional spoke SDP. An EVPN MPLS destination always has higher preference than a spoke SDP.

The **no** form of this command removes the MPLS instance from the service.

Parameters

bgp

Indicates the BGP instance identifier.

Values 1, 2

endpoint-name

Specifies the endpoint name for Epipe services, up to 32 characters.

Platforms

7705 SAR Gen 2

mpls

Syntax

mpls

Context

[Tree] (config>service>vprn>bgp-ipvpn mpls)

Full Context

configure service vprn bgp-ipvpn mpls

Description

Commands in this context configure the BGP IPVPN parameters.

Platforms

7705 SAR Gen 2

mpls

Syntax

[**no**] **mpls**

Context

[Tree] (config>router mpls)

Full Context

configure router mpls

Description

Commands in this context configure MPLS parameters. MPLS is not enabled by default and must be explicitly enabled (**no shutdown**).

The **no** form of this command deletes this MPLS protocol instance and removes all configuration parameters for this MPLS instance.

You must remove all SDP bindings and use the **shutdown** command to administratively disable the MPLS instance before deleting it.

Platforms

7705 SAR Gen 2

mpls

Syntax

mpls [**lsp** *lsp-name*] [**sender** *source-address*] [**endpoint** *endpoint-address*] [**tunnel-id** *tunnel-id*] [**lsp-id** *lsp-id*]

no mpls

Context

[Tree] (debug>router mpls)

Full Context

debug router mpls

Description

This command enables and configures debugging for MPLS.

Parameters

lsp *lsp-name*

Specifies the LSP name up to 64 characters in length.

sender *source-address*

Specifies the IP address of the sender.

endpoint *endpoint-address*

Specifies the far-end IP address.

tunnel-id *tunnel-id*

Specifies the MPLS SDP ID.

Values 0 to 4294967295

lsp-id *lsp-id*

Specifies the LSP ID.

Values 1 to 65535

Platforms

7705 SAR Gen 2

mpls

Syntax

[**no**] **mpls**

Context

[Tree] (config>oam-pm>session>ip>tunnel mpls)

Full Context

configure oam-pm session ip tunnel mpls

Description

Commands in this context configure the MPLS packet tunneling options for the session. Configure the **tunnel oam-pm session ip router-instance** to Base to configure commands in the MPLS context. When entering a context under MPLS, the system removes any previous configuration of any sibling context. You can only configure one of the contexts for each OAM-PM session.

The **no** form of this command deletes the **mpls** context and all configurations under it.

Platforms

7705 SAR Gen 2

18.70 mpls-echo-request-downstream-map

mpls-echo-request-downstream-map

Syntax

mpls-echo-request-downstream-map {**dsmap** | **ddmap**}

Context

[Tree] (config>test-oam mpls-echo-request-downstream-map)

Full Context

configure test-oam mpls-echo-request-downstream-map

Description

This command specifies which format of the downstream mapping TLV to use in all LSP trace packets and LDP tree trace packets originated on this node. The Downstream Mapping (DSMAP) TLV is the original format in RFC 4379 (obsoleted by RFC 8029) and is the default value. The new Downstream Detailed Mapping (DDMAP) TLV is the new enhanced format specified in RFC 6424 and RFC 8029.

This command applies to LSP trace of an RSVP P2P LSP, a MPLS-TP LSP, or LDP unicast FEC, and to LDP tree trace of a unicast LDP FEC. It does not apply to LSP trace of an RSVP P2MP LSP which always uses the DDMAP TLV.

The global DSMAP/DDMAP setting impacts the behavior of both OAM LSP trace packets and SAA test packets of type **lsp-trace** and is used by the sender node when one of the following events occurs:

1. An SAA test of type **lsp-trace** is created (not modified) and no value is specified for the per-test **downstream-map-tlv** {**dsmap** | **ddmap** | **none**} option. In this case, the SAA test **downstream-map-tlv** value defaults to the global **mpls-echo-request-downstream-map** value.
2. An OAM test of type **lsp-trace** test is executed and no value is specified for the per-test **downstream-map-tlv** {**dsmap** | **ddmap** | **none**} option. In this case, the OAM test **downstream-map-tlv** value defaults to the global **mpls-echo-request-downstream-map** value.

A consequence of the rules above is that a change to the value of **mpls-echo-request-downstream-map** option does not affect the value inserted in the downstream mapping TLV of existing tests.

Following are the details of the processing of the new DDMAP TLV:

1. When either the DSMAP TLV or the DDMAP TLV is received in an echo request message, the responder node includes the same type of TLV in the echo reply message with the proper downstream interface information and label stack information.
2. If an echo request message without a Downstream Mapping TLV (DSMAP or DDMAP) expires at a node which is not the egress for the target FEC stack, the responder node always includes the DSMAP TLV in the echo reply message. This can occur in the following cases:
   a. The user issues a LSP trace from a sender node with a **min-ttl** value higher than 1 and a **max-ttl** value lower than the number of hops to reach the egress of the target FEC stack. This is the sender node behavior when the global configuration or the per-test setting of the DSMAP/DDMAP is set to DSMAP.
   b. The user issues a LSP ping from a sender node with a **ttl** value lower than the number of hops to reach the egress of the target FEC stack. This is the sender node behavior when the global configuration of the DSMAP/DDMAP is set to DSMAP.
   c. The behavior in (a) is changed when the global configuration or the per-test setting of the Downstream Mapping TLV is set to DDMAP. The sender node includes in this case the DDMAP TLV with the Downstream IP address field set to the all-routers multicast address as per Section 3.4 of RFC 8029. The responder node then bypasses the interface and label stack validation and replies with a DDMAP TLV with the correct downstream information for the target FEC stack.
3. A sender node never includes the DSMAP or DDMAP TLV in an lsp-ping message.

In addition to performing the same features as the DSMAP TLV, the new DDMAP TLV addresses the following scenarios:

1. Full validation of an LDP FEC stitched to a BGP IPv4 label route. In this case, the LSP trace message is inserted from the LDP LSP segment or from the stitching point.
2. Full validation of a BGP IPv4 label route stitched to an LDP FEC. This includes the case of explicit configuration of the LDP-BGP stitching in which the BGP label route is active in Route Table Manager (RTM) and the case of a BGP IPv4 label route resolved to the LDP FEC due to the IGP route of the same prefix active in RTM. In this case, the LSP trace message is inserted from the BGP LSP segment or from the stitching point.
3. Full validation of an LDP FEC which is stitched to a BGP LSP and stitched back into an LDP FEC. In this case, the LSP trace message is inserted from the LDP segments or the or from the stitching points.
4. Full validation of an LDP FEC tunneled over an RSVP LSP using LSP trace.

To properly check a target FEC which is stitched to another FEC (stitching FEC) of the same or a different type, or which is tunneled over another FEC (tunneling FEC), it is necessary for the responding nodes to provide details about the FEC manipulation back to the sender node. This is achieved via the use of the new FEC stack change sub-TLV in the Downstream Detailed Mapping TLV (DDMAP) defined in RFC 6424.

When the user configures the use of the DDMAP TLV on a trace for an LSP that does not undergo stitching or tunneling operation in the network, the procedures at the sender and responder nodes are the same as in the case of the DSMAP TLV.

This feature however introduces changes to the target FEC stack validation procedures at the sender and responder nodes in the case of LSP stitching and LSP hierarchy. These changes pertain to the processing of the new FEC stack change sub-TLV in the new DDMAP TLV and the new return code of value 15 `Label switched with FEC change`.

The **no** form of this command reverts to the default behavior of using the DSMAP TLV in a LSP trace packet and LDP tree trace packet.

Default

mpls-echo-request-downstream-map dsmap

Parameters

dsmap

Specifies that the DSMAP TLV should be used in all LSP trace packets and LDP tree trace packets originating on the node.

ddmap

Specifies that the DDMAP TLV should be used in all LSP trace packets and LDP tree trace packets originating on the node.

Platforms

7705 SAR Gen 2

Output

The following output is an example of mpls-echo-request-downstream-map information.

Output Example: LDP-over-RSVP

```
A          B          D          F          E          C
o -------- o -------- o -------- o -------- o -------- o
| \______/ | \____________________________/ | \______/ |
 \  RSVP  /               LDP                \  RSVP  /
  \______/                                    \______/
    LDP                                         LDP

Testing LDP FEC of Node C with DSMAP TLV
----------------------------------------
*A:Dut-A#
*A:Dut-A# oam lsp-trace prefix 10.20.1.3/32 downstream-map-tlv dsmap detail
lsp-trace to 10.20.1.3/32: 0 hops min, 0 hops max, 104 byte packets
1  10.20.1.2  rtt=3.90ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.4.4 ifaddr=10.10.4.4 iftype=ipv4Numbered MRU=1500
           label[1]=131068 protocol=3(LDP)
2  10.20.1.4  rtt=5.69ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.9.6 ifaddr=10.10.9.6 iftype=ipv4Numbered MRU=1500
           label[1]=131066 protocol=3(LDP)
```

```
3  10.20.1.6  rtt=7.88ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.10.5 ifaddr=10.10.10.5 iftype=ipv4Numbered MRU=1500
           label[1]=131060 protocol=3(LDP)
4  10.20.1.5  rtt=23.2ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.5.3 ifaddr=10.10.5.3 iftype=ipv4Numbered MRU=1496
           label[1]=131071 protocol=3(LDP)
5  10.20.1.3  rtt=12.0ms rc=3(EgressRtr) rsc=1
*A:Dut-A#


Testing LDP FEC of Node C with DDMAP TLV
----------------------------------------
*A:Dut-A# oam lsp-trace prefix 10.20.1.3/32 downstream-map-tlv ddmap detail
lsp-trace to 10.20.1.3/32: 0 hops min, 0 hops max, 136 byte packets
1  10.20.1.2  rtt=4.00ms rc=3(EgressRtr) rsc=2
1  10.20.1.2  rtt=3.48ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.4.4 ifaddr=10.10.4.4 iftype=ipv4Numbered MRU=1500
           label[1]=131068 protocol=3(LDP)
2  10.20.1.4  rtt=5.34ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.9.6 ifaddr=10.10.9.6 iftype=ipv4Numbered MRU=1500
           label[1]=131066 protocol=3(LDP)
3  10.20.1.6  rtt=7.78ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.10.5 ifaddr=10.10.10.5 iftype=ipv4Numbered MRU=1500
           label[1]=131060 protocol=3(LDP)
4  10.20.1.5  rtt=12.8ms rc=15(LabelSwitchedWithFecChange) rsc=1
     DS 1: ipaddr=10.10.5.3 ifaddr=10.10.5.3 iftype=ipv4Numbered MRU=1496
           label[1]=131054 protocol=4(RSVP-TE)
           label[2]=131071 protocol=3(LDP)
           fecchange[1]=PUSH fectype=RSVP IPv4 prefix=10.20.1.3 remotepeer=10.10.5.3
5  10.20.1.3  rtt=12.8ms rc=3(EgressRtr) rsc=2
5  10.20.1.3  rtt=13.4ms rc=3(EgressRtr) rsc=1
*A:Dut-A#



D          F          E          C          A          B
o -------- o -------- o -------- o -------- o -------- o
  \_____________________/ | \_____________________/ | \______/ |
          LDP             \      RSVP ECA       / \   RSVP   /
                           \_____________________/    \______/
                                    LDP                LDP

Testing LDP FEC of Node B with DDMAP TLV
----------------------------------------
*A:Dut-D#
*A:Dut-D# oam lsp-trace prefix 10.20.1.2/32 downstream-map-tlv ddmap detail
lsp-trace to 10.20.1.2/32: 0 hops min, 0 hops max, 108 byte packets
1  10.20.1.6  rtt=3.17ms rc=8(DSRtrMatchLabel) rsc=1
     DS 1: ipaddr=10.10.10.5 ifaddr=10.10.10.5 iftype=ipv4Numbered MRU=1500
           label[1]=131065 protocol=3(LDP)
2  10.20.1.5  rtt=8.27ms rc=15(LabelSwitchedWithFecChange) rsc=1
     DS 1: ipaddr=10.10.5.3 ifaddr=10.10.5.3 iftype=ipv4Numbered MRU=1496
           label[1]=131068 protocol=4(RSVP-TE)
           label[2]=131065 protocol=3(LDP)
           fecchange[1]=PUSH fectype=RSVP IPv4 prefix=10.20.1.1 remotepeer=10.10.5.3
3  10.20.1.3  rtt=9.50ms rc=8(DSRtrMatchLabel) rsc=2
     DS 1: ipaddr=10.10.2.1 ifaddr=10.10.2.1 iftype=ipv4Numbered MRU=1500
           label[1]=131068 protocol=4(RSVP-TE)
4  10.20.1.1  rtt=10.4ms rc=3(EgressRtr) rsc=2
4  10.20.1.1  rtt=10.2ms rc=15(LabelSwitchedWithFecChange) rsc=1
     DS 1: ipaddr=10.10.1.2 ifaddr=10.10.1.2 iftype=ipv4Numbered MRU=1496
           label[1]=131066 protocol=4(RSVP-TE)
           label[2]=131071 protocol=3(LDP)
           fecchange[1]=PUSH fectype=RSVP IPv4 prefix=10.20.1.2 remotepeer=10.10.1.2
```

```
5  10.20.1.2  rtt=13.7ms rc=3(EgressRtr) rsc=2
5  10.20.1.2  rtt=13.6ms rc=3(EgressRtr) rsc=1
*A:Dut-D#
```

Output Example: LDP-BGP Stitching

```
A         B          C           D         E         F
o ------- o -------- o --------- o ------- o ------- o
  \_____/ |            \_______/           | \_____/
    LDP   |              RSVP              |   LDP
          |\_____________________________/|
          |               LDP              |
           \_____________________________/
                          BGP

Testing LDP FEC of Node F with DSMAP TLV
----------------------------------------
*A:Dut-A# *A:Dut-A# oam lsp-trace prefix 10.20.1.6/32 downstream-map-
tlv dsmap detail lsp-trace to 10.20.1.6/32: 0 hops min, 0 hops max, 104 byte packets
1  10.20.1.2  rtt=2.65ms rc=8(DSRtrMatchLabel) rsc=1
2  10.20.1.3  rtt=4.89ms rc=8(DSRtrMatchLabel) rsc=1
3  10.20.1.4  rtt=6.49ms rc=5(DSMappingMismatched) rsc=1
*A:Dut-A#

Testing LDP FEC of Node F with DDMAP TLV
----------------------------------------
*A:Dut-A# oam lsp-trace prefix 10.20.1.6/32 downstream-map-tlv ddmap detail lsp-
trace to 10.20.1.6/32: 0 hops min, 0 hops max, 108 byte packets
1  10.20.1.2  rtt=3.50ms rc=15(LabelSwitchedWithFecChange) rsc=1
      DS 1: ipaddr=10.10.3.3 ifaddr=10.10.3.3 iftype=ipv4Numbered MRU=1496
            label[1]=131068 protocol=3(LDP)
            label[2]=131060 protocol=2(BGP)
            fecchange[1]=POP  fectype=LDP IPv4 prefix=10.20.1.6 remotepeer=0.0.0.0
(Unknown)
            fecchange[2]=PUSH fectype=BGP IPv4 prefix=10.20.1.6 remotepeer=10.20.1.5
            fecchange[3]=PUSH fectype=LDP IPv4 prefix=10.20.1.5 remotepeer=10.10.3.3
2  10.20.1.3  rtt=6.53ms rc=15(LabelSwitchedWithFecChange) rsc=2
      DS 1: ipaddr=10.10.11.4 ifaddr=10.10.11.4 iftype=ipv4Numbered MRU=1496
          label[1]=131060 protocol=4(RSVP-TE)
          label[2]=131070 protocol=3(LDP)
          label[3]=131060 protocol=2(BGP)
          fecchange[1]=PUSH fectype=RSVP IPv4 prefix=10.20.1.4 remotepeer=10.10.11.4
3  10.20.1.4  rtt=7.94ms rc=3(EgressRtr) rsc=3
3  10.20.1.4  rtt=6.69ms rc=8(DSRtrMatchLabel) rsc=2
      DS 1: ipaddr=10.10.6.5 ifaddr=10.10.6.5 iftype=ipv4Numbered MRU=1500
            label[1]=131071 protocol=3(LDP)
            label[2]=131060 protocol=2(BGP)
4  10.20.1.5  rtt=10.1ms rc=3(EgressRtr) rsc=2
4  10.20.1.5  rtt=8.97ms rc=15(LabelSwitchedWithFecChange) rsc=1
      DS 1: ipaddr=10.10.10.6 ifaddr=10.10.10.6 iftype=ipv4Numbered MRU=1500
           label[1]=131071 protocol=3(LDP)
           fecchange[1]=POP  fectype=BGP IPv4 prefix=10.20.1.6 remotepeer=0.0.0.0
(Unknown)
           fecchange[2]=PUSH fectype=LDP IPv4 prefix=10.20.1.6 remotepeer=10.10.10.6
5  10.20.1.6  rtt=11.8ms rc=3(EgressRtr) rsc=1 *A:Dut-A#


A         B          C           D         E
o ------- o -------- o --------- o ---3--- o
  \_____/ |            \_______/           |
    LDP   |              RSVP              |
          |\_____________________________/|
          |               LDP              |
```

```
           \_____________________________/
                         BGP
Testing BGP Label Route of Node E with DDMAP TLV
------------------------------------------------
*A:Dut-B# oam lsp-trace prefix 11.20.1.5/32 bgp-label downstream-map-
tlv ddmap detail lsp-trace to 11.20.1.5/32: 0 hops min, 0 hops max, 124 byte packets
1  10.20.1.3  rtt=2.35ms rc=15(LabelSwitchedWithFecChange) rsc=2
      DS 1: ipaddr=10.10.11.4 ifaddr=10.10.11.4 iftype=ipv4Numbered MRU=1496
          label[1]=131060 protocol=4(RSVP-TE)
          label[2]=131070 protocol=3(LDP)
          label[3]=131070 protocol=2(BGP)
          fecchange[1]=PUSH fectype=RSVP IPv4 prefix=10.20.1.4 remotepeer=10.10.11.4
2  10.20.1.4  rtt=4.17ms rc=3(EgressRtr) rsc=3
2  10.20.1.4  rtt=4.50ms rc=8(DSRtrMatchLabel) rsc=2
      DS 1: ipaddr=10.10.6.5 ifaddr=10.10.6.5 iftype=ipv4Numbered MRU=1500
            label[1]=131071 protocol=3(LDP)
            label[2]=131070 protocol=2(BGP)
3  10.20.1.5  rtt=7.78ms rc=3(EgressRtr) rsc=2
3  10.20.1.5  rtt=6.80ms rc=3(EgressRtr) rsc=1 *A:Dut-B#

B          C           D         E         F
o -------- o --------- o ---3--- o ---3--- o
|            \_______/           | \_____/
|              RSVP              |
|\_______________________________/|
|              LDP               |
 \_______________________________/
              BGP

Testing with DDMAP TLV LDP FEC of Node F when stitched to a BGP Label Route
---------------------------------------------------------------------------

*A:Dut-B# oam lsp-trace prefix 10.20.1.6/32 bgp-label downstream-map-
tlv ddmap detail lsp-trace to 10.20.1.6/32: 0 hops min, 0 hops max, 124 byte packets
1  10.20.1.3  rtt=3.21ms rc=15(LabelSwitchedWithFecChange) rsc=2
      DS 1: ipaddr=10.10.11.4 ifaddr=10.10.11.4 iftype=ipv4Numbered MRU=1496
          label[1]=131060 protocol=4(RSVP-TE)
          label[2]=131070 protocol=3(LDP)
          label[3]=131060 protocol=2(BGP)
          fecchange[1]=PUSH fectype=RSVP IPv4 prefix=10.20.1.4 remotepeer=10.10.11.4
2  10.20.1.4  rtt=5.50ms rc=3(EgressRtr) rsc=3
2  10.20.1.4  rtt=5.37ms rc=8(DSRtrMatchLabel) rsc=2
      DS 1: ipaddr=10.10.6.5 ifaddr=10.10.6.5 iftype=ipv4Numbered MRU=1500
          label[1]=131071 protocol=3(LDP)
          label[2]=131060 protocol=2(BGP)
3  10.20.1.5  rtt=7.82ms rc=3(EgressRtr) rsc=2
3  10.20.1.5  rtt=6.11ms rc=15(LabelSwitchedWithFecChange) rsc=1
      DS 1: ipaddr=10.10.10.6 ifaddr=10.10.10.6 iftype=ipv4Numbered MRU=1500
          label[1]=131071 protocol=3(LDP)
          fecchange[1]=POP  fectype=BGP IPv4 prefix=10.20.1.6 remotepeer=0.0.0.0
(Unknown)
          fecchange[2]=PUSH fectype=LDP IPv4 prefix=10.20.1.6 remotepeer=10.10.10.6
4  10.20.1.6  rtt=10.2ms rc=3(EgressRtr) rsc=1 *A:Dut-B#
```

18.71 mpls-label

mpls-label

Syntax

mpls-label *value*

no mpls-label

Context

[Tree] (conf>router>segment-routing>sr-policies>policy>seg-list>segment mpls-label)

Full Context

configure router segment-routing sr-policies static-policy segment-list segment mpls-label

Description

This command configures the MPLS label value this is associated with a segment.

The **no** form of this command removes the label value.

Default

no mpls-label

Parameters

value

Specifies the MPLS label value.

Values 0 to 1048575

Platforms

7705 SAR Gen 2

18.72 mpls-labels

mpls-labels

Syntax

mpls-labels

Context

[Tree] (config>router mpls-labels)

Full Context

configure router mpls-labels

Description

This command creates a context for the configuration of global parameters related to MPLS labels.

Platforms

7705 SAR Gen 2

18.73 mpls-time-stamp-format

mpls-time-stamp-format

Syntax

mpls-time-stamp-format {**rfc4379** | **unix**}

Context

[Tree] (config>test-oam mpls-time-stamp-format)

Full Context

configure test-oam mpls-time-stamp-format

Description

This command configures the format of the timestamp used by for lsp-ping, lsp-trace, p2mp-lsp-ping and p2mp-lsp-trace, vccv-ping, vccv-trace, and lsp-trace.

If **rfc4379** is selected, then the timestamp is in seconds and microseconds since 1900, otherwise it is in seconds and microseconds since 1970.

Changing this system-wide setting does not affect tests that are currently in progress, but SAAs starts to use the new timestamp when they are restarted. When an SR OS receives an echo request, it replies with the locally configured timestamp format, and does not try to match the timestamp format of the incoming echo request message.

Default

mpls-time-stamp-format unix

Parameters

rfc4379

Specifies the RFC 4379 (obsoleted by RFC 8029) time stamp format. The timestamp's **seconds** field holds the integral number of seconds since 1-Jan-1900 00:00:00 UTC. The

timestamp's **microseconds** field contains a microseconds value in the range 0 to 999999. This setting is used to inter-operate with network elements which are fully compliant with RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*, (such as an SR OS system with the same setting, or any other RFC 4379 compliant router).

unix

Specifies the Unix time stamp format. The time stamps **seconds** field holds a Unix time, the integral number of seconds since 1-Jan-1970 00:00:00 UTC. The time stamps **microseconds** field contains a microseconds value in the range 0 to 999999. This setting is used to inter-operate with network elements which send and expect a 1970-based timestamp in MPLS Echo Request/Reply PDUs (such as an SR OS system with the same setting, or an SR OS system running software earlier than R8.0 R4).

Platforms

7705 SAR Gen 2

18.74 mrib

mrib

Syntax

mrib [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]

no mrib

Context

[Tree] (debug>router>pim mrib)

Full Context

debug router pim mrib

Description

This command enables debugging for PIM MRIB.

The **no** form of this command disables debugging for PIM MRIB.

Parameters

grp-ip-address

Debugs information associated with the specified PIM MRIB.

Values multicast group address (ipv4, ipv6)

ip-address

Debugs information associated with the specified PIM MRIB.

Values source address (ipv4, ipv6)

detail

Debugs detailed MRIB information.

Platforms

7705 SAR Gen 2

18.75 mrouter-port

mrouter-port

Syntax

[**no**] **mrouter-port**

Context

[Tree] (config>service>vpls>spoke-sdp>mld-snooping mrouter-port)

[Tree] (config>service>vpls>bind>mld-snooping mrouter-port)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping mrouter-port)

[Tree] (config>service>vpls>sap>igmp-snooping mrouter-port)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping mrouter-port)

Full Context

configure service vpls spoke-sdp mld-snooping mrouter-port

configure service vpls allow-ip-int-bind mld-snooping mrouter-port

configure service vpls spoke-sdp igmp-snooping mrouter-port

configure service vpls sap igmp-snooping mrouter-port

configure service vpls mesh-sdp igmp-snooping mrouter-port

Description

This command specifies whether a multicast router is attached behind this SAP, SDP, or routed VPLS IP interface.

Configuring these objects as an mrouter-port will have a double effect. Firstly, all multicast traffic received on another SAP, SDP, or routed VPLS IP interface will be copied to this SAP, SDP, or routed VPLS IP interface. Secondly, IGMP/MLD reports generated by the system as a result of a router joining or leaving a multicast group, will be sent to this SAP, SDP, or routed VPLS IP interface.

If two multicast routers exist in the network, one of them will become the active querier. While the other multicast router (non-querier) stops sending IGMP queries, it should still receive reports to keep its multicast trees up-to-date. To support this, the mrouter-port should be enabled on all SAPs, SDPs, or routed VPLS IP interfaces connecting to a multicast router.

Note:
The IGMP version to be used for the reports (v1, v2, or v3) can only be determined after an initial query has been received. Until such time, no reports are sent on the SAP, spoke-SDP, or routed VPLS IP interface, even if **mrouter-port** is enabled.

If the **send-queries** command is enabled on this SAP or spoke-SDP, the **mrouter-port** parameter cannot be set.

When PIM-snooping is enabled within a VPLS service, all IP multicast traffic and PIM messages will be sent to any SAP or SDP binding configured with an IGMP-snooping mrouter port. This occurs even without IGMP-snooping enabled, but is not supported in a BGP-VPLS or M-VPLS service.

The **no** form of this command reverts to the default.

Default

no mrouter-port

Platforms

7705 SAR Gen 2

18.76 msap-defaults

msap-defaults

Syntax

msap-default

Context

[Tree] (config>service>vpls>sap msap-defaults)

Full Context

configure service vpls sap msap-defaults

Description

This command configures MSAP authentication defaults.

Platforms

7705 SAR Gen 2

18.77 msg

msg

Syntax

[**no**] **msg**

Context

[Tree] (debug>router>pim msg)

Full Context

debug router pim msg

Description

This command enables debugging for PIM messaging.

The **no** form of this command disables debugging for PIM messaging.

Platforms

7705 SAR Gen 2

msg

Syntax

msg [**detail**]

no msg

Context

[Tree] (debug>router>pcep>pcc>conn msg)

[Tree] (debug>router>pcep>pcc msg)

Full Context

debug router pcep pcc connection msg

debug router pcep pcc msg

Description

This command enables debugging for PCC or connection messaging events.

The **no** form of this command disables debugging.

Parameters

detail

Keyword used to specify detailed information about PCC or connection messaging events.

Platforms

7705 SAR Gen 2

18.78 msg-pacing

msg-pacing

Syntax

[**no**] **msg-pacing**

Context

[Tree] (config>router>rsvp msg-pacing)

Full Context

configure router rsvp msg-pacing

Description

This command enables RSVP message pacing in which the specified number of RSVP messages, specified in the **max-burst** command, are sent in a configured interval, specified in the **period** command. A count is kept of the messages that were dropped because the output queue for the interface used for message pacing was full.

Default

no msg-pacing

Platforms

7705 SAR Gen 2

18.79 mss-adjust-group

mss-adjust-group

Syntax

mss-adjust-group *bb-group-id* **segment-size** *segment-size*

no mss-adjust-group

Context

[Tree] (config>router mss-adjust-group)

[Tree] (config>service>vprn mss-adjust-group)

Full Context

configure router mss-adjust-group

configure service vprn mss-adjust-group

Description

This command associates the MSS adjust group consisting of multiple ISAs with the routing context in which the application requiring TCP MSS adjust resides.

Parameters

bb-group-id

Specifies the group used for TCP MSS adjust

segment-size

Specifies the value to put into the TCP Maximum Segment Size (MSS) option if it is not already present, or if the present value is higher

Values 160 to 10240

Platforms

7705 SAR Gen 2

18.80 mst-instance

mst-instance

Syntax

mst-instance *mst-inst-number*

Context

[Tree] (config>service>vpls>sap>stp mst-instance)

Full Context

configure service vpls sap stp mst-instance

Description

Commands in this context configure MSTI related parameters at SAP level. This context can be open only for existing mst-instances defined at the service level (see **config>service>vpls>stp mst-instance**).

Parameters

mst-inst-number

Specifies an existing Multiple Spanning Tree Instance number.

Values 1 to 4094

Platforms

7705 SAR Gen 2

mst-instance

Syntax

mst-instance *mst-inst-number* [**create**]

no mst-instance [*mst-inst-number*]

Context

[Tree] (config>service>vpls>stp mst-instance)

Full Context

configure service vpls stp mst-instance

Description

This command creates the context to configure MST instance (MSTI) related parameters. Up to 16 instances will be supported by MSTP. The instance 0 is mandatory by protocol and therefore, it cannot be created by the CLI. The software will maintain this instance automatically.

Parameters

mst-inst-number

Specifies the Multiple Spanning Tree instance

Values 1 to 4094

Platforms

7705 SAR Gen 2

18.81 mst-max-hops

mst-max-hops

Syntax

mst-max-hops *hops-count*

no mst-max-hops

Context

[Tree] (config>service>vpls>stp mst-max-hops)

Full Context

configure service vpls stp mst-max-hops

Description

This command specifies the number of hops in the region before BPDU is discarded and the information held for the port is aged out. The root bridge of the instance sends a BPDU (or M-record) with remaining-hop-count set to configured <*max-hops*>. When a bridge receives the BPDU (or M-record), it decrements the received remaining-hop-count by 1 and propagates it in BPDU (or M-record) it generates.

The **no** form of this command sets the *hops-count* to its default value.

Default

mst-max-hops 20

Parameters

hops-count

Specifies the maximum number of hops.

Values 1 to 40

Platforms

7705 SAR Gen 2

18.82 mst-name

mst-name

Syntax

mst-name *region-name*

no mst-name

Context

[Tree] (config>service>vpls>stp mst-name)

Full Context

configure service vpls stp mst-name

Description

This command defines an MST region name. Two bridges are considered as a part of the same MST region as soon as their configuration of the MST region name, the MST-revision and VLAN-to-instance assignment is identical.

The **no** form of this command removes *region-name* from the configuration.

Default

no mst-name

Parameters

region-name

Specifies an MST-region name up to 32 characters in length.

Platforms

7705 SAR Gen 2

18.83 mst-path-cost

mst-path-cost

Syntax

mst-path-cost *inst-path-cost*

no mst-path-cost

Context

[Tree] (config>service>vpls>sap>stp>mst-instance mst-path-cost)

Full Context

configure service vpls sap stp mst-instance mst-path-cost

Description

This commands specifies path-cost within a specified instance, expressing probability that a specified port will be put into the forwarding state in case a loop occurs (the highest value expresses lowest priority).

The **no** form of this command sets port-priority to its default value.

Default

The path-cost is proportional to link speed.

Parameters

inst-path-cost

Specifies the contribution of this port to the MSTI path cost of paths toward the spanning tree regional root that include this port.

Values 1 to 200000000

Platforms

7705 SAR Gen 2

18.84 mst-port-priority

mst-port-priority

Syntax

mst-port-priority *stp-priority*

no mst-port-priority

Context

[Tree] (config>service>vpls>sap>stp>mst-instance mst-port-priority)

Full Context

configure service vpls sap stp mst-instance mst-port-priority

Description

This commands specifies the port priority within a specified instance, expressing probability that a specified port will be put into the forwarding state if a loop occurs.

The **no** form of this command sets port-priority to its default value.

Default

mst-port-priority 128

Parameters

stp-priority

Specifies the value of the port priority field.

Platforms

7705 SAR Gen 2

18.85 mst-priority

mst-priority

Syntax

mst-priority *bridge-priority*

no mst-priority

Context

[Tree] (config>service>vpls>stp>mst-instance mst-priority)

Full Context

configure service vpls stp mst-instance mst-priority

Description

This command specifies the bridge priority for this specific Multiple Spanning Tree Instance for this service. The *bridge-priority* value reflects likelihood that the switch will be chosen as the regional root switch (65535 represents the least likely). It is used as the highest 4 bits of the Bridge ID included in the MSTP BPDUs generated by this bridge.

The priority can only take on values that are multiples of 4096 (4k). If a value is specified that is not a multiple of 4K, then the value will be replaced by the closest multiple of 4K, which is lower than the value entered.

All instances created by the **configure service vpls stp mst-instance vlan-range** command and not having explicit definition of bridge-priority inherit the default value.

The **no** form of this command sets the bridge-priority to its default value.

Default

mst-priority 32768

Parameters

bridge-priority

Specifies the priority of this specific Multiple Spanning Tree Instance for this service.

Values 0 to 65535

Platforms

7705 SAR Gen 2

18.86 mst-revision

mst-revision

Syntax

mst-revision *revision-number*

no mst-revision

Context

[Tree] (config>service>vpls>stp mst-revision)

Full Context

configure service vpls stp mst-revision

Description

This command defines the MST configuration revision number. Two bridges are considered as a part of the same MST region as soon as their configuration of MST-region name, MST-revision and VLAN-to-instance assignment is identical.

The **no** form of this command returns MST configuration revision to its default value.

Default

mst-revision 0

Parameters

revision-number

Specifies the MSTP region revision number to define the MSTP region.

Values 0 to 65535

Platforms

7705 SAR Gen 2

18.87 mtu

mtu

Syntax

mtu *bytes*

no mtu

Context

[Tree] (config>service>vprn>router-advert>if mtu)

[Tree] (config>router>router-advert>if mtu)

Full Context

configure service vprn router-advertisement interface mtu

configure router router-advertisement interface mtu

Description

This command specifies the value to be placed in link MTU options sent by the router on this interface.

The **no** form of this command reverts to the default.

Default

no mtu — The MTU option is not sent in the router advertisement messages.

Parameters

bytes

Specifies the advertised MTU value in bytes for this interface.

Values 1280 to 9800 (for **config**>**router**>**router-advert**>**if** and **config**>**service**>**vprn**>**router-advert**>**if** contexts only)

1280 to 9212 (for **subscriber management** context, **ies** and **vprn service subscriber-interface** contexts)

Platforms

7705 SAR Gen 2

mtu

Syntax

mtu *mtu-bytes*

no mtu

Context

[Tree] (config>port>ethernet mtu)

Full Context

configure port ethernet mtu

Description

This command configures the maximum payload MTU size for an Ethernet port, PPP-enabled port or sub-port and Frame Relay-enabled port or subport. The Ethernet port level MTU parameter indirectly defines

the largest physical packet the port can transmit or the far-end Ethernet port can receive. Packets that cannot be fragmented at egress and exceed the MTU are discarded.

The value specified for the MTU includes the destination MAC address, source MAC address, the Ethertype or Length field and the complete Ethernet payload. The MTU value does not include the preamble, start of frame delimiter or the trailing CRC.

PoS channels use the MTU to define the largest PPP payload a PoS frame may contain. A significant difference between SONET/SDH PoS channel and Ethernet physical MTU values the overhead considered part of the framing method and the overhead considered to be part of the application using the frame. In Ethernet, the preamble, start of frame delimiter and the CRC are considered part of the framing overhead and not part of the frame payload. For a PoS channel, the HDLC framing overhead is not included in the physical MTU; only the PPP and PPP payload are included. If the port mode or encapsulation type is changed, the MTU assumes the default values of the new mode or encapsulation type.

The **no** form of this command restores the default values.

Default

The default MTU value depends on the (sub-)port type, mode and encapsulation and are listed in Table 70: Default MTU Values:

Table 70: Default MTU Values

Type	Mode	Encap Type	Default (Bytes)
10/100, Gig, or 10GigE	Access	null	1514
10/100, Gig, or 10GigE	Access	dot1q	1518
10/100, Gig, or 10GigE	Access	q-in-q	1522
SONET/SDH or TDM	Access	mpls	1506
SONET/SDH or TDM	Access	bcp-null	1518
SONET/SDH or TDM	Access	bcp-dot1q	1522
SONET/SDH or TDM	Access	ipcp	1502
SONET/SDH or TDM	Access	frame-relay	1578
ATM, SONET/SDH or TDM	Access	atm	1524
10/100 or 100FX Ethernet	Network	null	1514
10/100 or 100FX Ethernet	Network	dot1q	1518
SONET/SDH	Network	ppp-auto	1524

Parameters

mtu-bytes

Sets the maximum allowable size of the MTU, expressed as an integer.

Values	512 to 9212	config>port>ethernet
	512 to 9800	config>port>ethernet (for FP4-based connector ports)
	512 to 9208	config>port>sonet-sdh>path
	512 to 9208	config>port>tdm>ds1>channel-group
	512 to 9208	config>port>tdm>ds3
	512 to 9208	config>port>tdm>e1>channel-group
	512 to 9208	config>port>tdm>e3

Platforms

7705 SAR Gen 2

mtu

Syntax

mtu *mtu-bytes*

no mtu

Context

[Tree] (config>service>vprn>ospf3>area>if mtu)

[Tree] (config>service>vprn>ospf>area>if mtu)

Full Context

configure service vprn ospf3 area interface mtu

configure service vprn ospf area interface mtu

Description

This command configures the OSPF packet size used on this interface. If this parameter is not configured, OSPF derives the MTU value from the MTU configured (default or explicitly) in the following contexts:

config>**port**>**ethernet**, **config**>**port**>**sonet-sdh**>**path**, **config**>**port**>**tdm**>**t3-e3**, **config**>**port**>**tdm**>**t1-e1**>**channel-group**

If this parameter is configured, the smaller value between the value configured here and the MTU configured (default or explicitly) in an above-mentioned contexts is used.

To determine the actual packet size, add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured in this command.

The **no** form of this command reverts to default value derived from the MTU configured in the **config>port** context.

Default

no mtu

Parameters

mtu-bytes

Specifies the MTU to be used by OSPF for this logical interface, in bytes.

Values 512 to 9786

Platforms

7705 SAR Gen 2

mtu

Syntax

mtu *value*

no mtu

Context

[Tree] (config>service>vprn>nat>outside mtu)

Full Context

configure service vprn nat outside mtu

Description

This command configures the Maximum Transmission Unit (MTU) for downstream traffic flowing through this router (as outside NAT router). The system fragments IP datagrams exceeding the MTU.

The **no** form of the command reverts to the default.

Default

no mtu

Parameters

value

Specifies the MTU for downstream traffic.

Values 512 to 9000

Platforms

7705 SAR Gen 2

mtu

Syntax

mtu *mtu-size*

no mtu

Context

[Tree] (config>router>nat>outside mtu)

Full Context

configure router nat outside mtu

Description

This command configures the MTU for downstream traffic flowing through this router (as outside NAT router). The system fragments IP datagrams exceeding the MTU.

Default

no mtu

Parameters

mtu-size

Specifies the MTU for downstream traffic.

Values 512 to 9000

Platforms

7705 SAR Gen 2

mtu

Syntax

mtu *bytes*

no mtu

Context

[Tree] (config>router>ospf3>area>interface mtu)

[Tree] (config>router>ospf>area>interface mtu)

Full Context

configure router ospf3 area interface mtu

configure router ospf area interface mtu

Description

This command configures the OSPF packet size used on this interface. If this parameter is not configured OSPF derives the MTU value from the MTU configured (default or explicitly) in the following contexts:

- **config>port>ethernet**
- **config>port>sonet-sdh>path**
- **config>port>tdm>t3-e3**
- **config>port>tdm>t1-e1>channel-group**

If this parameter is configured, the smaller value between the value configured here and the MTU configured (default or explicitly) in an above-mentioned context is used.

To determine the actual packet size, add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured in this command.

The **no** form of this command reverts to the default derived from the MTU configured in the **config>port** context.

Default

no mtu

Parameters

bytes

Specifies the MTU to be used by OSPF for this logical interface in bytes.

Values 512 to 9786 in the **config**>**router**>**ospf**>**area**>**interface** context.

1280 to 9786 in the **config**>**router**>**ospf3**>**area**>**interface** context.

Platforms

7705 SAR Gen 2

18.88 mtu-over-head

mtu-over-head

Syntax

mtu-over-head *mtu-value*

no mtu-over-head

Context

[Tree] (config>service>vprn>pim mtu-over-head)

Full Context

configure service vprn pim mtu-over-head

Description

This commands subtracts the specified value from the MVPN MTU to allow a BIER header to be added without exceeding the network MTU.

Default

no mtu-over-head

Parameters

mtu-value

Specifies the value subtracted from the MVPN MTU.

Values 44, 76, 140, 268, 536

Platforms

7705 SAR Gen 2

18.89 multi-active

multi-active

Syntax

[**no**] **multi-active**

Context

[Tree] (config>isa>tunnel-grp multi-active)

Full Context

configure isa tunnel-group multi-active

Description

This command enables configuring multiple active MS-ISA in the tunnel-group. IPsec traffic will be load balanced to configured active MS-ISAs.

Operational notes:

- A shutdown of group and removal of all existing configured tunnels of the tunnel-group are needed before provisioning command "multi-active".
- If the tunnel-group is admin-up with "multi-active" configured then the configuration of "primary" and "backup" are not allowed.
- The active-mda-number must be =< total number of ISA configured.
  - If active-mda-number is less than total number of ISA configured then the delta number of ISA will become backup ISA.

Default
no multi-active

Platforms
7705 SAR Gen 2

18.90 multi-chassis

multi-chassis

Syntax
multi-chassis

Context
[Tree] (config>redundancy multi-chassis)

Full Context
configure redundancy multi-chassis

Description
Commands in this context configure multi-chassis parameters.

Platforms
7705 SAR Gen 2

18.91 multi-chassis-shunt-interface

multi-chassis-shunt-interface

Syntax
multi-chassis-shunt-interface *ip-int-name* [**create**]
no multi-chassis-shunt-interface *ip-int-name*

Context
[Tree] (config>router>ipsec multi-chassis-shunt-interface)
[Tree] (config>service>vprn>ipsec multi-chassis-shunt-interface)

Full Context

configure router ipsec multi-chassis-shunt-interface

configure service vprn ipsec multi-chassis-shunt-interface

Description

Commands in this context configure a multi-chassis IPsec shunt interface.

The **no** form of this command removes the interface name from the configuration.

Parameters

ip-int-name

Specifies the shunt interface name, up to 32 characters.

create

Keyword used to create the command instance.

Platforms

7705 SAR Gen 2

multi-chassis-shunt-interface

Syntax

multi-chassis-shunt-interface *ip-int-name*

no multi-chassis-shunt-interface *ip-int-name*

Context

[Tree] (config>router>ipsec>mc-shunt-profile>peer multi-chassis-shunt-interface)

[Tree] (config>service>vprn>ipsec>mc-shunt-profile>peer multi-chassis-shunt-interface)

Full Context

configure router ipsec multi-chassis-shunting-profile peer multi-chassis-shunt-interface

configure service vprn ipsec multi-chassis-shunting-profile peer multi-chassis-shunt-interface

Description

This command associates a multi-chassis-shunt-interface for the peer. The specified interface shunts traffic to the peer.

The **no** form of this command removes association from the configuration.

Parameters

ip-int-name

Specifies the shunt interface name, up to 32 characters.

Platforms

7705 SAR Gen 2

18.92 multi-chassis-shunting-profile

multi-chassis-shunting-profile

Syntax

multi-chassis-shunting-profile *name* [**create**]

no multi-chassis-shunting-profile *name*

Context

[Tree] (config>router>ipsec multi-chassis-shunting-profile)

[Tree] (config>service>vprn>ipsec multi-chassis-shunting-profile)

Full Context

configure router ipsec multi-chassis-shunting-profile

configure service vprn ipsec multi-chassis-shunting-profile

Description

Commands in this context configure a multi-chassis IPsec shunting profile.

The **no** form of this command removes the name from the configuration.

Parameters

name

Specifies the profile name of a MC shunting profile, up to 32 characters.

create

Keyword used to create the command instance.

Platforms

7705 SAR Gen 2

multi-chassis-shunting-profile

Syntax

multi-chassis-shunting-profile *name*

no multi-chassis-shunting-profile

Context

[Tree] (config>service>vprn>if multi-chassis-shunting-profile)

[Tree] (config>service>ies>if multi-chassis-shunting-profile)

Full Context

configure service vprn interface multi-chassis-shunting-profile

configure service ies interface multi-chassis-shunting-profile

Description

This command associates an existing multi-chassis IPsec shunting profile with the service interface.

The **no** form of this command removes the name from the configuration.

Parameters

name

Specifies the profile name of a **multi-chassis-shunting profile**, up to 32 characters.

Platforms

7705 SAR Gen 2

18.93 multi-homed-prefix

multi-homed-prefix

Syntax

[**no**] **multi-homed-prefix**

Context

[Tree] (config>router>ospf>lfa multi-homed-prefix)

Full Context

configure router ospf loopfree-alternates multi-homed-prefix

Description

This command enables multi-homed prefix LFA for OSPF routes (IP FRR) and SR-OSPF node SID tunnels.

This feature makes use of the multi-homed prefix model described in RFC 8518 to compute a backup IP next hop using an alternate ABR or ASBR for external prefixes and to an alternate router owner for local anycast prefixes.

This feature further enhances the multi-homed prefix backup path calculation beyond RFC 8518 with the addition of repair tunnels that make use of a PQ node or a P-Q set to reach the alternate exit ABR or ASBR of external prefixes or the alternate owner router of local anycast prefixes.

The computed IP next-hop based backup path is added to OSPF routes of external /32 prefixes (OSPFv2 routes types 3, 4, 5, and 7) and local /32 anycast prefixes in the RTM if the prefix is not protected by base LFA or if the user set leaf preference value to **all**. The user must enable the **ip-fast-reroute** leaf to have these backup paths programmed into the FIB in data path.

The computed IP next hop or repair tunnel based backup path is also programmed for SR-OSPF node SID tunnels of external /32 prefixes and to /32 prefixes in same area as the computing node S and which are advertised by multiple routers (anycast prefix) in both algorithm 0 and flexible-algorithm numbers.

The **no** form of this command disables multi-homed prefix LFA.

Default

no multi-homed-prefix

Platforms

7705 SAR Gen 2

multi-homed-prefix

Syntax

[**no**] **multi-homed-prefix**

Context

[Tree] (config>router>isis>lfa multi-homed-prefix)

Full Context

configure router isis loopfree-alternates multi-homed-prefix

Description

This command enables multihomed prefix LFA for IS-IS routes (IP FRR), SR-ISIS tunnels, and SRv6-ISIS tunnels.

This feature uses the multihomed prefix model described in RFC 8518 to compute a backup IP next hop using an alternate ABR or ASBR for external prefixes and to an alternate router owner for local anycast prefixes.

This feature further enhances the multihomed prefix backup path calculation beyond RFC 8518 with the addition of repair tunnels that make use of a PQ node or a P-Q set to reach the alternate exit ABR or ASBR of external prefixes or the alternate owner router of intra-area anycast prefixes.

The computed IP next hop-based backup path is added to IS-IS routes of external /32 or /128 prefixes and intra-area /32 or /128 anycast prefixes in the RTM if the prefix is not protected by base LFA or if the user set leaf **preference** command option to **all**. The user must enable the **ip-fast-reroute** leaf to have these backup paths programmed into the FIB in datapath.

The computed IP next hop or repair tunnel-based backup path is also programmed for:

1. SR-ISIS node SID tunnels of external /32 IPv4 prefixes and /128 IPv6 prefixes, and node SID tunnels of intra-area /32 IPv4 anycast prefixes and /128 anycast IPv6 prefixes, in both algorithm 0 and flexible-algorithms

2. SRv6-ISIS locator routes and tunnels of external prefixes and of intra-area anycast prefixes of any size, in both algorithm 0 and flexible algorithm numbers

As a result, an SR-TE LSP, an SR-MPLS policy, or an SRv6 policy which uses an SR-ISIS SID or an SRv6-ISIS SID of those same prefixes in its configured or computed SID list benefits from the multi-homed prefix LFA protection.

Once the IP next-hop based multihomed prefix LFA is enabled, the extensions to compute an SR-TE repair tunnel for the multihomed prefix LFA in the case of SR-ISIS and SRv6-ISIS are automatically enabled if the user also enabled TI-LFA or Remote LFA. The computation reuses the SID list of the primary path or of the TI-LFA or Remote LFA backup path of the alternate ABR or ASBR or alternate owner router.

The **no** form of this command disables multihomed prefix LFA.

Default

no multi-homed-prefix

Platforms

7705 SAR Gen 2

18.94 multi-instance

multi-instance

Syntax

[**no**] **multi-instance**

Context

[Tree] (config>router>ospf multi-instance)

Full Context

configure router ospf multi-instance

Description

This command enables OSPF multi-instance RFC 6549, *OSPFv2 Multi-Instance Extensions,* support in the BASE router. This support is enabled per instance and allows flexibility when migrating a specific instance from the classic OSPFv2 to a multi-instance OSPFv2.

The **no** form of this command disables OSPF multi-instance support in the BASE router.

Default

no multi-instance

Platforms

7705 SAR Gen 2

18.95 multi-path

multi-path

Syntax

multi-path

Context

[Tree] (config>service>vprn>bgp multi-path)

Full Context

configure service vprn bgp multi-path

Description

This command configures ECMP multipath parameters to apply to address families that support BGP multipath.

Platforms

7705 SAR Gen 2

multi-path

Syntax

multi-path

Context

[Tree] (config>router>bgp multi-path)

Full Context

configure router bgp multi-path

Description

This command configures ECMP multipath parameters to apply to address families that support BGP multipath.

Platforms

7705 SAR Gen 2

18.96 multi-service-site

multi-service-site

Syntax

[**no**] **multi-service-site** *customer-site-name*

Context

[Tree] (config>service>ies>if>sap multi-service-site)

Full Context

configure service ies interface sap multi-service-site

Description

This command creates a new customer site or edits an existing customer site with the *customer-site-name* parameter. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).

The scheduler policy association with the customer site normally prevents the scheduler policy from being deleted until after the scheduler policy is removed from the customer site. The multi-service-site object generates a log message indicating that the association was deleted due to scheduler policy removal.

When the multi-service customer site is created, an ingress and egress scheduler policy association does not exist. This does not prevent the site from being assigned to a chassis slot or prevent service SAP assignment. After the site has been created, the ingress and egress scheduler policy associations can be assigned or removed at any time.

The **no** form of this command removes the value from the configuration.

Default

n/a — Each customer site must be explicitly created.

Parameters

customer-site-name

Each customer site must have a unique name within the context of the customer. If *customer-site-name* already exists for the customer ID, the CLI context changes to that site name for the purpose of editing the site scheduler policies or assignment. Any modifications made to an existing site affects all SAPs associated with the site. Changing a scheduler policy association may cause new schedulers to be created and existing policers and queues on the SAPs to no longer be orphaned. Existing schedulers on the site may cease to exist, causing policers and queues relying on that scheduler to be orphaned.

If the *customer-site-name* does not exist, it is assumed that an attempt is being made to create a site of that name in the customer ID context. The success of the command execution depends on the following:

The maximum number of customer sites defined for the chassis slot has not been met.

The *customer-site-name* is valid.

The **create** keyword is included in the command line syntax (if the system requires it).

When the maximum number of customer sites has been exceeded a configuration error occurs, the command will not execute and the CLI context will not change.

If the *customer-site-name* is invalid, a syntax error occurs, the command will not execute and the CLI context will not change.

Values Valid names consist of any string, up to 32 characters, composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

multi-service-site

Syntax

multi-service-site *customer-site-name*

no multi-service-site

Context

[Tree] (config>service>vprn>if>sap multi-service-site)

[Tree] (config>service>vpls>sap multi-service-site)

Full Context

configure service vprn interface sap multi-service-site

configure service vpls sap multi-service-site

Description

This command associates the SAP with a *customer-site-name*. If the specified *customer-site-name* does not exist in the context of the service customer ID an error occurs and the command is not executed. If *customer-site-name* exists, the current and future defined queues on the SAP (ingress and egress) attempts to use the scheduler hierarchies created within *customer-site-name* as parent schedulers.

This command is mutually exclusive with the SAP ingress and egress scheduler policy commands. If a scheduler policy has been applied to either the ingress or egress nodes on the SAP, the **multi-service-site** command fails without executing. The locally applied scheduler policies must be removed prior to executing the **multi-service-site** command.

The **no** form of this command removes the SAP from any multi-service customer site the SAP belongs to. Removing the site can cause existing or future policers and queues to enter an orphaned state.

Parameters

customer-site-name

Specifies an existing customer site name, up to 32 characters. If the *customer-site-name* exists and local scheduler policies have not been applied to the SAP, the current and future policers queues defined on the SAP looks for their parent schedulers within the scheduler hierarchies defined in the customer-site-name.

Platforms

7705 SAR Gen 2

multi-service-site

Syntax

multi-service-site *customer-site-name*

no multi-service-site

Context

[Tree] (config>service>epipe>sap multi-service-site)

Full Context

configure service epipe sap multi-service-site

Description

This command associates the SAP with a *customer-site-name*. If the specified *customer-site-name* does not exist in the context of the service customer ID an error occurs and the command will not execute. If *customer-site-name* exists, the current and future defined queues on the SAP (ingress and egress) will attempt to use the scheduler hierarchies created within *customer-site-name* as parent schedulers.

The **no** form of this command removes the SAP from any multi-service customer site the SAP belongs to. Removing the site can cause existing or future queues to enter an orphaned state.

Parameters

customer-site-name

The customer-site-name must exist in the context of the customer-id defined as the service owner. If customer-site-name exists and local scheduler policies have not been applied to the SAP, the current and future queues defined on the SAP will look for their parent schedulers within the scheduler hierarchies defined on customer-site-name.

Values Any valid customer-site-name created within the context of the customer-id.

Platforms

7705 SAR Gen 2

multi-service-site

Syntax

multi-service-site *customer-site-name* [**create**]

no multi-service-site *customer-site-name*

Context

[Tree] (config>service>cust multi-service-site)

Full Context

configure service customer multi-service-site

Description

This command creates a new customer site or edits an existing customer site with the *customer-site-name* parameter. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).

The scheduler policy association with the customer site normally prevents the scheduler policy from being deleted until after the scheduler policy is removed from the customer site. The multi-service-site object will generate a log message indicating that the association was deleted due to scheduler policy removal.

When the multi-service customer site is created, an ingress and egress scheduler policy association does not exist. This does not prevent the site from being assigned to a chassis slot or prevent service SAP assignment. After the site has been created, the ingress and egress scheduler policy associations can be assigned or removed at any time.

Parameters

customer-site-name

Specifies the customer site name. Each customer site must have a unique name within the context of the customer. If *customer-site-name* already exists for the customer ID, the CLI context changes to that site name for the purpose of editing the site scheduler policies or assignment. Any modifications made to an existing site will affect all SAPs associated with the site. Changing a scheduler policy association may cause new schedulers to be created and existing policers and queues on the SAPs to no longer be orphaned. Existing schedulers on the site may cease to exist, causing policers and queues relying on that scheduler to be orphaned.

If the *customer-site-name* does not exist, it is assumed that an attempt is being made to create a site of that name in the customer ID context. The success of the command execution depends on the following:

- The maximum number of customer sites defined for the chassis has not been met.

- The *customer-site-name* is valid.
- The **create** keyword is included in the command line syntax (if the system requires).

When the maximum number of customer sites has been exceeded a configuration error occurs; the command will not execute and the CLI context will not change.

If the *customer-site-name* is invalid, a syntax error occurs; the command will not execute and the CLI context will not change.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

18.97 multi-topology

multi-topology

Syntax

[**no**] **multi-topology**

Context

[Tree] (config>service>vprn>isis multi-topology)

Full Context

configure service vprn isis multi-topology

Description

This command enables IS-IS multi-topology support.

The **no** form of this command disables IS-IS multi-topology.

Default

no multi-topology

Platforms

7705 SAR Gen 2

multi-topology

Syntax

[**no**] **multi-topology**

Context

[Tree] (config>router>isis multi-topology)

Full Context

configure router isis multi-topology

Description

This command enables IS-IS multi-topology support.

Default

no multi-topology

Platforms

7705 SAR Gen 2

multi-topology

Syntax

multi-topology mt2

no multi-topology

Context

[Tree] (config>router>isis>segment-routing multi-topology)

Full Context

configure router isis segment-routing multi-topology

Description

This command configures SR-MPLS for SR-ISIS MT, which enables Segment Routing in MT2.

The **no** form of this command disables Segment Routing in MT2.

Default

no multi-topology

Platforms

7705 SAR Gen 2

18.98 multicast

multicast

Syntax

multicast [**key-id** *key-id* | **authentication-keychain** *keychain-name*] [**version** *version*]

no multicast

Context

[Tree] (config>system>time>ntp multicast)

Full Context

configure system time ntp multicast

Description

This command configures NTP the node to transmit multicast packets on the CPM/CCM MGMT port. Broadcast and multicast messages can easily be spoofed; authentication is strongly recommended.

The **no** form of this command removes the multicast address from the configuration.

Parameters

key-id

Specifies the configured authentication key and authentication type used by this version to transmit NTP packets. If this command is omitted from the configuration, packets are sent unencrypted.

Values 1 to 255

keychain-name

Identifies the keychain name, up to 32 characters.

version

Specifies the NTP version number that is generated by this node. This parameter does not need to be configured when in client mode in which case all three versions are accepted.

Values 2 to 4

Default 4

Platforms

7705 SAR Gen 2

18.99 multicast-import

multicast-import

Syntax

[**no**] **multicast-import**

Context

[Tree] (config>service>vprn>isis multicast-import)

Full Context

configure service vprn isis multicast-import

Description

This command enables ISIS to submit routes into the multicast Route Table Manager (RTM).

The **no** form of this command disables the submission of routes into the multicast RTM.

Default

no multicast-import

Platforms

7705 SAR Gen 2

multicast-import

Syntax

[**no**] **multicast-import**

Context

[Tree] (config>service>vprn>ospf3 multicast-import)

[Tree] (config>service>vprn>ospf multicast-import)

Full Context

configure service vprn ospf3 multicast-import

configure service vprn ospf multicast-import

Description

This command enables the submission of routes into the multicast Route Table Manager (RTM) by OSPF.

The **no** form of this command disables the submission of routes into the multicast RTM.

Default

no multicast-import

Platforms

7705 SAR Gen 2

multicast-import

Syntax

[**no**] **multicast-import** [{**both** | **ipv4** | **ipv6**}]

Context

[Tree] (config>router>isis multicast-import)

Full Context

configure router isis multicast-import

Description

This command enables the submission of routes into the multicast Route Table Manager (RTM) by IS-IS.

The **no** form of this command disables the submission of routes into the multicast RTM.

Default

no multicast-import

Parameters

both

Allows submission of both IPv4 and IPv6 routes.

ipv4

Allows submission of IPv4 routes only.

ipv6

Allows submission of IPv6 routes only.

Platforms

7705 SAR Gen 2

multicast-import

Syntax

[**no**] **multicast-import**

Context

[Tree] (config>router>ospf multicast-import)

[Tree] (config>router>ospf3 multicast-import)

Full Context

configure router ospf multicast-import

configure router ospf3 multicast-import

Description

This command enables the submission of routes into the multicast Route Table Manager (RTM) by OSPF.

The **no** form of this command disables the submission of routes into the multicast RTM.

Default

no multicast-import

Platforms

7705 SAR Gen 2

18.100 multicast-leave-sync-propagation

multicast-leave-sync-propagation

Syntax

multicast-leave-sync-propagation *time*

Context

[Tree] (config>service>system>bgp-evpn multicast-leave-sync-propagation)

Full Context

configure service system bgp-evpn multicast-leave-sync-propagation

Description

This command configures the additional amount of time that the system waits before removing a multicast state that was synchronized in an Ethernet Segment via Multicast Join or Leave Synch routes. This value represents a delta corresponding to the time it takes for a BGP advertisement to propagate to ES peers.

The node triggering the route computes the maximum response time as the product of the locally configured values, Last Member Query Count and Last Member Query Interval (this value is taken from the **config**>**service**>**vpls**>**sap**>**igmp-snooping**>**last-member-query-interval** or **config**>**service**>**vpls**>**spoke-sdp**>**igmp-snooping**> **last-member-query-interval** commands depending on the Ethernet Segment being used), and adds the delta value to the Maximum Response Time. Increasing the Maximum Response Time by this value can help minimize the churn of removing and recreating the state on the node.

The maximum response time value should be configured consistently in all ES peers. For example, in a scenario where a maximum response time of five seconds is advertised by PE-A and there is a delay of four seconds in the BGP propagation to PE-B, the timer could already expire on PE-A while PE-B is still in LMQ time and can still receive joins (which would recreate state in A after a join synch route from B). To minimize this situation, adding an extra delta timer on PE-A, reduces the potential churn of PE-A removing and recreating the state.

Default

multicast-leave-sync-propagation 5

Parameters

time

Specifies the multicast leave sync propagation delay time, in seconds.

Values 0 to 300

Default 5

Platforms

7705 SAR Gen 2

18.101 multicast-network-domain

multicast-network-domain

Syntax

multicast-network-domain *multicast-network-domain*

no multicast-network-domain

Context

[Tree] (config>service>ies>if multicast-network-domain)

Full Context

configure service ies interface multicast-network-domain

Description

This command is used to enable efficient multicast replication over a spoke SDP. Multicast traffic is copied to only a subset of network interfaces that may be used as egress for a spoke SDP. A network domain is defined by associating multiple interfaces to a logical group that may participate in multicast replication for a spoke SDP.

The **no** form of command disables efficient multicast replication to a network domain for a spoke SDP and traffic is replicated to all forwarding complexes.

Default

no multicast-network-domain

Platforms

7705 SAR Gen 2

18.102 multicast-policer

multicast-policer

Syntax

multicast-policer *policer-id* [**fp-redirect-group**]

no multicast-policer

Context

[Tree] (config>qos>sap-ingress>fc multicast-policer)

Full Context

configure qos sap-ingress fc multicast-policer

Description

Within a **sap-ingress** QoS policy forwarding class context, the **multicast-policer** command is used to map packets that match the forwarding class and are considered multicast in nature to the specified policer-id. The specified policer-id must already exist within the **sap-ingress** QoS policy. While the system is determining the forwarding class of a packet, it is also looking up its forwarding destination based on the ingress service type and the service instance forwarding records. Two basic types of services support multicast packets: routed services (IES and VPRN) and L2 multipoint services (VPLS, I-VPLS, and B-VPLS). For the routed service types, a multicast packet is destined to an IPv4 or IPv6 multicast address. For the L2 multipoint services, a multicast packet is a packet destined to a multicast MAC address (multicast bit set in the destination MAC address but not the ff:ff:ff:ff:ff:ff broadcast address). The VPLS services also support two other multipoint forwarding types (broadcast and unknown), which are considered separate from the multicast forwarding type.

If ingress forwarding logic has resolved a packet to the multicast forwarding type within the forwarding class, it will be mapped to either an ingress multipoint queue (using the **multicast** *queue-id* or **multicast** *queue-id* **group** *ingress-queue-group* commands) or an ingress policer (**multicast-policer** *policer-id*). The **multicast** and **multicast-policer** commands within the forwarding class context are mutually exclusive. By default, the multicast forwarding type is mapped to the SAP ingress default multipoint queue. If the **multicast-policer** *policer-id* command is executed, any previous policer mapping or queue mapping for the multicast forwarding type within the forwarding class is overridden if the policer mapping is successful.

A policer defined within the **sap-ingress** policy is not actually created on an ingress SAP or a subscriber using an **sla-profile** where the policy is applied until at least one forwarding type (unicast, broadcast, unknown, or multicast) from one of the forwarding classes is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or multiservice site, or ingress policing is not

supported on the port associated with the SAP or subscriber or multiservice site, the initial forwarding class forwarding type mapping will fail.

The **multicast-policer** command is ignored for instances of the policer applied to SAPs subscribers or multiservice site where broadcast packets are not supported.

When the multicast forwarding type within a forwarding class is mapped to a policer, the multicast packets classified to the subclasses within the forwarding class are also mapped to the policer.

The **no** form of this command is used to restore the mapping of the multicast forwarding type within the forwarding class to the default multipoint queue. If all forwarding class forwarding types had been removed from the default multipoint queue, the queue will not exist on the SAPs subscribers or multiservice site associated with the QoS policy, and the **no multicast-policer** command will cause the system to attempt to create the default multipoint queue on each object. If the system cannot create the queue on each instance, the **no multicast-policer** command will fail and the multicast forwarding type within the forwarding class will continue its mapping to the existing *policer-id*. If the **no multicast-policer** command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscriber will be lost.

Parameters

policer-id

When the forwarding class **multicast-policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the **sap-ingress** QoS policy.

Values 1 to 63

fp-redirect-group

Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

Platforms

7705 SAR Gen 2

18.103 multicast-queue

multicast-queue

Syntax

multicast-queue *queue-id* [**group** *queue-group-name*]

no multicast-queue

Context

[Tree] (config>qos>sap-ingress>fc multicast-queue)

Full Context

configure qos sap-ingress fc multicast-queue

Description

This command overrides the default multicast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. When the forwarding class mapping is executed, all multicast traffic on a SAP using this policy is forwarded using the *queue-id*.

The multicast forwarding type includes the **unknown** unicast forwarding type and the **broadcast** forwarding type unless each is explicitly defined to a different multipoint queue. When the unknown and broadcast forwarding types are left as default, they will track the defined queue for the multicast forwarding type.

The **no** form of this command sets the multicast forwarding type *queue-id* back to the default queue for the forwarding class. If the **broadcast** and **unknown** forwarding types were not explicitly defined to a multipoint queue, they will also be set back to the default multipoint queue (queue 11).

Parameters

queue-id

The *queue-id* parameter specified must be an existing, multipoint queue defined in the config>qos>sap-ingress context.

Values Any valid multipoint queue-ID in the policy including 2 through 32.

Default 11

group *queue-group-name*

This optional parameter is used to redirect the forwarding type within the forwarding class to the specified queue-id within the queue-group-name. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the specified queue group. The *queue-group-name* are configured in the **config>qos>queue-group-templates** egress and ingress contexts.

Platforms

7705 SAR Gen 2

multicast-queue

Syntax

multicast-queue *queue-id*

no multicast-queue

Context

[Tree] (config>qos>network-queue>fc multicast-queue)

Full Context

configure qos network-queue fc multicast-queue

Description

This command overrides the default multicast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. When the forwarding class mapping is executed, all multicast traffic using this policy is forwarded using the *queue-id*.

The multicast forwarding type includes the **unknown** unicast forwarding type and the **broadcast** forwarding type, unless each is explicitly defined to a different multipoint queue. When the unknown and broadcast forwarding types are left as default, they will track the defined queue for the multicast forwarding type.

The **no** form of this command sets the multicast forwarding type *queue-id* back to the default queue for the forwarding class. If the **broadcast** and **unknown** forwarding types were not explicitly defined to a multipoint queue, they will also be set back to the default multipoint queue (queue 11).

Resource Utilization

When a multipoint queue is created and at least one forwarding class is mapped to the queue using the **multipoint-queue** command, a single ingress multipoint hardware queue is created per instance of the applied network-queue policy, using the queue-policy command at the ingress network FP level. Multipoint queues are not created at egress and the multipoint queues defined in the network-queue policy are ignored when the policy is applied to an egress port.

Parameters

queue-id

Specifies any valid multipoint queue-ID in the policy. The *queue-id* parameter specified must be an existing, multipoint queue defined in the **config>qos>network-queue>queue** context.

Values 1 to 16

Default 11

Platforms

7705 SAR Gen 2

18.104 multicast-senders

multicast-senders

Syntax

multicast-senders {**auto** | **always** | **never**}

no multicast-senders

Context

[Tree] (config>service>vprn>pim>if multicast-senders)

Full Context

configure service vprn pim interface multicast-senders

Description

This command configures the way subnet matching is done for incoming data packets on this interface. An IP multicast sender is an user entity to be authenticated in a receiving host.

Parameters

auto

Subnet matching is automatically performed for incoming data packets on this interface.

always

Subnet matching is always performed for incoming data packets on this interface.

never

Subnet matching is never performed for incoming data packets on this interface.

Platforms

7705 SAR Gen 2

multicast-senders

Syntax

multicast-senders {**auto** | **always** | **never**}

no multicast-senders

Context

[Tree] (config>router>pim>interface multicast-senders)

Full Context

configure router pim interface multicast-senders

Description

This command configures how traffic from directly-attached multicast sources should be treated on broadcast interfaces. It can also be used to treat all traffic received on an interface as traffic coming from a directly-attached multicast source. This is particularly useful if a multicast source is connected to a point-to-point or unnumbered interface.

The **no** form of this command reverts to the default value.

Default

multicast-senders auto

Parameters

auto

Specifies that, on broadcast interfaces, the forwarding plane performs subnet-match check on multicast packets received on the interface to determine if the packet is from a directly-attached source. On unnumbered/point-to-point interfaces, all traffic is implicitly treated as coming from a remote source.

always

Treats all traffic received on the interface as coming from a directly-attached multicast source.

never

Specifies that, on broadcast interfaces, traffic from directly-attached multicast sources will not be forwarded; however, traffic from a remote source will still be forwarded if there is a multicast state for it. On unnumbered/point-to-point interfaces, it means that all traffic received on that interface must not be forwarded.

Platforms

7705 SAR Gen 2

18.105 multicastclient

multicastclient

Syntax

multicastclient [**authenticate**]

no multicastclient

Context

[Tree] (config>system>time>ntp multicastclient)

Full Context

configure system time ntp multicastclient

Description

This command configures the node to receive multicast NTP messages on the CPM MGMT port. If **multicastclient** is not configured, received NTP multicast traffic will be ignored. Use the **show** command to view the state of the configuration.

The **no** construct of this message removes the multicast client for the specified interface from the configuration.

Parameters

authenticate

Specifies to make authentication a requirement (optional). If authentication is required, the authentication key-id received must have been configured in the **authentication-key** command, and that key-id type and key value must also match.

Platforms

7705 SAR Gen 2

18.106 multihop

multihop

Syntax

multihop *ttl-value*

no multihop

Context

[Tree] (config>service>vprn>bgp>group multihop)

[Tree] (config>service>vprn>bgp>group>neighbor multihop)

[Tree] (config>service>vprn>bgp multihop)

Full Context

configure service vprn bgp group multihop

configure service vprn bgp group neighbor multihop

configure service vprn bgp multihop

Description

This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGP peer multiple hops away.

This parameter is meaningful only when configuring EBGP peers. It is ignored if set for an IBGP peer.

The **no** form of this command is used to convey to the BGP instance that the EBGP peers are directly connected.

The **no** form of this command reverts to default values.

Default

multihop 1 (EBGP peers are directly connected)

multihop 64 (IBGP)

Parameters

ttl-value

Specifies the TTL value, expressed as a decimal integer.

Values 1 to 255

Platforms

7705 SAR Gen 2

multihop

Syntax

multihop *ttl-value*

no multihop

Context

[Tree] (config>router>bgp>group multihop)

[Tree] (config>router>bgp>group>neighbor multihop)

[Tree] (config>router>bgp multihop)

Full Context

configure router bgp group multihop

configure router bgp group neighbor multihop

configure router bgp multihop

Description

This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGP peer multiple hops away.

The **no** form of this command is used to convey to the BGP instance that the EBGP peers are directly connected.

The **no** form of this command used at the global level reverts to default.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

multihop 1 — EBGP peers are directly connected.

multihop 64 — IBGP

Parameters

ttl-value

Specifies the TTL value, expressed as a decimal integer.

Values 1 to 255

Platforms

7705 SAR Gen 2

18.107 multipath-eligible

multipath-eligible

Syntax

[**no**] **multipath-eligible**

Context

[Tree] (config>service>vprn>bgp>group multipath-eligible)

Full Context

configure service vprn bgp group multipath-eligible

Description

This command specifies that a BGP neighbor or the set of BGP neighbors in a peer group should be part of a selective multipath set. Selective multipaths are only supported by the ipv4, label-ipv4, ipv6, and label-ipv6 address families.

If no candidate multipath route for an IP prefix came from a multipath-eligible peer, multipaths are selected without further constraints.

If the best route for an IP prefix is received from a neighbor marked as multipath-eligible, other routes for the same prefix are not eligible to be used as multipaths unless they also came from peers marked as multipath-eligible.

If the best route for an IP prefix did not come from a multipath-eligible peer but there is at least one candidate multipath route for the same prefix from a multipath-eligible peer, multipath is not used.

The **no** form of this command marks a neighbor or group as non-multipath eligible. The effect of this depends on whether other neighbors and groups are marked as multipath eligible.

Default

no multipath-eligible

Platforms

7705 SAR Gen 2

multipath-eligible

Syntax

[**no**] **multipath-eligible**

Context

[Tree] (config>router>bgp>group>neighbor multipath-eligible)

[Tree] (config>router>bgp>group multipath-eligible)

Full Context

configure router bgp group neighbor multipath-eligible

configure router bgp group multipath-eligible

Description

This command specifies that a BGP neighbor or the set of BGP neighbors in a peer group should be part of a selective multipath set. Selective multipaths are only supported by the **ipv4**, **label-ipv4**, **ipv6**, and **label-ipv6** address families.

If no candidate multipath route for an IP prefix came from a multipath-eligible peer then multipaths are selected without further constraints.

If the best route for an IP prefix is received from a neighbor marked as multipath-eligible, then other routes for the same prefix are not eligible to be used as multipaths unless they also came from peers marked as multipath-eligible.

If the best route for an IP prefix did not come from a multipath-eligible peer but there is at least one candidate multipath route for the same prefix from a multipath-eligible peer then multipath is not used.

The **no** form of this command marks a neighbor or group as non-multipath eligible. The effect of this depends on whether other neighbors and groups are marked as multipath eligible.

Default

no multipath-eligible

Platforms

7705 SAR Gen 2

18.108 multiple-option

multiple-option

Syntax

multiple-option {**true** | **false**}

no multiple-option

Context

[Tree] (config>filter>ip-filter>entry>match multiple-option)

Full Context

configure filter ip-filter entry match multiple-option

Description

This command configures matching packets that contain one or more than one option fields in the IP header as an IP filter match criterion.

The **no** form of the command removes the checking of the number of option fields in the IP header as a match criterion.

Default

no multiple-option

Parameters

true

Specifies matching on IP packets that contain more than one option field in the header.

false

Specifies matching on IP packets that do not contain multiple option fields present in the header.

Platforms

7705 SAR Gen 2

18.109 multiplier

multiplier

Syntax

multiplier *multiplier*

no multiplier

Context

[Tree] (config>router>bfd>bfd-template multiplier)

Full Context

configure router bfd bfd-template multiplier

Description

This command specifies the detect multiplier for a BFD session. If a BFD control packet is not received for a period of *multiplier* x *receive-interval* (the parameter value of the **receive-interval** command), the session is declared down.

The **no** form of this command reverts to the default value.

Default

multiplier 3

Parameters

multiplier

Specifies the multiplier.

Values 3 to 20

Default 3

Platforms

7705 SAR Gen 2

18.110 mvpn-rtcache

mvpn-rtcache

Syntax

mvpn-rtcache [**group** *grp-ip-address*] [**peer** *ip-address*]

no mvpn-rtcache

Context

[Tree] (debug>router>pim mvpn-rtcache)

Full Context

debug router pim mvpn-rtcache

Description

This command enables debugging for the PIM MVPN route cache.

The **no** form of this command disables debugging for the PIM MVPN route cache.

Parameters

grp-ip-address

Debugs information associated with the specified group.

Values multicast group address (ipv4, ipv6) or zero

peer-ip-address

Debugs information associated with the specified peer.

Values peer address (ipv4, ipv6)

Platforms

7705 SAR Gen 2

18.111 mvrp-control

mvrp-control

Syntax

[**no**] **mvrp-control**

Context

[Tree] (config>service>vpls>vpls-group mvrp-control)

Full Context

configure service vpls vpls-group mvrp-control

Description

This command enables MVRP control in the VPLS instances instantiated using the templates for the specified vpls-group. That means the flooding FDB will be created empty and will be populated with endpoints whenever MVRP receives a declaration and a registration on a specific endpoint. Also the VLAN ID associated by the control VPLS with the instantiated VPLS will be declared on service activation by MVRP on all virtual MVRP ports in the control VPLS. Service activation takes place when at least one other SAP is provisioned and brought up under the data VPLS. This is usually a customer facing SAP or a SAP leading outside of the MVRP controlled domain.

The **no** form of this command disallows MVRP control over this VPLS. The VPLS will be created with a regular FDB and will become as a result active upon creation time. Command change is allowed only when the related vpls-group is in shutdown state.

Default

no mvrp-control

Platforms

7705 SAR Gen 2

19 n Commands

19.1 nak-non-matching-subnet

nak-non-matching-subnet

Syntax

[**no**] **nak-non-matching-subnet**

Context

[Tree] (config>service>vprn>dhcp>server>pool nak-non-matching-subnet)

[Tree] (config>router>dhcp>server>pool nak-non-matching-subnet)

Full Context

configure service vprn dhcp local-dhcp-server pool nak-non-matching-subnet

configure router dhcp local-dhcp-server pool nak-non-matching-subnet

Description

When this command is enabled, if the local DHCPv4 server receives a DHCP request with option 50 (client requested a previously allocated message as described in section 3.2 of RFC 2131, *Dynamic Host Configuration Protocol*) and the address allocation algorithm uses a pool that does not have option 50, the system returns a DHCP NAK. Otherwise, the system drops the DHCP packet.

The **no** form of this command reverts to the default.

Default

no nak-non-matching-subnet

Platforms

7705 SAR Gen 2

19.2 name

name

Syntax

name *system-name*

no name

Context

[Tree] (config>system name)

Full Context

configure system name

Description

This command creates a system name string for the device.

For example, system-name parameter ALA-1 for the **name** command configures the device name as ALA-1.

```
ABC>config>system# name "ALA-1"
ALA-1>config>system#
```

Only one system name can be configured. If multiple system names are configured, the last one encountered overwrites the previous entry.

The **no** form of the command reverts to the default value.

Default

no name

Parameters

system-name

Specifies the system name as a character string. The string may be up to 64 characters. Any printable, seven-bit ASCII characters can be used within the string. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

name

Syntax

name *name-string* **value** *value-string*

name *name-string* **address** *ip-address*

name *name-string* **decimal** *decimal*

name *name-string* **number** *value-number*

name *name-string* **prefix** *ip-prefix/ip-prefix-length*

no name *name-string*

Context

[Tree] (config>router>policy-options>policy-statement>entry>from>policy-variables name)

[Tree] (config>router>policy-options>global-variables name)

Full Context

configure router policy-options policy-statement entry from policy-variables name

configure router policy-options global-variables name

Description

This command configures routing policies that are often reused across BGP peers of a common type (transit, peer, customer, and so on). Using global variables allows a user to have a single variable that is consistent across all peers of a type, while retaining the flexibility to reference different policy functions (prefixes, prefix-lists, community lists, and so on) with unique names.

Depending on the parameter referenced, specify the correct type as follows:

- value-string: **as-path, as-path-group, community, prefix-list, damping**
- ip-address: **next-hop**
- value-number: **aigp-metric, as-path-prepend, local-preference, metric, origin, origin-validation, preference, tag, type**

The **no** form of this command removes the global variable.

Parameters

name-string

Specifies the name of the global variable, with the variable delimited by at-signs (@) at the beginning and the end of the name. Allowed values are any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

value-string

The value of the policy variable. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

value-number

Specifies the numerical value of the policy variable.

Values 0 to 4294967295

ip-address

Specifies the IP address of the policy variable.

Values		
	ipv4-address	a.b.c.d
	ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x - [0 to FFFF]H

d - [0 to 255]D

decimal

Specifies the decimal value of the policy variable.

Values 0.000 to 4294967295.000

ip-prefix/ip-prefix-length

Specifies the IP prefix and prefix length of the policy variable.

Values		
	ip-prefix/ip-prefix-length	*ipv4-prefix*/*ipv4-prefix-length* | *ipv6-prefix*/*ipv6-prefix-length*
	ipv4-prefix	a.b.c.d (host bits must be 0)
	ipv4-prefix-length	[0 to 32]
	ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x - [0 to FFFF]H
		d - [0 to 255]D
	ipv6-prefix-length	[0 to 128]

Platforms

7705 SAR Gen 2

19.3 nas-identifier

nas-identifier

Syntax

[**no**] **nas-identifier**

Context

[Tree] (config>ipsec>rad-auth-plcy>include nas-identifier)

[Tree] (config>ipsec>rad-acct-plcy>include nas-identifier)

Full Context

configure ipsec radius-authentication-policy include-radius-attribute nas-identifier

configure ipsec radius-accounting-policy include-radius-attribute nas-identifier

Description

This command enables the generation of the **nas-identifier** RADIUS attribute.

Default

no nas-identifier

Platforms

7705 SAR Gen 2

19.4 nas-ip-addr

nas-ip-addr

Syntax

[**no**] **nas-ip-addr**

Context

[Tree] (config>ipsec>rad-auth-plcy>include nas-ip-addr)

[Tree] (config>ipsec>rad-acct-plcy>include nas-ip-addr)

Full Context

configure ipsec radius-authentication-policy include-radius-attribute nas-ip-addr

configure ipsec radius-accounting-policy include-radius-attribute nas-ip-addr

Description

This command enables the generation of the NAS IP address attribute.

Default

no nas-ip-addr

Platforms

7705 SAR Gen 2

19.5 nas-port-id

nas-port-id

Syntax

[**no**] **nas-port-id**

Context

[Tree] (config>ipsec>rad-acct-plcy>include nas-port-id)

[Tree] (config>ipsec>rad-auth-plcy>include nas-port-id)

Full Context

configure ipsec radius-accounting-policy include-radius-attribute nas-port-id

configure ipsec radius-authentication-policy include-radius-attribute nas-port-id

Description

This command enables the generation of the **nas-port-id** RADIUS attribute. Optionally, the value of this attribute (the SAP-id) can be prefixed by a fixed string and suffixed by the circuit-id or the remote-id of the client connection. If a suffix is configured, but no corresponding data is available, the suffix used will be 0/0/0/0/0/0.

Default

no nas-port-id

Platforms

7705 SAR Gen 2

19.6 nat

nat

Syntax

[**no**] **nat**

Context

[Tree] (config>router nat)

[Tree] (config>service>vprn nat)

Full Context

configure router nat

configure service vprn nat

Description

This command enables a NAT instance for the specified router or service.

The **no** form of this command disables the NAT instance.

Platforms

7705 SAR Gen 2

nat

Syntax

nat [**nat-policy** *nat-policy-name*]

Context

[Tree] (config>filter>ip-filter>entry>action nat)

Full Context

configure filter ip-filter entry action nat

Description

This command enables NAT traffic diversion based on IPv4 filters (LSN44) or IPv6 filters (DS-Lite, NAT64). The filter contains a matching condition based on any combination of the 5 tuple. Traffic is diverted to NAT based on such defined matching condition. Filter fields outside of the 5 tuples are not valid and it will be ignored in filter based traffic diversion to NAT.

The pool selection for the outside IP address and port along with other mapping characteristics can be specified by the means on the NAT policy.

Parameters

nat-type

Specifies the NAT type.

nat-policy-name

Specifies the NAT policy name, up to 32 characters.

Platforms

7705 SAR Gen 2

19.7 nat-group

nat-group

Syntax

nat-group *nat-group-id* [**create**]

no nat-group *nat-group-id*

Context

[Tree] (config>isa nat-group)

Full Context

configure isa nat-group

Description

This command configures an ISA NAT group.

The **no** form of the command removes the ID from the configuration.

Parameters

nat-group-id

Specifies the ISA NAT group ID.

Values 1 to 4

create

Keyword used to create the NAT group.

Platforms

7705 SAR Gen 2

19.8 nat-policy

nat-policy

Syntax

nat-policy *nat-policy-name*

no nat-policy

Context

[Tree] (config>router>nat>inside nat-policy)

[Tree] (config>service>vprn>nat>inside nat-policy)

Full Context

configure router nat inside nat-policy

configure service vprn nat inside nat-policy

Description

This command configures the NAT policy that is used for large-scale NAT in this service. If a **nat-policy** is not configured, then the default **nat-policy** is used.

The **no** form of the command removes the policy name from the configuration.

Parameters

nat-policy-name

Specifies the NAT policy name, up to 32 characters.

Platforms

7705 SAR Gen 2

nat-policy

Syntax

nat-policy *nat-policy-name* [**create**]

no nat-policy *nat-policy-name*

Context

[Tree] (config>service>nat nat-policy)

Full Context

configure service nat nat-policy

Description

This command configures a NAT policy.

Parameters

nat-policy-name

Specifies the NAT policy name, up to 32 characters.

create

Keyword used to create the NAT policy.

Platforms
7705 SAR Gen 2

19.9 nat-port-forwarding

nat-port-forwarding

Syntax
nat-port-forwarding

Context
[Tree] (config>system>persistence nat-port-forwarding)

Full Context
configure system persistence nat-port-forwarding

Description
This command configures NAT port forwarding persistence parameters.

Platforms
7705 SAR Gen 2

19.10 nat-traversal

nat-traversal

Syntax
nat-traversal [**force**] [**keep-alive-interval** *keep-alive-interval*] [**force-keep-alive**]
no nat-traversal

Context
[Tree] (config>ipsec>ike-policy nat-traversal)

Full Context
configure ipsec ike-policy nat-traversal

Description
This command specifies whether NAT-T (Network Address Translation Traversal) is enabled, disabled or in forced mode.

The **no** form of this command reverts the parameters to the default.

Default

no nat-traversal

Parameters

force

Forces to enable NAT-T

keep-alive-interval *keep-alive-interval*

Specifies the keep-alive interval in seconds.

Values 120 to 600

force-keep-alive

When specified, the keep-alive does not expire.

Platforms

7705 SAR Gen 2

19.11 nbr

nbr

Syntax

nbr [**detail**]

no nbr

Context

[Tree] (debug>router>rsvp>event nbr)

Full Context

debug router rsvp event nbr

Description

This command debugs neighbor events.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about neighbor events.

Platforms
7705 SAR Gen 2

19.12 nd

nd

Syntax
nd

Context
[Tree] (config>service>vprn>if>vpls>evpn nd)
[Tree] (config>service>ies>if>vpls>evpn nd)

Full Context
configure service vprn interface vpls evpn nd
configure service ies interface vpls evpn nd

Description
Commands in this context configure ND host route parameters.

Platforms
7705 SAR Gen 2

19.13 nd-host-route

nd-host-route

Syntax
nd-host-route

Context
[Tree] (config>service>vprn>if>ipv6 nd-host-route)

Full Context
configure service vprn interface ipv6 nd-host-route

Description

Commands in this context populate ND host route entries.

Platforms

7705 SAR Gen 2

19.14 nd-learn-unsolicited

nd-learn-unsolicited

Syntax

nd-learn-unsolicited {**global** | **link-local** | **both**}

no nd-learn-unsolicited

Context

[Tree] (config>service>ies>if>ipv6 nd-learn-unsolicited)

Full Context

configure service ies interface ipv6 nd-learn-unsolicited

Description

This command enables the ability to learn neighbor entries out of received unsolicited Neighbor Advertisement messages with or without the solicited flag set. The command can be enabled for global addresses, link-local addresses, or for both.

The **no** form of this command makes the router use standard RFC 4861 behavior, as described below, for learning of neighbor entries.

- If an unsolicited NA, regardless of the S flag, is received from a neighbor that is not yet in the ND cache, the NA is ignored.
- If an NS, RS, RA, or Redirect message with a Link Layer Address (MAC) is received from a neighbor that is not yet in the ND cache, a new neighbor entry is created in the cache to store the received Link Layer MAC. The neighbor is put in the stale state.

Parameters

global

Learns global neighbor entries out of received unsolicited Neighbor Advertisement messages.

link-local

Learns link local neighbor entries out of received unsolicited Neighbor Advertisement messages.

both

Learns both global and link local neighbor entries out of received unsolicited Neighbor Advertisement messages.

Platforms

7705 SAR Gen 2

nd-learn-unsolicited

Syntax

nd-learn-unsolicited {**global** | **link-local** | **both**}

no nd-learn-unsolicited

Context

[Tree] (config>service>vprn>if>ipv6 nd-learn-unsolicited)

Full Context

configure service vprn interface ipv6 nd-learn-unsolicited

Description

This command enables the ability to learn neighbor entries out of received unsolicited Neighbor Advertisement messages, with or without the solicited flag set. The command can be enabled for global addresses, link-local addresses, or for both.

The **no** form of this command makes the router follow standard RFC 4861 behavior for learning of neighbor entries.

- If an unsolicited NA (regardless of the S flag) is received from a neighbor that is not yet in the ND cache, the NA is ignored in line with RFC 4861.
- If an NS, RS, RA, or Redirect message with a Link Layer Address (MAC) is received from a neighbor that is not yet in the ND cache, a new neighbor entry is created in the cache to store the received Link Layer MAC. The neighbor is put in the STALE state. This is the standard RFC behavior.

Parameters

global

Learns global neighbor entries out of received unsolicited Neighbor Advertisement messages.

link-local

Learns link local neighbor entries out of received unsolicited Neighbor Advertisement messages.

both

Learns both global and link local neighbor entries out of received unsolicited Neighbor Advertisement messages.

Platforms

7705 SAR Gen 2

nd-learn-unsolicited

Syntax

nd-learn-unsolicited {**global** | **link-local** | **both**}

no nd-learn-unsolicited

Context

[Tree] (config>router>if>ipv6 nd-learn-unsolicited)

Full Context

configure router interface ipv6 nd-learn-unsolicited

Description

This command enables the ability to learn neighbor entries out of received unsolicited Neighbor Advertisement messages, with or without the solicited flag set. The command can be enabled for global addresses, link-local addresses, or for both.

The **no** form of this command makes the router follow standard RFC 4861 behavior for learning of neighbor entries.

- If an unsolicited NA (regardless of the S flag) is received from a neighbor that is not yet in the ND cache, the NA is ignored in line with RFC 4861.
- If an NS, RS, RA, or Redirect message with a Link Layer Address (MAC) is received from a neighbor that is not yet in the ND cache, a new neighbor entry is created in the cache to store the received Link Layer MAC. The neighbor is put in the STALE state. This is the standard RFC behavior.

Parameters

global

Learns global neighbor entries out of received unsolicited Neighbor Advertisement messages. This parameter is relevant only to global IPv6 addresses.

link-local

Learns link local neighbor entries out of received unsolicited Neighbor Advertisement messages.

both

Learns both global and link local neighbor entries out of received unsolicited Neighbor Advertisement messages.

Platforms

7705 SAR Gen 2

19.15 nd-proactive-refresh

nd-proactive-refresh

Syntax

nd-proactive-refresh {**global** | **link-local** | **both**}

no nd-proactive-refresh

Context

[Tree] (config>service>ies>if>ipv6 nd-proactive-refresh)

Full Context

configure service ies interface ipv6 nd-proactive-refresh

Description

This command enables a proactive refresh of the neighbor entries. When enabled, at the stale timer expiration, the router sends a NUD message to the host (regardless of the existence of traffic to the IP address on the IOM), so the entry can be refreshed or removed.

This behavior is different from ARP, where the refresh is sent 30 seconds prior to the entry's age out time. The refresh can be optionally enabled for global addresses, link-local addresses, or both.

The **no** form of this command disables the proactive behavior and the router only refreshes an entry if there is traffic that needs to be sent to the IP address.

Parameters

global

Refreshes global neighbor entries.

link-local

Refreshes link local neighbor entries.

both

Refreshes both global and link local neighbor entries.

Platforms

7705 SAR Gen 2

nd-proactive-refresh

Syntax

nd-proactive-refresh {**global** | **link-local** | **both**}

no nd-proactive-refresh

Context

[Tree] (config>service>vprn>if>ipv6 nd-proactive-refresh)

Full Context

configure service vprn interface ipv6 nd-proactive-refresh

Description

This command enables a proactive refresh of the neighbor entries. When enabled, at the stale timer expiration, the router sends an NUD message to the host (regardless of the existence of traffic to the IP address on the IOM), so the entry can be refreshed or removed.

This behavior is different from ARP, where the refresh is sent 30 seconds prior to the entry's age out time. The refresh can be optionally enabled for global addresses, link-local addresses, or both.

The **no** form of this command disables the proactive behavior and the router only refreshes an entry if there is traffic that needs to be sent to the IP address.

Parameters

global

Refreshes global neighbor entries. This parameter is relevant only to global IPv6 addresses.

link-local

Refreshes link local neighbor entries. This parameter is relevant only to global IPv6 addresses.

both

Refreshes both global and link local neighbor entries. This parameter is relevant only to global IPv6 addresses.

Platforms

7705 SAR Gen 2

nd-proactive-refresh

Syntax

nd-proactive-refresh {**global** | **link-local** | **both**}

no nd-proactive-refresh

Context

[Tree] (config>router>if>ipv6 nd-proactive-refresh)

Full Context

configure router interface ipv6 nd-proactive-refresh

Description

This command enables a proactive refresh of the neighbor entries. When enabled, at the stale timer expiration, the router sends an NUD message to the host (regardless of the existence of traffic to the IP address on the IOM), so the entry can be refreshed or removed.

This behavior is different from ARP, where the refresh is sent 30 seconds prior to the entry's age out time. The refresh can be optionally enabled for global addresses, link-local addresses, or both.

The **no** form of this command disables the proactive behavior and the router only refreshes an entry if there is traffic that needs to be sent to the IP address.

Parameters

global

Refreshes global neighbor entries. This parameter is relevant only to global IPv6 addresses.

link-local

Refreshes link local neighbor entries. This parameter is relevant only to global IPv6 addresses.

both

Refreshes both global and link local neighbor entries. This parameter is relevant only to global IPv6 addresses.

Platforms

7705 SAR Gen 2

19.16 nd-router-preference

nd-router-preference

Syntax

nd-router-preference {**medium** | **high** | **low**}

no nd-router-preference

Context

[Tree] (config>service>vprn>router-advert>if nd-router-preference)

[Tree] (config>router>router-advert>if nd-router-preference)

Full Context

configure service vprn router-advertisement interface nd-router-preference

configure router router-advertisement interface nd-router-preference

Description

This command configures the default router preference for Router Advertisement (RA) and allows IPv6 hosts to discover and select a default gateway address by listening to RAs.

This feature provides basic traffic engineering functionality for host devices. When this command is applied, the router advertises the respective router preference to the connected host to assist in its selection of the most appropriate default gateway on a link.

This extension is backward compatible, both for routers (setting the router preference bits) and hosts (interpreting the router preference bits). These bits are ignored by hosts that do not implement the RFC 4191 functionality by configuring this command. Similarly, hosts that do not implement the RFC 4191 functionality interpret the values sent by devices that do not implement the RFC 4191 extension with the **medium** preference option.

The **no** form of this command configures this command to the default value.

Default

nd-router-preference medium

Parameters

medium

Specifies the router advertises a medium default gateway preference.

high

Specifies the router advertises a high default gateway preference.

low

Specifies the router advertises a low default gateway preference.

Platforms

7705 SAR Gen 2

19.17 neid

neid

Syntax

neid *hex-string*

no neid

Context

[Tree] (config>system>ned>profile neid)

Full Context

configure system network-element-discovery profile neid

Description

This command configures the NEID for this profile.

The **no** form of this command deletes the NEID for this profile.

Parameters

hex-string

A hexadecimal string that consists of a subnet ID and basic ID. The first 8 high-order bits indicate the subnet ID and range from 0x1 to 0xFE. The 16 low-order bits indicate the basic ID and ranges from 0x0001 to 0xFFFE. The NEID cannot be configured as 0x90006 to 0x9FF06 or 0x9bff0.

Values 0x10001 to 0xFEFFFE

Platforms

7705 SAR Gen 2

19.18 neighbor

neighbor

Syntax

[**no**] **neighbor** *ip-int-name*

Context

[Tree] (config>router>rip>group neighbor)

[Tree] (config>router>ripng>group neighbor)

[Tree] (config>service>vprn>rip>group neighbor)

Full Context

configure router rip group neighbor

configure router ripng group neighbor

configure service vprn rip group neighbor

Description

This command creates a context for configuring a RIP neighbor interface. By default, group interfaces are not activated with RIP, unless explicitly configured. The BNG only learns RIP routes from IPv4 host on the group interface. The RIP neighbor group interface defaults to **none**. The send operation is unchangeable for group-interface.

The **no** form of this command deletes the RIP interface configuration for this group interface. The shutdown command in the **config>router>rip>group group-name>neighbor** context can be used to disable an interface without removing the configuration for the interface.

Default

no neighbor

Parameters

ip-int-name

Specifies the IP interface name. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

If the IP interface name does not exist or does not have an IP address configured, an error message will be returned.

Platforms

7705 SAR Gen 2

neighbor

Syntax

neighbor *ipv6-address mac-address*

no neighbor *ipv6-address*

Context

[Tree] (config>service>ies>if>ipv6 neighbor)

Full Context

configure service ies interface ipv6 neighbor

Description

This command configures IPv6-to-MAC address mapping on the IES interface.

Parameters

ipv6-address

The IPv6 address of the interface for which to display information.

Values		
	ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x - [0..FFFF]H
		d - [0..255]D

mac-address

Specifies the 48-bit MAC address for the IPv6-to-MAC address mapping in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

Platforms

7705 SAR Gen 2

neighbor

Syntax

[**no**] **neighbor** *ip-address*

Context

[Tree] (config>router>bgp>group neighbor)

Full Context

configure router bgp group neighbor

Description

This command creates a BGP peer/neighbor instance within the context of the BGP group.

This command can be issued repeatedly to create multiple peers and their associated configuration.

The **no** form of this command is used to remove the specified neighbor and the entire configuration associated with the neighbor. The neighbor must be administratively **shutdown** before attempting to delete it. If the neighbor is not shutdown, the command will not result in any action except a warning message on the console indicating that neighbor is still administratively up.

Default

no neighbor

Parameters

ip-address

Specifies the IP address of the BGP peer router in dotted decimal notation.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x [-interface]
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H
- d: [0 to 255]D

- interface: 32 characters maximum, mandatory for link local addresses

Platforms

7705 SAR Gen 2

neighbor

Syntax

[**no**] **neighbor** *ip-address*

Context

[Tree] (config>service>vprn>bgp>group neighbor)

Full Context

configure service vprn bgp group neighbor

Description

This command creates a BGP peer/neighbor instance within the context of the BGP group.

This command can be issued repeatedly to create multiple peers and their associated configuration.

The **no** form of this command is used to remove the specified neighbor and the entire configuration associated with the neighbor. The neighbor must be administratively **shutdown** before attempting to delete it. If the neighbor is not shut down, the command will not result in any action except a warning message on the console indicating that neighbor is still administratively up.

Parameters

ip-address

The IP address of the BGP peer router in dotted decimal notation.

Values

ipv4-address	a.b.c.d (host bits must be 0)
ipv6-address	x:x:x:x:x:x:x:x[-interface]
	x:x:x:x:x:x:d.d.d.d[-interface]
	x: [0 to FFFF]H
	d: [0 to 255]D
	interface: 32 characters maximum, mandatory for link local addresses

Platforms

7705 SAR Gen 2

neighbor

Syntax

neighbor *ipv6-address mac-address*

no neighbor *ipv6-address*

Context

[Tree] (config>service>vprn>if>ipv6 neighbor)

Full Context

configure service vprn interface ipv6 neighbor

Description

This command configures IPv6-to-MAC address mapping on the interface.

Parameters

ipv6-address

Specifies the IPv6 address on the interface.

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x [0 to FFFF]H
	d [0 to 255]D

mac-address

Specifies the 48-bit MAC address for the static ARP in the form *aa*:*bb*: *cc*:*dd*:*ee*:*ff* or *aa-bb-cc -dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

Platforms

7705 SAR Gen 2

neighbor

Syntax

[**no**] **neighbor** *ip-address*

Context

[Tree] (config>service>vprn>ospf>area>if neighbor)

[Tree] (config>service>vprn>ospf3>area>if neighbor)

Full Context

configure service vprn ospf area interface neighbor

configure service vprn ospf3 area interface neighbor

Description

This command configures an OSPF non-broadcast multi-access (NBMA) neighbor. The OSPF interface must be configured as an NBMA interface with the **interface-type non-broadcast** command. An NBMA network has no broadcast or multicast capabilities, so the router cannot discover its neighbors dynamically. All neighbors must be configured statically with the **neighbor** command.

In addition to configuring the OSPF NBMA neighbor's IP address, the neighbor's MAC address may need to be configured with the **config>service>vprn>interface>static-arp** command for OSPFv2 neighbors using its IPv4 address, and the **config>service>vprn>interface>ipv6>neighbor** command for OSPFv3 neighbors using its IPv6 link-local address.

The **no** form of this command removes the **neighbor** configuration.

Default

No OSPF NBMA neighbors are configured.

Parameters

ip-address

Specifies the OSPFv2 neighbor's IPv4 address or the OSPFv3 neighbor's IPv6 link-local address.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x [-interface]
	x:x:x:x:x:x:d.d.d.d [-interface]
	x: [0..FFFF]H
	d: [0..255]D
	interface —32 characters max, for link local addresses.

Platforms

7705 SAR Gen 2

neighbor

Syntax

neighbor *ipv6-address mac-address*

no neighbor *ipv6-address*

Context

[Tree] (config>router>if>ipv6 neighbor)

Full Context

configure router interface ipv6 neighbor

Description

This command configures an IPv6-to-MAC address mapping on the interface. Use this command if a directly attached IPv6 node does not support ICMPv6 neighbor discovery, or for some reason, a static address must be used. This command can only be used on Ethernet media.

The *ipv6-address* must be on the subnet that was configured from the IPv6 **address** command or a link-local address.

Parameters

ipv6-address

The IPv6 address assigned to a router interface.

Values

ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 to FFFF]H
	d:	[0 to 255]D

mac-address

Specifies the MAC address for the neighbor in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx.

Platforms

7705 SAR Gen 2

neighbor

Syntax

neighbor [*ip-int-name*]

no neighbor

Context

[Tree] (debug>router>ip neighbor)

Full Context

debug router ip neighbor

Description

This command enables IPv6 neighbor debugging.

Parameters

ip-int-name

Specifies the IP interface name.

Platforms

7705 SAR Gen 2

neighbor

Syntax

[**no**] **neighbor** *ipv4-address*

[**no**] **neighbor** *ipv6-address*

Context

[Tree] (config>router>ospf3>area>interface neighbor)

[Tree] (config>router>ospf>area>interface neighbor)

Full Context

configure router ospf3 area interface neighbor

configure router ospf area interface neighbor

Description

This command configures an OSPF non-broadcast multi-access (NBMA) neighbor. The OSPF interface must be configured as an NBMA interface with the **interface-type non-broadcast** command. An NBMA network has no broadcast or multicast capabilities, so the router cannot discover its neighbors dynamically. All neighbors must be configured statically with the **neighbor** command.

In addition to configuring the IP address of the OSPF NBMA neighbor, the MAC address of the neighbor may need to be configured with the **config>router>interface>static-arp** command for OSPFv2 neighbors using its IPv4 address, and the **config>router>interface>ipv6>neighbor** command for OSPFv3 neighbors using its IPv6 link-local address.

The **no** form of this command removes the **neighbor** configuration.

Default

no neighbor

Parameters

ipv4-address

Specifies the IPv4 address of the OSPFv2 neighbor.

Values ipv4-address — a.b.c.d

ipv6-address

Specifies the IPv6 link-local address of the OSPFv3 neighbor.

Values

ipv6-address: x:x:x:x:x:x:x:x [-interface]

x:x:x:x:x:x:d.d.d.d [-interface]

x: [0..FFFF]H

d: [0..255]D

interface — 32 characters maximum for link local addresses.

Platforms

7705 SAR Gen 2

neighbor

Syntax

neighbor [*ip-int-name* | *ip-address*]

neighbor [*ip-int-name*] [*router-id*]

no neighbor

Context

[Tree] (debug>router>ospf3 neighbor)

[Tree] (debug>router>ospf neighbor)

Full Context

debug router ospf3 neighbor

debug router ospf neighbor

Description

This command enables debugging for an OSPF or OSPF3 neighbor.

Parameters

ip-int-name

Specifies the neighbor interface name.

ip-address

Specifies neighbor information for the neighbor identified by the specified IP address, in the **debug**>**router**>**ospf** context.

router-id

Specifies neighbor information for the neighbor identified by the specified router ID, in the **debug**>**router**>**ospf3** context.

Platforms

7705 SAR Gen 2

neighbor

Syntax

neighbor {*ip-address* | **prefix-list** *name*}

no neighbor

Context

[Tree] (config>router>policy-options>policy-statement>entry>to neighbor)

[Tree] (config>router>policy-options>policy-statement>entry>from neighbor)

Full Context

configure router policy-options policy-statement entry to neighbor

configure router policy-options policy-statement entry from neighbor

Description

This command specifies the neighbor address as found in the source address of the actual join and prune message as a filter criterion. If no neighbor is specified, any neighbor is considered a match.

The **no** form of the of the command removes the neighbor IP match criterion from the configuration.

Default

no neighbor

Parameters

ip-address

Specifies the neighbor IP address in dotted decimal notation.

Values ipv4-address:

- a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x:x [-interface]
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H

- d: [0 to 255]D
- interface: 32 characters maximum, mandatory for link local addresses

prefix-list *name*

Specifies the prefix-list name. Allowed values are any string up to 64 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

The *name* specified must already be defined.

Platforms

7705 SAR Gen 2

19.19 neighbor-limit

neighbor-limit

Syntax

neighbor-limit *limit* [**log-only**] [**threshold** *percent*]

no neighbor-limit

Context

[Tree] (config>service>ies>if>ipv6 neighbor-limit)

Full Context

configure service ies interface ipv6 neighbor-limit

Description

This command configures the maximum amount of dynamic IPv6 neighbor entries that can be learned on an IP interface.

When the number of dynamic neighbor entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations is dropped. Entries that have already been learned is refreshed.

The **no** form of this command removes the **neighbor-limit**.

Default

no neighbor-limit

Parameters

log-only

Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit is learned.

percent

The threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

limit

The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic neighbor learning is disabled and no dynamic neighbor entries are learned.

Values 0 to 102400

Platforms

7705 SAR Gen 2

neighbor-limit

Syntax

neighbor-limit *limit* [**log-only**] [**threshold** *percent*]

no neighbor-limit

Context

[Tree] (config>service>vprn>if>ipv6 neighbor-limit)

Full Context

configure service vprn interface ipv6 neighbor-limit

Description

This command configures the maximum amount of dynamic IPv6 neighbor entries that can be learned on an IP interface.

When the number of dynamic neighbor entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations will be dropped. Entries that have already been learned will be refreshed.

The **no** form of this command removes the **neighbor-limit**.

Default

neighbor-limit 90

Parameters

log-only

Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned.

percent

The threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

limit

The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic neighbor learning is disabled and no dynamic neighbor entries are learned.

Values 0 to 102400

Platforms

7705 SAR Gen 2

neighbor-limit

Syntax

neighbor-limit *limit* [**log-only**] [**threshold** *percent*]

no neighbor-limit

Context

[Tree] (config>router>if>ipv6 neighbor-limit)

Full Context

configure router interface ipv6 neighbor-limit

Description

This command configures the maximum amount of dynamic IPv6 neighbor entries that can be learned on an IP interface.

When the number of dynamic neighbor entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations will be dropped. Entries that have already been learned will be refreshed.

The **no** form of this command removes the neighbor-limit.

Default

no neighbor-limit

Parameters

limit

The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic neighbor learning is disabled and no dynamic neighbor entries are learned.

Values 0 to 102400

log-only

Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned.

percent

The threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

Platforms

7705 SAR Gen 2

19.20 neighbor-liveness-time

neighbor-liveness-time

Syntax

neighbor-liveness-time *interval*

no neighbor-liveness-time

Context

[Tree] (config>router>ldp>graceful-restart neighbor-liveness-time)

Full Context

configure router ldp graceful-restart neighbor-liveness-time

Description

This command configures the neighbor liveness time.

The **no** form of this command returns the default value.

Default

no neighbor-liveness (which equals a value of 120 seconds)

Parameters

interval

Specifies the length of time in seconds.

Values 5 to 300

Platforms

7705 SAR Gen 2

19.21 neighbor-trust

neighbor-trust

Syntax

neighbor-trust [**vpn-ipv4**] [**vpn-ipv6**] [**evpn**]

no neighbor-trust

Context

[Tree] (config>router>bgp neighbor-trust)

Full Context

configure router bgp neighbor-trust

Description

This command enables a label security feature for prefixes of a VPN family at an inter-AS boundary.

This label security feature allows the configuration of a router, acting in a PE, ASBR, or both roles, to accept packets of VPN-IP or EVPN prefixes only from direct EBGP neighbors to which it advertised a service label.

The untrusted state identifies the participating interfaces. The router supports a maximum of 15 network interfaces that can participate in this feature.

At a high level, BGP tracks each direct EBGP neighbor over an untrusted interface to which it sent a prefix label. For each of those prefixes, BGP programs a bitmap in the ILM record that indicates, on per-untrusted interface basis, whether the matching received packets must be forwarded or dropped.

The **no** form of this command disables the inter-AS security feature for the VPN family.

Parameters

vpn-ipv4

Keyword to enable the inter-AS label security for VPN IPv4 family.

vpn-ipv6

Keyword to enable the inter-AS label security for VPN IPv6 family.

evpn

Keyword to enable the inter-AS label security for EVPN family.

Platforms

7705 SAR Gen 2

19.22 neip

neip

Syntax

neip

Context

[Tree] (config>system>ned>profile neip)

Full Context

configure system network-element-discovery profile neip

Description

Commands in this context configure the NEIP.

Platforms

7705 SAR Gen 2

19.23 netbios-name-server

netbios-name-server

Syntax

netbios-name-server *ip-address* [*ip-address*]

no netbios-name-server

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options netbios-name-server)

[Tree] (config>service>vprn>dhcp>server>pool>options netbios-name-server)

[Tree] (config>router>dhcp>server>pool>options netbios-name-server)

Full Context

configure subscriber-mgmt local-user-db ipoe host options netbios-name-server

configure service vprn dhcp local-dhcp-server pool options netbios-name-server

configure router dhcp local-dhcp-server pool options netbios-name-server

Description

This command configures up to four Network Basic Input/Output System (NetBIOS) name server IP addresses for a DHCP client.

The **no** form of this command removes the IP address from the netbios-name-server configuration.

Parameters

ip-address

Specifies up to four NetBIOS name server IP addresses. The address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

Platforms

7705 SAR Gen 2

19.24 netbios-node-type

netbios-node-type

Syntax

netbios-node-type *netbios-node-type*

no netbios-node-type

Context

[Tree] (config>service>vprn>dhcp>server>pool>options netbios-node-type)

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options netbios-node-type)

[Tree] (config>router>dhcp>server>pool>options netbios-node-type)

Full Context

configure service vprn dhcp local-dhcp-server pool options netbios-node-type

configure subscriber-mgmt local-user-db ipoe host options netbios-node-type

configure router dhcp local-dhcp-server pool options netbios-node-type

Description

This command configures the Network Basic Input/Output System (NetBIOS) node type.

The **no** form of this command removes the NetBIOS node type parameters from the configuration.

Parameters

netbios-node-type

Specifies the netbios node type.

Values B — Broadcast node uses broadcasting to query nodes on the network for the owner of a NetBIOS name.

P — Peer-to-peer node uses directed calls to communicate with a known NetBIOS name server for the IP address of a NetBIOS machine name.

M — Mixed node uses broadcast queries to find a node, and if that fails, queries a known P-node name server for the address.

H — Hybrid node is the opposite of the M-node action so that a directed query is executed first, and if that fails, a broadcast is attempted.

Platforms
7705 SAR Gen 2

19.25 netconf

netconf

Syntax
netconf

Context
[Tree] (debug>system netconf)

Full Context
debug system netconf

Description
Commands in this context debug NETCONF.

Platforms
7705 SAR Gen 2

netconf

Syntax
netconf

Context
[Tree] (config>system>security>profile netconf)

Full Context
configure system security profile netconf

Description
This command authorizes various netconf capabilities for the user.

Platforms
7705 SAR Gen 2

netconf

Syntax
netconf

Context
[Tree] (config>system>security>management-interface netconf)

Full Context
configure system security management-interface netconf

Description
Commands in this context configure hash-control for the Netconf interface.

Platforms
7705 SAR Gen 2

19.26 netconf-stream

netconf-stream

Syntax
netconf-stream *stream-name*
no netconf-steam

Context
[Tree] (config>log>log-id netconf-stream)

Full Context
configure log log-id netconf-stream

Description

This command is used to associate a NETCONF stream name with a log ID. The NETCONF stream name must be unique per SR OS device. For the same log ID, **to netconf** must be configured for a subscription to that NETCONF stream name to be accepted. A **netconf-stream** cannot be set to "NETCONF" as "NETCONF" is reserved for log-id 101. If a **netconf-stream** is changed, active subscriptions to the changed stream name are terminated by SR OS.

The **no** form of this command removes a NETCONF stream name from a log ID. Active subscriptions to the removed stream name are terminated by SR OS.

Parameters

stream-name

Specifies a NETCONF stream name, up to 32 characters.

Platforms

7705 SAR Gen 2

19.27 network

network

Syntax

network

Context

[Tree] (config>port network)

[Tree] (config>card>mda network)

Full Context

configure port network

configure card mda network

Description

This command enables the network context to configure egress and ingress pool policy parameters.

On the MDA level, network egress pools are only allocated on channelized MDAs.

Platforms

7705 SAR Gen 2

network

Syntax

network

Context

[Tree] (config>card>fp>ingress network)

Full Context

configure card fp ingress network

Description

This command specifies the CLI node that contains the network forwarding-plane parameters.

Platforms

7705 SAR Gen 2

network

Syntax

network

Context

[Tree] (config>port>ethernet network)

Full Context

configure port ethernet network

Description

This command enables access to the context to configure network port parameters.

Platforms

7705 SAR Gen 2

network

Syntax

network

Context

[Tree] (config>service>vprn network)

Full Context

configure service vprn network

Description

Commands in this context configure network parameters for the VPRN service.

Platforms

7705 SAR Gen 2

network

Syntax

network *network-policy-id* [**create**] [**name** *name*]

no network *network-policy-id*

Context

[Tree] (config>qos network)

Full Context

configure qos network

Description

This command creates or edits a QoS network policy. The network policy defines the treatment that IP or MPLS packets receive as they ingress and egress the network port.

The QoS network policy consists of an ingress and egress component. The ingress component of the policy defines how DiffServ code points and MPLS EXP bits are mapped to internal forwarding class and profile state. The forwarding class and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the router. The mapping on each network interface defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the network interface.

The egress component of the network QoS policy defines the queuing parameters associated with each forwarding class. Each of the forwarding classes defined within the system automatically creates a queue on each network interface. This queue gets all the parameters defined within the default network QoS policy 1 until an explicit policy is defined for the network interface access uplink port. If the egressing packet originated on an ingress SAP, or the remarking parameter is defined for the egress interface, the egress QoS policy also defines the IP DSCP, dot1p/DE, or MPLS EXP bit marking based on the forwarding class and the profile state.

Network **policy-id 1** exists as the default policy that is applied to all network interfaces by default. The network **policy-id 1** cannot be modified or deleted. It defines the default DSCP-to-FC mapping and MPLS EXP-to-FC mapping for the ingress. For the egress, it defines six forwarding classes that represent individual queues and the packet marking criteria.

Network **policy-id 1** exists as the default policy that is applied to all network ports by default. This default policy cannot be modified or deleted. It defines the default DSCP-to-FC mapping and default unicast meters for ingress IP traffic. For the egress, it defines the forwarding class to dot1p and DSCP values and the packet marking criteria.

If a new network policy is created (for instance, policy-id 3), only the default action and egress forwarding class parameters are identical to the default policy. A new network policy does not contain the default DSCP-to-FC and MPLS-EXP-to-FC mapping for network QoS policy of type **ip-interface** or the DSCP-to-FC mapping (for network QoS policy of type **port**). The default network policy can be copied (use the copy command) to create a new network policy that includes the default ingress DSCP-to-FC and MPLS EXP-to-FC mapping (as appropriate). Parameters can be modified, or the **no** form of this command can be used to remove an object from the configuration.

Any changes made to an existing policy, using any of the sub-commands, will be applied immediately to all network interfaces where this policy is applied. For this reason, when many changes are required on a policy, it is highly recommended that the policy be copied to a work area policy-id. That work-in-progress policy can be modified until complete, then written over the original policy-id. Use the config qos copy command to maintain policies in this manner.

The **no** form of this command deletes the network policy. A policy cannot be deleted until it is removed from all entities where it is applied. The default network **policy** *policy-id* 1 cannot be deleted.

Default

network 1 — System Default Network Policy 1

Parameters

network-policy-id

The policy-id uniquely identifies the policy on the router.

Values 1 to 65535

Default 1

create

Required parameter when creating a QoS network policy.

name name

A name that is saved as part of the configuration data. If a name is not specified at creation time, then SR OS assigns a string version of the network policy identifier as the name.

Values A string up to 64 characters

Platforms

7705 SAR Gen 2

19.28 network-domain

network-domain

Syntax

[**no**] **network-domain** *network-domain-name*

Context

[Tree] (config>router>network-domains network-domain)

Full Context

configure router network-domains network-domain

Description

This command creates network-domains that can be associated with individual interfaces and SDPs.

Default

network-domain "default”

Parameters

network-domain-name

Specifies the network domain name, up to 32 characters.

Platforms

7705 SAR Gen 2

network-domain

Syntax

[**no**] **network-domain** *network-domain-name*

Context

[Tree] (config>router>if network-domain)

Full Context

configure router interface network-domain

Description

This command assigns a given interface to a given network-domain. The network-domain is then taken into account during sap-ingress queue allocation for VPLS SAP.

The network-domain association can only be done in a base-routing context. Associating a network domain with an loop-back or system interface will be rejected. Associating a network-domain with an interface that has no physical port specified will be accepted, but will have no effect as long as a corresponding port, or LAG, is defined.

Single interfaces can be associated with multiple network-domains.

Default

network-domain "default”

Platforms

7705 SAR Gen 2

network-domain

Syntax

network-domain *network-domain-name*

no network-domain

Context

[Tree] (config>service>sdp network-domain)

Full Context

configure service sdp network-domain

Description

This command assigns a given SDP to a given network-domain. The network-domain is then taken into account during sap-ingress queue allocation for VPLS SAP.

The network-domain association can only be done in a base-routing context. Associating a network domain with an loop-back or system interface will be rejected. Associating a network-domain with an interface that has no physical port specified will be accepted, but will have no effect as long as a corresponding port, or LAG, is undefined.

A single SDP can only be associated with a single network-domain.

Default

network-domain "default"

Platforms

7705 SAR Gen 2

19.29 network-domains

network-domains

Syntax

network-domains

Context

[Tree] (config>router network-domains)

Full Context

configure router network-domains

Description

This command opens context for defining network-domains. This command is applicable only in the base routing context.

Platforms

7705 SAR Gen 2

19.30 network-element-discovery

network-element-discovery

Syntax

network-element-discovery

Context

[Tree] (config>system network-element-discovery)

Full Context

configure system network-element-discovery

Description

Commands in this context configure the network-element discovery parameters and MIB table generation.

Platforms

7705 SAR Gen 2

19.31 network-interface

network-interface

Syntax

network-interface *interface-name* [**create**]

no network-interface *interface-name*

Context

[Tree] (config>service>vprn network-interface)

Full Context

configure service vprn network-interface

Description

This command configures a network interface in a VPRN that acts as a CSC interface to a CSC-CE in a Carrier Supporting Carrier IP VPN deployment model.

Parameters

interface-name

Specifies the name of the interface to be added.

create

Keyword used to create the network interface.

Platforms

7705 SAR Gen 2

19.32 network-queue

network-queue

Syntax

network-queue *policy-name* [**create**]

no network-queue *policy-name*

Context

[Tree] (config>qos network-queue)

Full Context

configure qos network-queue

Description

This command creates a context to configure a network queue policy. Network queue policies define the ingress network queuing at the FP network node level and on the Ethernet port and SONET/SDH path level to define network egress queuing.

Default

network-queue "default”

Parameters

policy-name

The name of the network queue policy.

Values Valid names consist of any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

create

Required keyword when creating a network queue policy.

Platforms

7705 SAR Gen 2

19.33 new-password-at-login

new-password-at-login

Syntax

[**no**] **new-password-at-login**

Context

[Tree] (config>system>security>user>console new-password-at-login)

Full Context

configure system security user console new-password-at-login

Description

This command forces the user to change a password at the next console login. The new password applies to FTP but the change can be enforced only by the console, SSH, or Telnet login.

The **no** form of this command does not force the user to change passwords.

Default
no new-password-at-login

Platforms
7705 SAR Gen 2

19.34 newline

newline

Syntax
[**no**] **newline**

Context
[Tree] (config>system>management-interface>cli>md-cli>environment>prompt newline)

Full Context
configure system management-interface cli md-cli environment prompt newline

Description
This command displays a new line before the first prompt line.

The **no** form of this command suppresses the new line before the first prompt line.

Default
newline

Platforms
7705 SAR Gen 2

19.35 next-header

next-header

Syntax
next-header *next-header*

no next-header

Context

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter>entry next-header)

Full Context

configure system security management-access-filter ipv6-filter entry next-header

Description

This command specifies the next header to match. The protocol type such as TCP, UDP or OSPF is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17). IPv6 Extension headers are identified by the next header IPv6 numbers as per RFC 2460.

Parameters

next-header

Specifies for IPv4 MAF the IP protocol field, and for IPv6 the next header type to be used in the match criteria for this Management Access Filter Entry.

Values

next-header:	0 to 255, protocol numbers accepted in DHB
keywords:	none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, drp, igmp, igp, ip, ipv6, ipv6-icmp, ipv6-no-nxt, isis, iso-ip, l2tp, spf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp

Platforms

7705 SAR Gen 2

19.36 next-hop

next-hop

Syntax

next-hop {*ip-address* | *ip-int-name* | *ipv6 address*}

Context

[Tree] (config>service>vprn>static-route-entry next-hop)

Full Context

configure service vprn static-route-entry next-hop

Description

This command specifies the directly connected next hop IP address or interface used to reach the destination. If the next hop is over an unnumbered interface or a point-to-point interface, the *ip-int-name* of the unnumbered or point-to-point interface (on this node) can be configured.

The configured *ip-address* can be either on the network side or the access side on this node. The address must be associated with a network directly connected to a network configured on this node.

Default

no next-hop

Parameters

ip-int-name, ipv4-address, ipv6-address

the IP-INT, IPv4, and IPv6 addresses

Values

ip-int-name	32 characters max
ipv4-address	a.b.c.d
ipv6-address	x:x:x:x:x:x:x:x-[interface]
	x:x:x:x:x:x:d.d.d.d[-interface]
	x: [0 to FFFF]H
	d: [0 to 255]D
	interface: 32 characters maximum, mandatory for link local addresses

Platforms

7705 SAR Gen 2

next-hop

Syntax

next-hop *ip-address*

no next-hop

Context

[Tree] (config>router>mpls>fwd-policies>fwd-policy>nh-grp>pri next-hop)

[Tree] (config>router>mpls>fwd-policies>fwd-policy>nh-grp>bkup next-hop)

Full Context

configure router mpls forwarding-policies forwarding-policy next-hop-group primary-next-hop next-hop

configure router mpls forwarding-policies forwarding-policy next-hop-group backup-next-hop next-hop

Description

This command configures the address of primary or backup next hop of an NHG entry in a forwarding policy.

The **no** form of this command removes the address of primary or backup next hop of an NHG entry in a forwarding policy.

Parameters

ip-address

Specifies the destination IPv4 or IPv6 address.

Values		
	ipv4-address	a.b.c.d
	ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x - [0..FFFF]H
		d - [0..255]D

Platforms

7705 SAR Gen 2

next-hop

Syntax

next-hop {*ip-int-name* | *ip-address* | *ipv6-address*}

Context

[Tree] (config>router>static-route-entry next-hop)

Full Context

configure router static-route-entry next-hop

Description

This command specifies the directly connected next hop IP address or interface used to reach the destination. If the next hop is over a point-to-point unnumbered interface, the **ip-int-name** of the unnumbered point-to-point interface (on this node) can be configured.

The configured *ip-address* can be either on the network side or the access side on this node. The address must be associated with a network directly connected to a network configured on this node.

Default

no next-hop

Parameters

ip-int-name* | *ip-address* | *ipv6-address

Specifies the interface or IPv4/IPv6 address of the next hop.

Values

ip-int-name	32 characters max
ipv4-address	a.b.c.d
ipv6-address	x:x:x:x:x:x:x:x-[interface]
	x:x:x:x:x:x:d.d.d.d[-interface]
	x: [0..FFFF]H
	d: [0..255]D
	interface: 32 characters maximum, mandatory for link local addresses

Platforms

7705 SAR Gen 2

next-hop

Syntax

[**no**] **next-hop** *ip-address*

Context

[Tree] (config>vrrp>policy>priority-event>route-unknown next-hop)

Full Context

configure vrrp policy priority-event route-unknown next-hop

Description

This command enables an allowed next hop IP address to match the IP route prefix for a route-unknown priority control event.

If the next-hop IP address does not match one of the defined *ip-address*, the match is considered unsuccessful and the **route-unknown** event transitions to the set state.

The **next-hop** command is optional. If no **next-hop** *ip-address* commands are configured, the comparison between the RTM prefix return and the **route-unknown** IP route prefix are not included in the next hop information.

When more than one next hop IP addresses are eligible for matching, a **next-hop** command must be executed for each IP address. Defining the same IP address multiple times has no effect after the first instance.

The **no** form of the command removes the *ip-address* from the list of acceptable next hops when looking up the **route-unknown** prefix. If this *ip-address* is the last next hop defined on the **route-unknown** event, the returned next hop information is ignored when testing the match criteria. If the *ip-address* does not exist, the **no next-hop** command returns a warning error, but continues to execute if part of an **exec** script.

Default

no next-hop — No next hop IP address for the route unknown priority control event is defined.

Parameters

ip-address

The IP address for an acceptable next hop IP address for a returned route prefix from the RTM when looking up the **route-unknown** route prefix.

Values

ipv4-address:	a.b.c.d	
ipv6-address:	x:x:x:x:x:x:x:x[-interface]	
	x:	[0..FFFF]H
	interface:	32 chars maximum, mandatory for link local addresses

The link-local IPv6 address must have an interface name specified. The global IPv6 address must not have an interface name specified.

Platforms

7705 SAR Gen 2

next-hop

Syntax

next-hop *ip-address*

next-hop prefix-list *name*

no next-hop

Context

[Tree] (config>router>policy-options>policy-statement>entry>from next-hop)

Full Context

configure router policy-options policy-statement entry from next-hop

Description

This command enables BGP routes to be matched based on the BGP next-hop address. The match condition is evaluated against the IPv4 or IPv6 address in the NEXT_HOP or MP_REACH_NLRI attribute.

When the next-hop match is applied to VPN-IP routes, the Route Distinguisher (RD) is ignored.

A non-BGP route does not match a policy entry if it contains the **next-hop** command.

Default

no next-hop

Parameters

ip-address

An IPv4 or IPv6 address.

Values a.b.c.d or x:x:x:x:x:x:x:x or x:x:x:x:x:x:d.d.d.d

name

Specifies the name of a prefix-list (up to 64 characters).

prefix-list

Specifies that the BGP next hop should be matched against a prefix-list instead of an individual IP address.

Platforms

7705 SAR Gen 2

next-hop

Syntax

next-hop {*ip-address* | **peer-address**}

no next-hop

Context

[Tree] (config>router>policy-options>policy-statement>entry>action next-hop)

[Tree] (config>router>policy-options>policy-statement>default-action next-hop)

Full Context

configure router policy-options policy-statement entry action next-hop

configure router policy-options policy-statement default-action next-hop

Description

This command assigns the specified next hop IP address to routes matching the policy statement entry.

If a next-hop IP address is not specified, the next-hop attribute is not changed.

The **no** form of this command disables assigning a next hop address in the route policy entry.

Default

no next-hop

Parameters

ip-address

Specifies the next hop IP address in dotted decimal notation.

Values

ipv4-address:	a.b.c.d	
ipv6-address:	x:x:x:x:x:x:x:x	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 to FFFF]H
	d:	[0 to 255]D
param-name:	The next-hop parameter variable name.	
	Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@”.	

peer-address

Set the next-hop IP address to the peer’s IP address.

Platforms

7705 SAR Gen 2

19.37 next-hop-group

next-hop-group

Syntax

next-hop-group *index* [**resolution-type** { **direct** | **indirect**}]

no next-hop-group *index*

Context

[Tree] (config>router>mpls>fwd-policies>fwd-policy next-hop-group)

Full Context

configure router mpls forwarding-policies forwarding-policy next-hop-group

Description

This command configures an NHG entry in an MPLS forwarding policy.

Each NHG can have primary and backup next hops of the same type.

The **no** form of this command removes the NHG from the MPLS forwarding policy.

Parameters

index

Specifies the index value.

Values 1 to 32

direct

Specifies the direct resolution type.

indirect

Specifies the indirect resolution type.

Platforms

7705 SAR Gen 2

19.38 next-hop-reachability

next-hop-reachability

Syntax

[**no**] **next-hop-reachability**

Context

[Tree] (configure>router>bgp>group>neighbor>bfd-strict-mode next-hop-reachability)

[Tree] (configure>router>bgp>bfd-strict-mode next-hop-reachability)

[Tree] (configure>service>vprn>bgp>bfd-strict-mode next-hop-reachability)

[Tree] (configure>service>vprn>bgp>group>bfd-strict-mode next-hop-reachability)

[Tree] (configure>service>vprn>bgp>group>neighbor>bfd-strict-mode next-hop-reachability)

[Tree] (configure>router>bgp>group>bfd-strict-mode next-hop-reachability)

Full Context

configure router bgp group neighbor bfd-strict-mode next-hop-reachability

configure router bgp bfd-strict-mode next-hop-reachability

configure service vprn bgp bfd-strict-mode next-hop-reachability

configure service vprn bgp group bfd-strict-mode next-hop-reachability

configure service vprn bgp group neighbor bfd-strict-mode next-hop-reachability

configure router bgp group bfd-strict-mode next-hop-reachability

Description

This command configures the router to consider next-hop self routes belonging to specific address families received from a peer within scope of this command as having an unresolved next hop, provided that the following requirements are met:

- The BFD session to the peer is in a down state.
- There is a valid interface BFD configuration that applies to the peer.
- There is a valid BFD liveness configuration that applies to the peer.

The unresolved state is maintained until the BFD session state changes to up or administratively down, even if there is a resolving route or tunnel that matches the BGP next-hop address.

Routes received from one peer with a BGP next-hop address equal to the address of another peer are not affected by the BFD session to the other peer.

The behavior of the router when this command is enabled does not depend on whether Strict-BFD is used, as both features are independent.

Enabling this command only affects routes belonging to the following address families:

- IPv4
- IPv6
- IPv4 VPN
- IPv6 VPN
- labeled unicast IPv4
- labeled unicast IPv6
- EVPN
- IPv4 multicast
- IPv6 multicast
- IPv4 VPN multicast
- IPv6 VPN multicast

The **no** form of this command prevents the router from considering next-hop self routes belonging to the preceding address families as having an unresolved next hop if the BFD session goes down.

Default

no next-hop-reachability

Platforms

7705 SAR Gen 2

19.39 next-hop-resolution

next-hop-resolution

Syntax

next-hop-resolution

Context

[Tree] (config>service>vprn>bgp next-hop-resolution)

Full Context

configure service vprn bgp next-hop-resolution

Description

Commands in this context configure next-hop resolution parameters.

Platforms

7705 SAR Gen 2

next-hop-resolution

Syntax

next-hop-resolution

Context

[Tree] (config>router>bgp next-hop-resolution)

Full Context

configure router bgp next-hop-resolution

Description

Commands in this context configure next-hop resolution parameters.

Platforms

7705 SAR Gen 2

19.40 next-hop-self

next-hop-self

Syntax

[**no**] **next-hop-self**

Context

[Tree] (config>service>vprn>bgp>group>neighbor next-hop-self)

[Tree] (config>service>vprn>bgp>group next-hop-self)

Full Context

configure service vprn bgp group neighbor next-hop-self

configure service vprn bgp group next-hop-self

Description

This command configures the group or neighbor to always set the NEXTHOP path attribute to its own physical interface when advertising to a peer.

This is primarily used to avoid third-party route advertisements when connected to a multi-access network.

The **no** form of this command used at the group level allows third-party route advertisements in a multi-access network.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no next-hop-self — Third-party route advertisements are allowed.

Platforms

7705 SAR Gen 2

next-hop-self

Syntax

[**no**] **next-hop-self**

Context

[Tree] (config>router>bgp>group next-hop-self)

[Tree] (config>router>bgp>group>neighbor next-hop-self)

Full Context

configure router bgp group next-hop-self

configure router bgp group neighbor next-hop-self

Description

This command enables BGP to advertise routes to members of a group or to a specific neighbor using a local address of the BGP instance as the BGP next-hop address. Note that **next-hop-self** is set without exception, regardless of the route source (EBGP or IBGP) or its family. When used with VPN-IPv4 and VPN-IPv6 routes the **enable-rr-vpn-forwarding** command should also be configured.

The **no** form of this command uses protocol standard behavior to decide whether or not to set **next-hop-self** in advertised routes.

Default

no next-hop-self

Platforms

7705 SAR Gen 2

next-hop-self

Syntax

[**no**] **next-hop-self**

Context

[Tree] (config>router>policy-options>policy-statement>default-action next-hop-self)

[Tree] (config>router>policy-options>policy-statement>entry>action next-hop-self)

Full Context

configure router policy-options policy-statement default-action next-hop-self

configure router policy-options policy-statement entry action next-hop-self

Description

This command configures BGP to advertise routes that match a policy entry (or that match no other policy entry and, therefore, to which the default action applies) using a local address of the BGP instance as the BGP next-hop address. The command applies to IPv4, IPv6, label-IPv4, and label-IPv6 routes. It also applies to VPN-IPv4 and VPN-IPv6 routes, but only when used in conjunction with the **enable-rr-vpn-forwarding** command.

This command affects how routes are advertised to IBGP peers, regardless of whether or not they were learned from an IBGP or EBGP peer

The **no** form of this command uses protocol standard behavior to decide whether or not to set **next-hop-self** in advertised routes.

Default

no next-hop-self

Platforms

7705 SAR Gen 2

19.41 next-hop-unchanged

next-hop-unchanged

Syntax

next-hop-unchanged [**label-ipv4**] [**label-ipv6**] [**vpn-ipv4**] [**vpn-ipv6**] [**evpn**]

no next-hop-unchanged

Context

[Tree] (config>router>bgp>group next-hop-unchanged)

[Tree] (config>router>bgp>group>neighbor next-hop-unchanged)

Full Context

configure router bgp group next-hop-unchanged

configure router bgp group neighbor next-hop-unchanged

Description

This command enables unchanged BGP next-hops when sending BGP routes to peers in this group or neighbor.

The **no** form of this command disables unchanged BGP next-hops.

Default

no next-hop-unchanged

Parameters

evpn

Specifies BGP next hops are unchanged for the evpn address family.

label-ipv4

Specifies BGP next hops are unchanged for the label-ipv4 address family.

label-ipv6

Specifies BGP next hops are unchanged for the label-ipv6 address family.

vpn-ipv4

Specifies BGP next hops are unchanged for the vpn-ipv4 address family.

vpn-ipv6

Specifies BGP next hops are unchanged for the vpn-ipv6 address family.

Platforms

7705 SAR Gen 2

19.42 nh-type

nh-type

Syntax

nh-type {**ip** | **tunnel**}

no nh-type

Context

[Tree] (config>router>route-next-hop-policy>template nh-type)

Full Context

configure router route-next-hop-policy template nh-type

Description

This command configures the next-hop type constraint into the route next-hop policy template.

The user can select if tunnel backup next-hop or IP backup next-hop is preferred. The default in SR OS implementation is to prefer IP next-hop over tunnel next-hop. The implementation will fall back to the other type if no LFA next-hop of the preferred type is found.

When the route next-hop policy template is applied to an IP interface, all prefixes using this interface as a primary next-hop will follow the next-hop type preference specified in the template.

The **no** form deletes the next-hop type constraint from the route next-hop policy template.

Default

nh-type ip

Parameters

{ip | tunnel}

Specifies the two possible values for the next-hop type.

Default ip

Platforms

7705 SAR Gen 2

19.43 nmda

nmda

Syntax

nmda

Context

[Tree] (config>system>management-interface>yang-modules nmda)

Full Context

configure system management-interface yang-modules nmda

Description

Commands in this context configure the attributes for the Network Management Datastores Architecture (NMDA).

Platforms

7705 SAR Gen 2

19.44 nmda-support

nmda-support

Syntax

[**no**] **nmda-support**

Context

[Tree] (config>system>management-interface>yang-modules>nmda nmda-support)

Full Context

configure system management-interface yang-modules nmda nmda-support

Description

This command enables the advertisement of NMDA support over NETCONF through the use of YANG library 1.1.

The **no** form of this command disables NMDA advertisement over NETCONF and YANG library 1.0 is used.

Default
no nmda-support

Platforms
7705 SAR Gen 2

19.45 node-id-in-rro

node-id-in-rro

Syntax
[**no**] **node-id-in-rro** [**include** | **exclude**]

Context
[Tree] (config>router>rsvp node-id-in-rro)

Full Context
configure router rsvp node-id-in-rro

Description
This command enables the option to include node-id sub-object in RRO. Node-ID sub-object propagation is required to provide fast reroute protection for LSP that spans across multiple area domains.

If this option is disabled, then node-id is not included in RRO object.

Default
node-id-in-rro exclude

Platforms
7705 SAR Gen 2

19.46 node-protect

node-protect

Syntax
[**no**] **node-protect**

Context
[Tree] (config>router>mpls>lsp>fast-reroute node-protect)

[Tree] (config>router>mpls>lsp-template>fast-reroute node-protect)

Full Context

configure router mpls lsp fast-reroute node-protect

configure router mpls lsp-template fast-reroute node-protect

Description

This command enables or disables node and link protection on the specified LSP. Node protection ensures that traffic from an LSP traversing a neighboring router will reach its destination even if the neighboring router fails.

Default

node-protect (for a provisioned LSP)

no node-protect (for a P2P LSP template)

Platforms

7705 SAR Gen 2

node-protect

Syntax

node-protect [**max-pq-nodes** *value*]

no node-protect

Context

[Tree] (config>router>isis>loopfree-alternates>remote-lfa node-protect)

Full Context

configure router isis loopfree-alternates remote-lfa node-protect

Description

This command enables node-protect in which the router prefers a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node.

The **no** form of this command disables node-protect.

Default

no node-protect

Parameters

value

Specifies the maximum number of PQ nodes found in the LFA SPFs for which the node protection check is performed. The node-protect condition means the router must run the

original Remote LFA algorithm plus one extra forward SPF on behalf of each PQ node found, potentially after applying the **max-pq-cost** parameter, to check if the path from the PQ node to the destination does not traverse the protected node. Setting this parameter to a lower value means the LFA SPFs will use less computation time and resources but may result in not finding a node-protect repair tunnel.

Values 1 to 32

Default 16

Platforms

7705 SAR Gen 2

node-protect

Syntax

[**no**] **node-protect**

Context

[Tree] (config>router>isis>loopfree-alternates>ti-lfa node-protect)

Full Context

configure router isis loopfree-alternates ti-lfa node-protect

Description

This command enables node-protect in which the router prefers a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node.

The **no** form of this command disables node-protect.

Default

no node-protect

Platforms

7705 SAR Gen 2

node-protect

Syntax

node-protect [**max-pq-nodes** *value*]

no node-protect

Context

[Tree] (config>router>ospf>loopfree-alternates>remote-lfa node-protect)

Full Context

configure router ospf loopfree-alternates remote-lfa node-protect

Description

This command enables node-protect in which the router prefers a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node.

The **no** form of this command disables node-protect.

Default

no node-protect

Parameters

max-pq-nodes value

Specifies the maximum number of PQ nodes found in the LFA SPFs for which the node protection check is performed. The node-protect condition means the router must run the original Remote LFA algorithm plus one extra forward SPF on behalf of each PQ node found, potentially after applying the **max-pq-cost** parameter, to check if the path from the PQ node to the destination does not traverse the protected node. Setting this parameter to a lower value means the LFA SPFs will use less computation time and resources but may result in not finding a node-protect repair tunnel.

Values 1 to 32

Default 16

Platforms

7705 SAR Gen 2

node-protect

Syntax

[**no**] **node-protect**

Context

[Tree] (config>router>ospf>loopfree-alternates>ti-lfa node-protect)

Full Context

configure router ospf loopfree-alternates ti-lfa node-protect

Description

This command enables node-protect in which the router prefers a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node.

The **no** form of this command disables node-protect.

Default

no node-protect

Platforms

7705 SAR Gen 2

19.47 node-sid

node-sid

Syntax

node-sid index *index-value* [**clear-n-flag**]

node-sid label *label-value* [**clear-n-flag**]

no node-sid

Context

[Tree] (config>router>ospf>area>interface node-sid)

Full Context

configure router ospf area interface node-sid

Description

This command assigns a node SID index or label value to the prefix representing the primary address of a network interface of type system or loopback. A separate SID value can be configured for each IPv4 and IPv6 primary address of the interface. The secondary address of an IPv4 interface cannot be assigned a node SID index and does not inherit the SID of the primary IPv4 address.

In OSPFv2 and OSPFv3, the node SID is configured in the primary area but is inherited in any other area in which the interface is added as secondary.

This command fails if the network interface is not of type loopback or if the interface is defined in an IES or VPRN context. Assigning the same SID index or label value to the same interface in two different IGP instances is not allowed within the same node.

The value of the label or index SID is taken from the range configured for this IGP instance. When using the global mode of operation, the segment routing module checks that the same index or label value is not assigned to more than one loopback interface address. When using the per-instance mode of operation, this check is not required because the index, and therefore, the label ranges of IGP instances are not allowed to overlap.

The **clear-n-flag** option allows the user to clear the N-flag (node-sid flag) in an OSPF or OSPF3 prefix SID sub-TLV originated for the prefix of a loopback interface on the system. By default, the prefix SID sub-TLV for the prefix of a loopback interface is tagged as a node SID; that is, it belongs to this node only. However, to configure and advertise an anycast SID using the same loopback interface prefix on multiple nodes, the user must clear the N-flag to assure interoperability with third-party implementations, which may perform a strict check on the receive end and drop duplicate prefix SID sub-TLVs when the N-flag is set.

The SR OS implementation is relaxed on the receive end and accepts duplicate prefix SIDs with the N-flag set or clear. SR OS will resolve to the closest owner, or owners if ECMP, of the prefix SID cost-wise.

Parameters

index-value

Specifies the node SID index value.

Values 0 to 4294967295

label-value

Specifies the node SID label value.

Values 0 to 4294967295

clear-n-flag

Clears the node SID flag.

Default no clear-n-flag

Platforms

7705 SAR Gen 2

node-sid

Syntax

node-sid index [**0..4294967295**]

node-sid label [**1..4294967295**]

no node-sid

Context

[Tree] (config>router>ospf>area>if>flex-algo node-sid)

Full Context

configure router ospf area interface flex-algo node-sid

Description

This command configures a flexible algorithm-aware node SID label.

The **no** form of this command removes the configured node SID label.

Default

no node-sid

Platforms

7705 SAR Gen 2

node-sid

Syntax

node-sid

no node-sid

Context

[Tree] (config>router>segment-routing>sr-mpls>prefix-sids node-sid)

Full Context

configure router segment-routing sr-mpls prefix-sids node-sid

Description

This command sets the N-flag for the SR SID. The N-flag should be set when the prefix SID is a node SID for the primary prefix. If the N-flag is not set, the SR SID is an SR anycast SID.

The **no** form of this command removes the assigned node SID.

Default

no node-sid

Platforms

7705 SAR Gen 2

19.48 nokia-combined-modules

nokia-combined-modules

Syntax

[**no**] **nokia-combined-modules**

Context

[Tree] (config>system>management-interface>yang-modules nokia-combined-modules)

Full Context

configure system management-interface yang-modules nokia-combined-modules

Description

This command enables support of the "combined" Nokia SR OS YANG files for both configuration and state data in the NETCONF server.

When **management-interface configuration-mode** is set to **classic**, attempts to access (read or write) the configuration using the Nokia configuration modules or namespace via NETCONF results in errors, even if **nokia-combined-modules** or **nokia-submodules** is enabled.

This command and the **nokia-submodules** command cannot both be enabled at the same time.

The **no** form of this command disables support of the combined Nokia SR OS YANG files.

Default

nokia-combined-modules

Platforms

7705 SAR Gen 2

19.49 nokia-grpc-rpc-authorization

nokia-grpc-rpc-authorization

Syntax

[**no**] **nokia-grpc-rpc-authorization**

Context

[Tree] (config>system>security>tacplus>service-request nokia-grpc-rpc-authorization)

[Tree] (config>service>vprn>aaa>remote-servers>tacplus>service-request nokia-grpc-rpc-authorization)

Full Context

configure system security tacplus service-request nokia-grpc-rpc-authorization

configure service vprn aaa remote-servers tacplus service-request nokia-grpc-rpc-authorization

Description

This command enables the nokia-grpc-rpc-authorization service to be requested from the TACACS+ server after successful authentication.

The **no** form of this command disables the nokia-grpc-rpc-authorization service from being requested from the TACACS+ server.

Default

no nokia-grpc-rpc-authorization

Platforms

7705 SAR Gen 2

19.50 nokia-netconf-base-op-authorization

nokia-netconf-base-op-authorization

Syntax

[**no**] **nokia-netconf-base-op-authorization**

Context

[Tree] (config>system>security>tacplus>service-request nokia-netconf-base-op-authorization)

[Tree] (config>service>vprn>aaa>remote-servers>tacplus>service-request nokia-netconf-base-op-authorization)

Full Context

configure system security tacplus service-request nokia-netconf-base-op-authorization

configure service vprn aaa remote-servers tacplus service-request nokia-netconf-base-op-authorization

Description

This command enables the nokia-netconf-base-op-authorization service to be requested from the TACACS + server after successful authentication.

The **no** form of this command disables that the nokia-netconf-base-op-authorization service from being requested from the TACACS+ server.

Default

no nokia-netconf-base-op-authorization

Platforms

7705 SAR Gen 2

19.51 nokia-submodules

nokia-submodules

Syntax

[**no**] **nokia-submodules**

Context

[Tree] (config>system>management-interface>yang-modules nokia-submodules)

Full Context

configure system management-interface yang-modules nokia-submodules

Description

This command enables support of the alternative submodule-based packaging of the Nokia SR OS YANG files for both configuration and state data in the SR OS NETCONF server.

When **management-interface configuration-mode** is set to **classic**, attempts to access (read or write) the configuration using the Nokia configuration modules or namespace via NETCONF results in errors, even if **nokia-combined-modules** or **nokia-submodules** is enabled.

This command and the **nokia-combined-modules** command cannot both be enabled at the same time.

The **no** form of this command disables support of submodule-based packaging of the Nokia SR OS YANG files.

Default

no nokia-submodules

Platforms

7705 SAR Gen 2

19.52 nokia-user

nokia-user

Syntax

[**no**] **nokia-user**

Context

[Tree] (config>system>security>tacplus>service-request nokia-user)

[Tree] (config>service>vprn>aaa>remote-servers>tacplus>service-request nokia-user)

Full Context

configure system security tacplus service-request nokia-user

configure service vprn aaa remote-servers tacplus service-request nokia-user

Description

This command enables the nokia-netconf-base-op-authorization service to be requested from the TACACS + server after successful authentication

The **no** form of this command disables the nokia-netconf-base-op-authorization service from being requested from the TACACS+ server.

Default

no nokia-user

Platforms

7705 SAR Gen 2

19.53 non-dr-attract-traffic

non-dr-attract-traffic

Syntax

[**no**] **non-dr-attract-traffic**

Context

[Tree] (config>service>vprn>pim non-dr-attract-traffic)

Full Context

configure service vprn pim non-dr-attract-traffic

Description

This command specifies whether the router should ignore the designated router state and attract traffic even when it is not the designated router.

An operator can configure an interface (router or IES or VPRN interfaces) to IGMP and PIM. The interface IGMP state will be synchronized to the backup node if it is associated with the redundant peer port. The interface can be configured to use PIM which will cause multicast streams to be sent to the elected DR only. The DR will also be the router sending traffic to the DSLAM. Since it may be required to attract traffic to both routers a flag non-dr-attract-traffic can be used in the PIM context to have the router ignore the DR state and attract traffic when not DR. While using this flag, the router may not send the stream down to the DSLAM while not DR.

When enabled, the designated router state is ignored. When disabled, **no non-dr-attract-traffic**, the designated router value is honored.

Default

no non-dr-attract-traffic

Platforms

7705 SAR Gen 2

non-dr-attract-traffic

Syntax

[**no**] **non-dr-attract-traffic**

Context

[Tree] (config>router>pim non-dr-attract-traffic)

Full Context

configure router pim non-dr-attract-traffic

Description

This command specifies whether the router should ignore the designated router state and attract traffic even when it is not the designated router.

An operator can configure an interface (router or IES or VPRN interfaces) to IGMP and PIM. The interface state will be synchronized to the backup node if it is associated with the redundant peer port. The interface can be configured to use PIM which will cause multicast streams to be sent to the elected DR only. The DR will also be the router sending traffic to the DSLAM. Since it may be required to attract traffic to both routers a flag non-dr-attract-traffic can be used in the PIM context to have the router ignore the DR state and attract traffic when not DR. While using this flag, the router may not send the stream down to the DSLAM while not DR.

When enabled, the designated router state is ignored.

The **no** form of this command the designated router value is honored.

Default

no non-dr-attract-traffic

Platforms

7705 SAR Gen 2

19.54 notification

notification

Syntax

[**no**] **notification**

Context

[Tree] (config>port>ethernet>lldp>dstmac notification)

Full Context

configure port ethernet lldp dest-mac notification

Description

This command enables LLDP notifications.

The **no** form of this command disables LLDP notifications.

Default

no notification

Platforms

7705 SAR Gen 2

notification

Syntax

notification [**neighbor** *ip-address* | **group** *name*]

no notification

Context

[Tree] (debug>router>bgp notification)

Full Context

debug router bgp notification

Description

This command decodes and logs all sent and received notification messages in the debug log.

The **no** form of this command disables the debugging.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x [-interface] (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H
- d: [0 to 255]D
- interface: up to 32 characters for link local addresses

group *name*

Debugs only events affecting the specified peer group name, up to 64 characters, and associated neighbors.

Platforms

7705 SAR Gen 2

notification

Syntax

[**no**] **notification**

Context

[Tree] (config>lag>lldp-member-template>dstmac notification)

Full Context

configure lag lldp-member-template dest-mac notification

Description

This command enables LLDP notifications.

The **no** form of this command disables LLDP notifications.

Default

no notification

Platforms

7705 SAR Gen 2

19.55 notification-bundling

notification-bundling

Syntax

notification-bundling

Context

[Tree] (config>system>telemetry notification-bundling)

Full Context

configure system telemetry notification-bundling

Description

Commands in this context configure SubscribeResponse notification bundling.

Platforms

7705 SAR Gen 2

19.56 notification-interval

notification-interval

Syntax

notification-interval *time*

no notification-interval

Context

[Tree] (config>system>lldp notification-interval)

Full Context

configure system lldp notification-interval

Description

This command configures the minimum time between change notifications.

The **no** form of this command reverts to the default value.

Default

no notification-interval

Parameters

time

Specifies the minimum time, in seconds, between change notifications.

Values 5 to 3600

Default 5

Platforms

7705 SAR Gen 2

19.57 notify-dest-change

notify-dest-change

Syntax

[**no**] **notify-dest-change**

Context

[Tree] (config>filter>redirect-policy notify-dest-change)

Full Context

configure filter redirect-policy notify-dest-change

Description

This command instructs the system to send notifications (Log, SNMP, …) when the active destination of a redirect policy changes. No notification is sent when there are no more active destinations (as this is covered by a specific other notification). Notifications can be controlled (using the **config**>**log**>**event-control** command) using application ID *2017* and event-name *tFilterRPActiveDstChangeEvent*.

The **no** form of the command disables notification generation.

Default

no notify-dest-change

Platforms

7705 SAR Gen 2

19.58 nsp-proxy

nsp-proxy

Syntax

[**no**] **nsp-proxy**

Context

[Tree] (debug>system nsp-proxy)

Full Context

debug system nsp-proxy

Description

This command enables debugging for NSP proxy.

The **no** form of this command disables debugging for NSP proxy.

Default

no nsp-proxy

Platforms

7705 SAR Gen 2

19.59 nssa

nssa

Syntax

[**no**] **nssa**

Context

[Tree] (config>service>vprn>ospf3>area nssa)

[Tree] (config>service>vprn>ospf>area nssa)

Full Context

configure service vprn ospf3 area nssa

configure service vprn ospf area nssa

Description

This command creates the context to configure an OSPF Not So Stubby Area (NSSA) and adds/removes the NSSA designation from the area.

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is that an NSSA has the capability to flood external routes that it learns throughout its area and via an ABR to the entire OSPF domain.

Existing virtual links of a non-stub or NSSA area are removed when the designation is changed to NSSA or stub.

An area can be designated as stub or NSSA but never both at the same time.

By default, an area is not configured as an NSSA area.

The **no** form of this command removes the NSSA designation and configuration context from the area.

Default

no nssa — The OSPF area is not an NSSA.

Platforms

7705 SAR Gen 2

nssa

Syntax

[**no**] **nssa**

Context

[Tree] (config>router>ospf3>area nssa)

[Tree] (config>router>ospf>area nssa)

Full Context

configure router ospf3 area nssa

configure router ospf area nssa

Description

This command creates the context to configure an OSPF or OSPF3 Not So Stubby Area (NSSA) and adds/removes the NSSA designation from the area.

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is an NSSA has the capability to flood external routes that it learns throughout its area and via an ABR to the entire OSPF or OSPF3 domain.

Existing virtual links of a non-stub or NSSA area will be removed when the designation is changed to NSSA or stub.

An area can be designated as stub or NSSA but never both at the same time.

By default, an area is not configured as an NSSA area.

The **no** form of this command removes the NSSA designation and configuration context from the area.

Default

no nssa

Platforms

7705 SAR Gen 2

19.60 nssa-range

nssa-range

Syntax

nssa-range [*ip-address*]

no nssa-range

Context

[Tree] (debug>router>ospf3 nssa-range)

[Tree] (debug>router>ospf nssa-range)

Full Context

debug router ospf3 nssa-range

debug router ospf nssa-range

Description

This command enables debugging for an NSSA range.

Parameters

ip-address

Specifies the IPv4 or IPv6 address range to debug OSPF or OSPF3 leaks.

Values ipv4-address:

- a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

19.61 ntp

ntp

Syntax

[**no**] **ntp**

Context

[Tree] (config>service>vprn ntp)

Full Context

configure service vprn ntp

Description

Commands in this context configure Network Time Protocol (NTP) and its operation. It also enables NTP server mode within the VPRN routing instance so that the router will respond to NTP requests from external clients received inside the VPRN.

The **no** form of this command stops the execution of NTP and removes its configuration.

Platforms

7705 SAR Gen 2

ntp

Syntax

[**no**] **ntp**

Context

[Tree] (config>system>time ntp)

Full Context

configure system time ntp

Description

Commands in this context configure Network Time Protocol (NTP) and its operation. This protocol defines a method to accurately distribute and maintain time for network elements. Furthermore, this capability allows for the synchronization of clocks between the various network elements.

The **no** form of the command stops the execution of NTP and remove its configuration.

Default

ntp

Platforms

7705 SAR Gen 2

ntp

Syntax

ntp [**router** *router-instance*] [**interface** *ip-int-name*]

Context

[Tree] (debug>system ntp)

Full Context

debug system ntp

Description

This command enables and configures debugging for NTP.

The **no** form of the command disables debugging for NTP.

Parameters

router-instance

Specifies the router name or CPM router instance.

Values *router-name* | *vprn-svc-id*

router-name – "Base", "management"

vprn-svc-id – 1 to 2147483647

Default Base

ip-int-name

Specifies the name of the IP interface. The name can be up to 32 characters and must begin with a letter. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

19.62 ntp-reply

ntp-reply

Syntax

[**no**] **ntp-reply**

Context

[Tree] (config>service>ies>if>vrrp ntp-reply)

[Tree] (config>service>ies>if>ipv6>vrrp ntp-reply)

Full Context

configure service ies interface vrrp ntp-reply

configure service ies interface ipv6 vrrp ntp-reply

Description

This command enables the reception and response to NTP Requests directed at the VRRP virtual IP address. This behavior only applies the router currently acting as the master VRRP router.

The **no** form of this command disables NTP Requests from being processed.

Default

no ntp-reply

Platforms

7705 SAR Gen 2

ntp-reply

Syntax

[**no**] **ntp-reply**

Context

[Tree] (config>service>vprn>if>vrrp ntp-reply)

[Tree] (config>service>vprn>if>ipv6>vrrp ntp-reply)

Full Context

configure service vprn interface vrrp ntp-reply

configure service vprn interface ipv6 vrrp ntp-reply

Description

This command enables the reception and response to NTP Requests directed at the VRRP virtual IP address. This behavior only applies the router currently acting as the master VRRP router.

The **no** form of this command disables NTP Requests from being processed.

Default

no ntp-reply

Platforms

7705 SAR Gen 2

ntp-reply

Syntax

[**no**] **ntp-reply**

Context

[Tree] (config>router>if>vrrp ntp-reply)

[Tree] (config>router>if>ipv6>vrrp ntp-reply)

Full Context

configure router interface vrrp ntp-reply

configure router interface ipv6 vrrp ntp-reply

Description

This command enables the reception and response to NTP Requests directed at the VRRP virtual IP address. This behavior only applies the router currently acting as the master VRRP router.

The **no** form of this command disables NTP Requests from being processed.

Default

no ntp-reply

Platforms

7705 SAR Gen 2

19.63 ntp-server

ntp-server

Syntax

ntp-server [**authenticate**]

no ntp-server

Context

[Tree] (config>system>time>ntp ntp-server)

Full Context

configure system time ntp ntp-server

Description

This command configures the node to assume the role of an NTP server. Unless the **server** command is used, this node will function as an NTP client only and will not distribute the time to downstream network elements.

Default

no ntp-server

Parameters

authenticate

Specifies to make authentication a requirement (optional). If authentication is required, the authentication key-id received in a message must have been configured in the **authentication-key** command, and that key-id type and key value must also match.

The authentication key from the received messages will be used for the transmitted messages.

Platforms

7705 SAR Gen 2

19.64 number

number

Syntax

number {**eq** | **neq** | **lt** | **lte** | **gt** | **gte**} *event-id*

no number

Context

[Tree] (config>service>vprn>log>filter>entry>match number)

Full Context

configure service vprn log filter entry match number

Description

This command adds an SR OS application event number as a match criterion.

SR OS event numbers uniquely identify a specific logging event within an application.

Only one **number** command can be entered per event filter entry. The latest **number** command overwrites the previous command.

The **no** form of this command removes the event number as a match criterion.

Default

no event-number — No event ID match criterion is specified.

Parameters

eq | neq | lt | lte | gt | gte

Specifies the type of match. Valid operators are listed below.

Values

Table 71: Valid Operators

Operator	Note
eq	equal to
neq	not equal to
lt	less than
lte	less than or equal to
gt	greater than
gte	greater than or equal to

event-id

Specifies the event ID, expressed as a decimal integer.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

number

Syntax

number {**eq** | **neq** | **lt** | **lte** | **gt** | **gte**} *event-id*

no number

Context

[Tree] (config>log>filter>entry>match number)

Full Context

configure log filter entry match number

Description

This command adds an SR OS application event number as a match criterion.

SR OS event numbers uniquely identify a specific logging event within an application.

Only one **number** command can be entered per event filter entry. The latest **number** command overwrites the previous command.

The **no** form of this command removes the event number as a match criterion.

Parameters

eq | neq | lt | lte | gt | gte

Specifies the type of match. Valid operators are listed in Table 72: Valid Operators.

Table 72: Valid Operators

Operator	Notes
eq	equal to
neq	not equal to
lt	less than
lte	less than or equal to
gt	greater than
gte	greater than or equal to

event-id

The event ID, expressed as a decimal integer.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

19.65 number-down

number-down

Syntax

[**no**] **number-down** *number-of-lag-ports-down*

Context

[Tree] (config>vrrp>policy>priority-event>lag-port-down number-down)

Full Context

configure vrrp policy priority-event lag-port-down number-down

Description

This command creates a context to configure an event set threshold within a lag-port-down priority control event.

The **number-down** command defines a sub-node within the **lag-port-down** event and is uniquely identified with the *number-of-lag-ports-down* parameter. Each **number-down** node within the same **lag-port-down** event node must have a unique *number-of-lag-ports-down* value. Each **number-down** node has its own **priority** command that takes effect whenever that node represents the current threshold.

The total number of sub-nodes (uniquely identified by the *number-of-lag-ports-down* parameter) allowed in a single **lag-port-down** event is equal to the total number of possible physical ports allowed in a LAG.

A **number-down** node is not required for each possible number of ports that could be down. The active threshold is always the closest lower threshold. When the number of ports down equals a given threshold, that is the active threshold.

The **no** form of the command deletes the event set threshold. The threshold may be removed at any time. If the removed threshold is the current active threshold, the event set thresholds must be re-evaluated after removal.

Default

no number-down — No threshold for the LAG priority event is created.

Parameters

number-of-lag-ports-down

The number of LAG ports down to create a set event threshold. This is the active threshold when the number of down ports in the LAG equals or exceeds *number-of-lag-ports-down*, but does not equal or exceed the next highest configured *number-of-lag-ports-down*.

Values 1 to 64 (applies to 64-link LAG) 1 to 32 (applies to other LAGs)

Platforms

7705 SAR Gen 2

19.66 number-retries

number-retries

Syntax

number-retries *number-retries*

no number-retries

Context

[Tree] (config>service>vpls>mac-move number-retries)

[Tree] (config>service>template>vpls-template>mac-move number-retries)

Full Context

configure service vpls mac-move number-retries

configure service template vpls-template mac-move number-retries

Description

This command configures the number of times retries are performed for re-enabling the SAP/SDP.

Default

number-retries 3

Parameters

number-retries

Specifies number of retries for re-enabling the SAP/SDP. A zero (0) value indicates unlimited number of retries.

Values 0 to 255

Platforms

7705 SAR Gen 2

20 o Commands

20.1 oam

oam

Syntax

oam

Context

[Tree] (oam)

Full Context

oam

Description

Commands in this context use the OAM test suite.

Platforms

7705 SAR Gen 2

oam

Syntax

oam

Context

[Tree] (debug oam)

Full Context

debug oam

Description

This command enables OAM debugging.

Platforms

7705 SAR Gen 2

20.2 oam-pm

oam-pm

Syntax

oam-pm session *session-name* {**dm** | **dmm** | **lmm** | **slm** | **twamp-light**} { **start** | **stop**}

Context

[Tree] (oam oam-pm)

Full Context

oam oam-pm

Description

This command allows the operator to start and stop on-demand OAM-PM sessions.

Parameters

session-name

Identifies the session name, up to 32 characters, that the test is associated with.

dm

Specifies the MPLS delay measurement test that is affected by the command.

dmm

Specifies the DMM test that is affected by the command.

lmm

Specifies the LMM test that is affected by the command.

slm

Specifies the SLM test that is affected by the command.

twamp-light

Specifies the TWAMP-light test that is affected by the command.

start

Manually starts the test.

stop

Manually stops the test.

Platforms

7705 SAR Gen 2

oam-pm

Syntax

oam-pm

Context

[Tree] (config oam-pm)

Full Context

configure oam-pm

Description

This is the top level context that contains the configuration parameters that defines storage parameters (including binning structures), availability/resiliency and the individual proactive, and on-demand tests used to gather the performance/statistical information.

Platforms

7705 SAR Gen 2

20.3 ocsp

ocsp

Syntax

[**no**] **ocsp**

Context

[Tree] (debug>certificate ocsp)

Full Context

debug certificate ocsp

Description

This command enables debug output of the OCSP protocol for a CA profile.

The **no** form of this command disables the debug output.

Platforms

7705 SAR Gen 2

ocsp

Syntax

ocsp

Context

[Tree] (config>system>security>pki>ca-profile ocsp)

Full Context

configure system security pki ca-profile ocsp

Description

Commands in this context configure OCSP parameters.

Platforms

7705 SAR Gen 2

20.4 offer-time

offer-time

Syntax

offer-time [**min** *minutes*] [**sec** *seconds*]

no offer-time

Context

[Tree] (config>service>vprn>dhcp>server>pool offer-time)

[Tree] (config>router>dhcp>server>pool offer-time)

Full Context

configure service vprn dhcp local-dhcp-server pool offer-time

configure router dhcp local-dhcp-server pool offer-time

Description

This command configures the time interval during which a DHCP offer is valid.

The **no** form of this command reverts to the default.

Default

offer-time min 1

Parameters

time

Specifies the offer time.

Values

min *minutes*	0 to 10
sec *seconds*	0 to 59

Platforms

7705 SAR Gen 2

20.5 offset

offset

Syntax

offset *offset*

Context

[Tree] (config>system>time>dst-zone offset)

Full Context

configure system time dst-zone offset

Description

This command specifies the number of minutes that are added to the time when summer time takes effect. The same number of minutes is subtracted from the time when the summer time ends.

Default

offset 60

Parameters

offset

Specifies the number of minutes added to the time at the beginning of summer time and subtracted at the end of summer time, expressed as an integer.

Values 0 to 60

Default 60

Platforms

7705 SAR Gen 2

20.6 on-cac-failure

on-cac-failure

Syntax

[**no**] **on-cac-failure**

Context

[Tree] (config>router>rsvp>te-threshold-update on-cac-failure)

Full Context

configure router rsvp te-threshold-update on-cac-failure

Description

This command is used to enable a CAC failure-triggered IGP update.

The **no** form of this command should reset on-cac-failure to the default value and disable the CAC failure-triggered IGP update.

Default

no on-cac-failure

Platforms

7705 SAR Gen 2

20.7 on-link

on-link

Syntax

[**no**] **on-link**

Context

[Tree] (config>service>vprn>router-advert>if>prefix on-link)

Full Context

configure service vprn router-advertisement interface prefix on-link

Description

This command specifies whether the prefix can be used for onlink determination.

Default
on-link

Platforms
7705 SAR Gen 2

on-link

Syntax
[**no**] **on-link**

Context
[Tree] (config>router>router-advert>if>prefix on-link)

Full Context
configure router router-advertisement interface prefix on-link

Description
This command specifies whether the prefix can be used for on link determination.

Default
on link

Platforms
7705 SAR Gen 2

20.8 open

open

Syntax
open [**neighbor** *ip-address* | **group** *name*]
no open

Context
[Tree] (debug>router>bgp open)

Full Context
debug router bgp open

Description

This command decodes and logs all sent and received open messages in the debug log.

The **no** form of this command disables debugging.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x [-interface] (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H
- d: [0 to 255]D
- interface: up to 32 characters for link local addresses

group *name*

Debugs only events affecting the specified peer group name, up to 64 characters, and associated neighbors.

Platforms

7705 SAR Gen 2

20.9 oper-group

oper-group

Syntax

oper-group *name*

no oper-group

Context

[Tree] (config>service>epipe oper-group)

Full Context

configure service epipe oper-group

Description

This command associates an operational group to the status of the Epipe. When this oper-group is used in Epipes with static VXLAN or BGP-EVPN, the oper-group behaves as follows:

- The Epipe (and the oper-group) goes down if a SAP or spoke SDP goes operationally down because of an administrative shutdown, service shutdown, or non-DF status as a result of EVPN multi-homing single-active election.
- The Epipe (and oper-group) goes down if the Epipe's EVPN destination is removed (because of an EVPN AD per-EVI route withdrawal, for example).
- The Epipe (and oper-group) does not go down if a static VXLAN destination exists and the egress VTEP is not in the global route table.

The operational group must be monitored in a different service and not in the service where it is defined.

The **no** version of this command removes the oper-group association.

Parameters

name

Specifies the name of the **oper-group**, up to 32 characters.

Platforms

7705 SAR Gen 2

oper-group

Syntax

oper-group *name*

no oper-group

Context

[Tree] (config>service>vpls>bgp-evpn>mpls oper-group)

[Tree] (config>service>epipe>bgp-evpn>mpls oper-group)

Full Context

configure service vpls bgp-evpn mpls oper-group

configure service epipe bgp-evpn mpls oper-group

Description

This command adds the BGP EVPN MPLS, SRv6, or VXLAN instance or Ethernet Segment (ES) as a member of the operational group.

When configured on an ES, the state of the operational group depends on the state of the SAPs contained in the ES. The operational group transitions to up if at least one SAP in the ES is up. The operational group goes down when all the associated SAPs are operationally down. The ES operational group should be monitored on the LAG associated with the ES, along with single-active multi-homing, so that the NDF state can be signaled to the CE by LAG standby signaling.

When configured on a BGP EVPN instance, the operational group is up when it is either empty (meaning that the operational group has no members) or at least an EVPN destination is created under the EVPN instance added as member. When configured, no other SAP, SDP binding, or BGP EVPN instance can be added to the same operational group within the same or a different service.

The operational group is down when the following apply:

- on a BGP EVPN instance
- the service is disabled
- the BGP EVPN MPLS, VXLAN or SRv6 instance are disabled
- all the EVPN destinations in the instance are removed

Default

no oper-group

Parameters

name

Specifies the name of the operational group, up to 32 characters.

Platforms

7705 SAR Gen 2

oper-group

Syntax

oper-group *group-name*

no oper-group

Context

[Tree] (config>service>epipe>sap oper-group)

Full Context

configure service epipe sap oper-group

Description

This command configures the operational group identifier.

The no form of this command removes the group name from the configuration.

Parameters

group-name

Specifies the Operational-Group identifier up to 32 characters in length.

Platforms

7705 SAR Gen 2

oper-group

Syntax

oper-group *group-name*

no oper-group

Context

[Tree] (config>service>vpls>bgp>pw-template-binding oper-group)

[Tree] (config>service>vpls>sap oper-group)

[Tree] (config>service>vpls>spoke-sdp oper-group)

Full Context

configure service vpls bgp pw-template-binding oper-group

configure service vpls sap oper-group

configure service vpls spoke-sdp oper-group

Description

This command associates the context to which it is configured to the operational group specified in the *group-name*. The **oper-group** *oper-name* must be already configured under **config>service** context before its name is referenced in this command.

The **no** form of this command removes the association.

Default

no oper-group

Parameters

group-name

Specifies a character string of maximum 32 ASCII characters identifying the group instance.

Platforms

7705 SAR Gen 2

oper-group

Syntax

oper-group *group-name*

no oper-group

Context

[Tree] (config>service>ies>if>vrrp oper-group)

Full Context

configure service ies interface vrrp oper-group

Description

This command configures VRRP to associate with an operational group. When associated, VRRP notifies the operational group of its state changes so that other protocols can monitor it to provide a redundancy mechanism. When VRRP is the master router (MR), the operational group is up and is down for all other VRRP states.

The **no** form of this command removes the association.

Default

no oper-group

Parameters

group-name

Specifies the operational group identifier up to 32 characters in length.

Platforms

7705 SAR Gen 2

oper-group

Syntax

oper-group *group-name*

no oper-group

Context

[Tree] (config>service>vprn>if>vrrp oper-group)

[Tree] (config>service>vprn>if>ipv6>vrrp oper-group)

Full Context

configure service vprn interface vrrp oper-group

configure service vprn interface ipv6 vrrp oper-group

Description

This command configures VRRP to associate with an operational group. When associated, VRRP notifies the operational group of its state changes so that other protocols can monitor it to provide a redundancy mechanism. When VRRP is the master router (MR), the operational group is up and is down for all other VRRP states.

The **no** form of this command removes the association.

Default

no oper-group — No operational group is configured.

Parameters

group-name

Specifies the operational group identifier, up to 32 characters in length.

Platforms

7705 SAR Gen 2

oper-group

Syntax

oper-group *group-name*

no oper-group

Context

[Tree] (config>router>if>ipv6>vrrp oper-group)

[Tree] (config>router>if>vrrp oper-group)

Full Context

configure router interface ipv6 vrrp oper-group

configure router interface vrrp oper-group

Description

This command configures VRRP to associate with an operational group. When associated, VRRP notifies the operational group of its state changes so that other protocols can monitor it to provide a redundancy mechanism. When VRRP is the master router (MR), the operational group is up; the operational group is down for all other VRRP states.

The **no** form of the command removes the association.

Default

no oper-group — No operational group is configured.

Parameters

group-name

Specifies the operational group identifier, up to 32 characters.

Platforms

7705 SAR Gen 2

oper-group

Syntax

oper-group *group-name* [**create**]

no oper-group *group-name*

Context

[Tree] (config>service oper-group)

Full Context

configure service oper-group

Description

This command creates a system-wide group (operational group) name which can be used to associate a number of service objects (for example, SAPs or pseudowires). The status of the group is derived from the status of its members. The status of the group can then be used to influence the status of non-member objects. For example, when a group status is marked as down, the object(s) that monitor the group change their status accordingly.

The **no** form of the command removes the group. All the object associations need to be removed before the no form of the command can be executed.

Default

no oper-group

Parameters

group-name

Specifies the operational group identifier up to 32 characters in length.

create

This keyword is required when first creating the configuration context. Once the context is created, it is possible to navigate into the context without the **create** keyword.

Platforms

7705 SAR Gen 2

20.10 oper-members

oper-members

Syntax

oper-members *oper-members*

no oper-members

Context

[Tree] (config>service>vprn>isis>link-group>level oper-members)

Full Context

configure service vprn isis link-group level oper-members

Description

This command sets the threshold for the minimum number of operational links for the associated link-group. If the number of operational links drops below this threshold, the configured offsets are applied. For example, oper-members=3. The metric of the member interfaces is increased when the number of interfaces is lower than 3.

The **no** form of this command reverts the oper-members limit to 1.

Default

no oper-members

Parameters

oper-members

Specifies the number of operational members.

Values 0 to 8

Platforms

7705 SAR Gen 2

oper-members

Syntax

oper-members [*value*]

no oper-members

Context

[Tree] (config>router>isis>link-group>level oper-members)

Full Context

configure router isis link-group level oper-members

Description

This command sets the threshold for the minimum number of operational links for the associated link-group. If the number of operational links drops below this threshold, the configured offsets are applied. For example, oper-members=3. The metric of the member interfaces is increased when the number of interfaces is lower than 3.

The **no** form of this command reverts the **oper-members** limit to 1.

Default

oper-members 1

Parameters

value

Specifies the threshold for operational members.

Values 1 to 8

Platforms

7705 SAR Gen 2

20.11 operations

operations

Syntax

operations

Context

[Tree] (config>system>management-interface operations)

Full Context

configure system management-interface operations

Description

Commands in this context configure parameters associated with operational commands in model-driven interfaces.

Platforms

7705 SAR Gen 2

20.12 optical-line-system

optical-line-system

Syntax

optical-line-system

no optical-line-system

Context

[Tree] (config>port>transceiver optical-line-system)

Full Context

configure port transceiver optical-line-system

Description

This command enables the QSFP-LS pluggable optical line system, used in conjunction with 400G ZR/ZR+ coherent pluggable modules.

Note: A user cannot enable a transceiver as both a Digital Coherent Optic (DCO) and an Open Line System (OLS) at the same time.

The **no** form of this command reverts to the default.

Default

no optical-line-system

Platforms

7705 SAR Gen 2

20.13 optimal-route-reflection

optimal-route-reflection

Syntax

optimal-route-reflection

Context

[Tree] (config>router>bgp optimal-route-reflection)

Full Context

configure router bgp optimal-route-reflection

Description

This command creates the optimal route reflection context.

Platforms

7705 SAR Gen 2

20.14 option

option

Syntax

option *dhcp-option-number* {**present** | **absent**}

option *dhcp-option-number* **match hex** *hex-string* [**exact**] [**invert-match**]

option *dhcp-option-number* **match string** *ascii-string* [**exact**] [**invert-match**]

no option

Context

[Tree] (config>filter>dhcp-filter>entry option)

Full Context

configure filter dhcp-filter entry option

Description

This command configures match criteria for the DHCP filter policy entry.

The **no** form of this command reverts to the default.

Parameters

dhcp-option-number

Specifies the DHCP option number.

Values 0 to 255

present

Specifies that the related DHCP option must be present.

absent

Specifies that the related DHCP option must be absent.

hex-string

Specifies that the option must partially match a specified hex string.

Values 0x0 to 0xFFFFFFFF (up to 254 hex nibbles)

ascii-string

Specifies that the option must partially match a specified ASCII string, up to 127 characters.

exact

Specifies that this option requires an exact match of a hex or ASCII string.

invert-match

Requires the option not to partially match.

Platforms

7705 SAR Gen 2

option

Syntax

option *dhcp6-option-number* {**present** | **absent**}

option *dhcp6-option-number* **match hex** *hex-string* [**exact**] [**invert-match**]

option *dhcp6-option-number* **match string** *ascii-string* [**exact**] [**invert-match**]

no option

Context

[Tree] (config>filter>dhcp6-filter>entry option)

Full Context

configure filter dhcp6-filter entry option

Description

This command configures match criteria for the DHCP6 filter policy entry.

The **no** form of this command reverts to the default.

Parameters

dhcp-option-number

Specifies the DHCP6 option number.

Values 0 to 255

present

Specifies that the related DHCP6 option must be present.

absent

Specifies that the related DHCP6 option must be absent.

match hex *hex-string*

Specifies that the option must (partially) match a specified hex string.

Values 0x0 to 0xFFFFFFFF (up to 254 hex nibbles)

match string *ascii-string*

Specifies that the option must partially match a specified ASCII string, up to 127 characters.

exact

Specifies that this option requires an exact match of a hex or ASCII string.

invert-match

Requires the option not to partially match.

Platforms

7705 SAR Gen 2

option

Syntax

[**no**] **option**

Context

[Tree] (config>service>vprn>if>ipv6>dhcp6-relay option)

[Tree] (config>service>ies>if>ipv6>dhcp6-relay option)

Full Context

configure service vprn interface ipv6 dhcp6-relay option

configure service ies interface ipv6 dhcp6-relay option

Description

Commands in this context configure DHCPv6 relay information options.

The **no** form of this command disables DHCPv6 relay information options.

Platforms

7705 SAR Gen 2

option

Syntax

[**no**] **option**

Context

[Tree] (config>service>ies>if>dhcp option)

[Tree] (config>service>vpls>sap>dhcp option)

[Tree] (config>service>vprn>if>dhcp option)

[Tree] (config>router>if>dhcp option)

Full Context

configure service ies interface dhcp option

configure service vpls sap dhcp option

configure service vprn interface dhcp option

configure router interface dhcp option

Description

This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.

The **no** form of this command reverts to the default.

Default

no option

Platforms

7705 SAR Gen 2

option

Syntax

option {**basic** | **isis-enhanced**}

no option

Context

[Tree] (config>system>security>keychain>direction>bi>entry option)

Full Context

configure system security keychain direction bi entry option

Description

This command configures allows options to be associated with the authentication key.

Parameters

basic

Specifies that IS-IS should use RFC 5304 encoding of the authentication information. It is only applicable if used with the IS-IS protocol. All other protocols should ignore this configuration command.

isis-enhanced

Specifies that IS-IS should use RFC 5310 encoding of the authentication information. It is only applicable if used with the IS-IS protocol. All other protocols should ignore this configuration command.

Platforms

7705 SAR Gen 2

20.15 option-present

option-present

Syntax

option-present {**true** | **false**}

no option-present

Context

[Tree] (config>filter>ip-filter>entry>match option-present)

Full Context

configure filter ip-filter entry match option-present

Description

This command configures matching packets that contain any IP options in the IP header as an IP filter match criterion.

The **no** form of the command removes the checking of IP options in the IP header as a match criterion.

Default

no option-present

Parameters

true

Specifies matching on all IP packets that contain any IP options in the IP header. A match occurs for all packets that have any IP option present. An option field of zero is considered as no option present.

false

Specifies matching on IP packets that do not have any IP option present in the IP header. (an option field of zero). An option field of zero is considered as no option present.

Platforms

7705 SAR Gen 2

20.16 option60

option60

Syntax

option60 hex *hex-string*

option60 string *ascii-string*

no option60

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>host-ident option60)

Full Context

configure subscriber-mgmt local-user-db ipoe host host-identification option60

Description

This command specifies the Vendor-Identifying Vendor Option to match. Option 60 is encoded as Type-Length-Value (TLV). The *hex-string* portion of Option 60 in the received DHCP request is used for matching. Only the first 32 bytes can be defined here. If Option 60 from the message is longer, those bytes are ignored.

Note:
This command is only used when **option60** is configured as one of the **match-list** parameters.

The **no** form of this command removes **option60** from the configuration.

Parameters

hex-string

Specifies the hexadecimal format for this option.

Values 0x0 to 0xFFFFFFFF(maximum 64 hex nibbles)

ascii-string

Specifies the string format for this option, up to 32 characters.

Platforms

7705 SAR Gen 2

20.17 options

options

Syntax

options

Context

[Tree] (config>service>vprn>dhcp>server>pool options)

[Tree] (config>service>vprn>dhcp6>server>pool options)

[Tree] (config>router>dhcp6>server>pool options)

[Tree] (config>router>dhcp>server>pool options)

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host options)

[Tree] (config>router>dhcp6>server>pool>prefix options)

[Tree] (config>router>dhcp>server>pool>subnet options)

[Tree] (config>service>vprn>dhcp6>server>pool>prefix options)

Full Context

configure service vprn dhcp local-dhcp-server pool options

configure service vprn dhcp6 local-dhcp-server pool options

configure router dhcp6 local-dhcp-server pool options

configure router dhcp local-dhcp-server pool options

configure subscriber-mgmt local-user-db ipoe host options

configure router dhcp6 local-dhcp-server pool prefix options

configure router dhcp local-dhcp-server pool subnet options

configure service vprn dhcp6 local-dhcp-server pool prefix options

Description

Commands in this context configure pool options. The options defined here can be overruled by defining the same option in the local user database.

Platforms

7705 SAR Gen 2

options

Syntax

options

Context

[Tree] (config>system>persistence options)

Full Context

configure system persistence options

Description

This command enables the CLI context to configure persistence options parameters.

Platforms

7705 SAR Gen 2

options

Syntax

options

Context

[Tree] (config>service>vpls>sap>dhcp6>ldra options)

Full Context

configure service vpls sap dhcp6 ldra options

Description

Commands in this context configure forwarding path options.

Platforms

7705 SAR Gen 2

20.18 options6

options6

Syntax

options6

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host options6)

Full Context

configure subscriber-mgmt local-user-db ipoe host options6

Description

Commands in this context configure IPv6 DNS server information in the local user database.

Platforms

7705 SAR Gen 2

20.19 origin

origin

Syntax

origin {**igp** | **egp** | **incomplete** | **any** | **aaa** | **dynamic** | **static** | **bonding | pfcp**}

no origin

Context

[Tree] (config>router>policy-options>policy-statement>entry>from origin)

Full Context

configure router policy-options policy-statement entry from origin

Description

This command configures the match criteria for the origin attribute of the route. The origin attribute is applicable to BGP routes and to the following subscriber-management routes.

- Host routes (for example, IPv4 /32 address, or IPv6 SLAAC prefix) carry the origin attribute with the **aaa**, **dynamic**, or **static** options, depending on the address assignment method. For CUPS hosts, the

pfcp option is used for the origin attribute. Host routes can also be distinguished using the **sub-mgmt** option for the **protocol** command.

- Dynamically provisioned prefixes or loopback addresses carry the origin attribute with the **aaa** or **pfcp** options, depending on the protocol that provides the prefix and address. Dynamic routes can also be distinguished using the **direct** option for the **protocol** command.
- Statically configured prefixes under the subscriber interface do not have an origin attribute.These routes can be distinguished using the **direct** option for the **protocol** command.
- Framed routes for non-CUPS hosts do not have an origin attribute. Framed routes for CUPS hosts use the **pfcp** option for the origin attribute. Alternatively, framed routes can be distinguished using the **managed** option for the **protocol** command.

These values that are specific to subscriber-management routes are never carried in BGP updates as part of the BGP origin attribute and are not visible within the BGP process.

Default

no origin

Parameters

igp

Specifies path-matching information that originates within the local AS.

egp

Specifies path-matching information that originates in another AS.

incomplete

Configures path-matching information learned by another method.

any

Specifies to ignore this criteria.

aaa

Specifies to use the subscriber-host address that originates from AAA.

Values IPv4 — subscriber-management /32 host routes that originate from the RADIUS framed-ip-address VSA other than 255.255.255.254. The 255.255.255.254 returned by the RADIUS indicates that the BNG (NAS) should assign an IP address from its own pool.

IPv6 — subscriber-management routes that originate through framed-ipv6-prefix (SLAAC), delegated-ipv6-prefix (IA_PD) or alc-ipv6-address (IA_NA) RADIUS attributes. It is also applicable to VSA Alc-IPv6-Sub-If-Prefix, where the subscriber interface prefix can originate from RADIUS. This is valid for IPoE and PPPoE type hosts.

dynamic

Specifies to use the subscriber host address that originates from DHCP, DHCPv6, or the local address server.

Values IPv4 — subscriber-management /32 host routes that originate from the DHCP server (local or remote) or RADIUS framed-ip-address=255.255.255.254 (RFC 2865).

IPv6 — subscriber-management routes that are assigned via local DHCPv6 server pools whose name is obtained through the Alc-Delegated-IPv6-Pool (PD pool) and Framed-IPv6-Pool (NA pool) RADIUS attributes, or the local address server whose name is obtained through the Alc-SLAAC-IPv6-Pool (SLAAC pool) RADIUS attribute. This is valid for IPoE and PPPoE type hosts.

For IPoEv6 only, the pool name can also be obtained from ipv6-delegated-prefix-pool (PD pool) and ipv6-wan-address-pool (NA pool) from the LUDB.

static

Specifies to use the subscriber-host address that originates from the local user database.

Values IPv4 — subscriber-management /32 host routes that originate from the LUDB and also covers the RADIUS fallback category (RADIUS falls back to system defaults or to the LUDB).

IPv6 — subscriber-management routes that originate from the LUDB from ipv6-address (IA_NA) or ipv6-prefix (IA_PD), or ipv6-slaac-prefix (SLAAC).

bonding

Specifies to use bonding.

pfcp

Specifies to use routes learned using the PFCP protocol.

Platforms

7705 SAR Gen 2

origin

Syntax

origin {**igp** | **egp** | **incomplete** | *param-name*}

no origin

Context

[Tree] (config>router>policy-options>policy-statement>default-action origin)

[Tree] (config>router>policy-options>policy-statement>entry>action origin)

Full Context

configure router policy-options policy-statement default-action origin

configure router policy-options policy-statement entry action origin

Description

This command sets the BGP origin assigned to routes exported into BGP.

If the routes are exported into protocols other than BGP, this option is ignored.

The **no** form of this command disables setting the BGP origin for the route policy entry.

Default

no origin

Parameters

igp

Sets the path information as originating within the local AS.

egp

Sets the path information as originating in another AS.

incomplete

Sets the path information as learned by some other means.

param-name

The origin parameter variable name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@”.

Platforms

7705 SAR Gen 2

20.20 origin-invalid-unusable

origin-invalid-unusable

Syntax

[**no**] **origin-invalid-unusable**

Context

[Tree] (config>service>vprn>bgp>best-path-selection origin-invalid-unusable)

Full Context

configure service vprn bgp best-path-selection origin-invalid-unusable

Description

When this command is configured, all VPRN BGP routes that have an origin validation state of "Invalid” are considered unusable by the best path selection algorithm, meaning they are not used for forwarding, not advertised to BGP peers, and not eligible for export as a VPN-IP route.

With the default value, VPRN BGP routes with an origin validation state of "Invalid” are usable if they are selected.

Default

no origin-invalid-unusable

Platforms

7705 SAR Gen 2

origin-invalid-unusable

Syntax

[**no**] **origin-invalid-unusable**

Context

[Tree] (config>router>bgp>best-path-selection origin-invalid-unusable)

Full Context

configure router bgp best-path-selection origin-invalid-unusable

Description

When **origin-invalid-unusable** is configured, all routes that have an RPKI origin validation state of 'Invalid' are considered unusable by the best path selection algorithm, meaning they are not used for forwarding and not advertised to BGP peers.

With the default of **no origin-invalid-unusable**, routes with an RPKI origin validation state of 'Invalid' are compared to other 'usable' routes for the same prefix according to the BGP decision process.

Default

no origin-invalid-unusable

Platforms

7705 SAR Gen 2

20.21 origin-validation

origin-validation

Syntax

origin-validation

Context

[Tree] (config>router origin-validation)

Full Context

configure router origin-validation

Description

Commands in this context display origin validation information.

Platforms

7705 SAR Gen 2

20.22 origin-validation-state

origin-validation-state

Syntax

origin-validation-state state

no origin-validation-state

Context

[Tree] (config>router>policy-options>policy-statement>entry>from origin-validation-state)

Full Context

configure router policy-options policy-statement entry from origin-validation-state

Description

This command is used to match BGP routes on the basis of origin validation state:

- Valid (0)
- Not-Found (1)
- Invalid (2)

Default

no origin-validation-state

Parameters

valid

Marks the route as having an origin validation state of valid.

notFound

Marks the route as having an origin validation state of Not Found.

invalid

Marks the route as having an origin validation state of invalid.

Platforms

7705 SAR Gen 2

origin-validation-state

Syntax

origin-validation-state {*state* | *param-name*}

no origin-validation-state

Context

[Tree] (config>router>policy-options>policy-statement>default-action origin-validation-state)

[Tree] (config>router>policy-options>policy-statement>entry>action origin-validation-state)

Full Context

configure router policy-options policy-statement default-action origin-validation-state

configure router policy-options policy-statement entry action origin-validation-state

Description

This command is used to mark BGP IPv4 and IPv6 routes matching the **default-action** or a specific entry of a route policy with one of the 3 following origin validation states:

- Valid (0)
- Not-Found (1)
- Invalid (2)

Default

no origin-validation-state

Parameters

state

Specifies the default operational origin validation state for this policy statement.

Values valid — Marks the route as having an origin validation state of valid.

notFound — Marks the route as having an origin validation state of Not Found.

invalid — Marks the route as having an origin validation state of invalid.

param-name

Specifies the origin parameter variable name. Allowed values are any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@”.

Platforms

7705 SAR Gen 2

20.23 originate-default-route

originate-default-route

Syntax

originate-default-route [**type-nssa**] [**adjacency-check**]

no originate-default-route

Context

[Tree] (config>service>vprn>ospf>area>nssa originate-default-route)

[Tree] (config>service>vprn>ospf3>area>nssa originate-default-route)

Full Context

configure service vprn ospf area nssa originate-default-route

configure service vprn ospf3 area nssa originate-default-route

Description

This command specifies whether when configuring an NSSA with no summaries, the Area Border Router (ABR) injects a type-7 LSA default route into the NSSA area. The default behavior is to inject a type-3 LSA default route, but some older implementations expect a type-7 LSA default route.

When configuring an NSSA with no summaries, the ABR will inject a type 3 LSA default route into the NSSA area. Some older implementations expect a type 7 LSA default route.

The **no** form of this command disables origination of a default route.

Default

no originate-default-route — A default route is not originated.

Parameters

type-nssa

Specifies that a type 7 LSA should be used for the default route.

Configure this parameter to inject a type 7 LSA default route into an NSSA configured with no summaries, instead of a type 3 LSA.

To revert to a type 3 LSA, execute the **originate-default-route** command without the **type-nssa** parameter.

Default type 3 LSA default route

adjacency-check

Specifies whether adjacency checks are performed before originating a default route. If this parameter is configured, then no area 0 adjacency is required for the ABR to advertise the default route.

Default Adjacency checks are performed, and an area 0 adjacency is required for the ABR to advertise the default route

Platforms

7705 SAR Gen 2

originate-default-route

Syntax

originate-default-route [**type-7**] [**no-adjacency-check**]

originate-default-route [**type-nssa**] [**no-adjacency-check**]

no originate-default-route

Context

[Tree] (config>router>ospf>area>nssa originate-default-route)

[Tree] (config>router>ospf3>area>nssa originate-default-route)

Full Context

configure router ospf area nssa originate-default-route

configure router ospf3 area nssa originate-default-route

Description

This command enables the generation of a default route and its LSA type (3 or 7) into a Not So Stubby Area (NSSA) by an NSSA Area Border Router (ABR).

When configuring an NSSA with no summaries, the ABR will inject a type 3 LSA default route into the NSSA area. Some older implementations expect a type 7 LSA default route.

The **no** form of this command disables origination of a default route.

Default

no originate-default-route

Parameters

type-7

Specifies a type 7 LSA should be used for the default route in the **config**>**router**>**ospf**>**area**>**nssa** context.

Configure this parameter to inject a type-7 LSA default route instead the type 3 LSA into the NSSA configured with no summaries.

To revert to a type 3 LSA, enter **originate-default-route** without the **type-7** parameter.

Default Type 3 LSA default route.

type-nssa

Specifies an NSSA-LSA type should be used for the default route in the **config**>**router**>**ospf3**>**area**>**nssa** context.

no-adjacency-check

Specifies whether adjacency checks are performed before originating a default route. If this parameter is configured, then no area 0 adjacency is required for the ABR to advertise the default route.

Default Adjacency checks are performed, and an area 0 adjacency is required for the ABR to advertise the default route.

Platforms

7705 SAR Gen 2

20.24 originated-qos-marking

originated-qos-marking

Syntax

[**no**] **originated-qos-marking** *dscp-name*

Context

[Tree] (config>system>telemetry>persistent-subscriptions>subscription originated-qos-marking)

Full Context

configure system telemetry persistent-subscriptions subscription originated-qos-marking

Description

This command configures the QoS marking used for packets carrying telemetry notifications.

The **no** form of this command removes the QoS marking.

Parameters

dscp-name

Specifies the QoS marking name.

The *dscp-name* parameter is a 6-bit value. It must be one of the predefined DSCP names in the system.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25,

af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, e f, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

Platforms

7705 SAR Gen 2

originated-qos-marking

Syntax

originated-qos-marking *dscp-name*

no originated-qos-marking

Context

[Tree] (config>system>grpc-tunnel>destination-group>destination originated-qos-marking)

Full Context

configure system grpc-tunnel destination-group destination originated-qos-marking

Description

This command configures the QoS marking used for packets carrying gRPC tunnel packets.

The **no** form of this command removes the QoS marking.

Default

no originated-qos-marking

Parameters

dscp-name

Specifies the QoS marking name.

The *dscp-name* parameter is a 6-bit value. It must be one of the predefined DSCP names in the system.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, e f, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

Platforms

7705 SAR Gen 2

20.25 orphan-override

orphan-override

Syntax

orphan-override [**level** *priority-level*] [**weight** *weight*] [**cir-level** *cir-level*] [**cir-weight** *cir-weight*]

no orphan-override

Context

[Tree] (config>qos>port-scheduler-policy orphan-override)

Full Context

configure qos port-scheduler-policy orphan-override

Description

This command overrides the default orphan behavior for port schedulers created using the port scheduler policy. The default orphan behavior is to give all orphan queues and schedulers bandwidth after all other properly parented queues and schedulers. Orphans by default do not receive any within-CIR bandwidth and receive above-CIR bandwidth after priority levels 8 through 1 have been allocated. The orphan-override command accepts the same parameters as the port-parent command in the SAP egress and network queue policy contexts. The defined parameters are used as a default port-parent association for any queue or scheduler on the port that the port scheduler policy is applied.

Orphan queues and schedulers are identified as:

- Any queue or scheduler that does not have a port-parent or parent command applied
- Any queue that has a parent command applied, but the specified scheduler name does not exist on the queue's SAP, MSS, or SLA Profile instance.

A queue or scheduler may be properly parented to an upper level scheduler, but that scheduler may be orphaned. In this case, the queue or scheduler receives bandwidth from its parent scheduler based on the parent schedulers ability to receive bandwidth as an orphan.

Within-CIR Priority Level Parameters

The within-CIR parameters define which port priority level the orphan queues and schedulers should be associated with when receiving bandwidth for the queue or schedulers within-CIR offered load. The within-CIR offered load is the amount of bandwidth the queue or schedulers could use that is equal to or less than its defined or summed CIR value. The summed value is only valid on schedulers and is the sum of the within-CIR offered loads of the children attached to the scheduler. The parameters that control within-CIR bandwidth allocation for orphans are the orphan-override commands cir-level and cir-weight keywords. The cir-level keyword defines the port priority level that the scheduler or queue uses to receive bandwidth for its within-CIR offered load. The cir-weight is used when multiple queues or schedulers exist at the same port priority level for within-CIR bandwidth. The weight value defines the relative ratio that is used to distribute bandwidth at the priority level when more within-CIR offered load exists than the port priority level has bandwidth.

A cir-weight equal to zero (the default value) has special meaning and informs the system that the orphan queues and schedulers do not receive bandwidth from the within-CIR distribution. Instead, all bandwidth for the orphan queues and schedulers must be allocated from the port scheduler's above-CIR pass.

Above-CIR Priority Level Parameters

The above-CIR parameters define which port priority level the orphan queues and schedulers should be associated with when receiving bandwidth for the queue or schedulers above-CIR offered load. The above-CIR offered load is the amount of bandwidth the queue or schedulers could use that is equal to or less than its defined PIR value (based on the queue or schedulers rate command) less any bandwidth that was given to the queue or scheduler during the above-CIR scheduler pass. The parameters that control above-CIR bandwidth allocation for orphans are the orphan-override commands level and weight keywords. The **level** keyword defines the port priority level that the scheduler or queue uses to receive bandwidth for its above-CIR offered load. The weight is used when multiple queues or schedulers exist at the same port priority level for above-CIR bandwidth. The weight value defines the relative ratio that is used to distribute bandwidth at the priority level when more above-CIR offered load exists than the port priority level has bandwidth.

The **no** form of this command removes the orphan override port parent association for the orphan queues and schedulers on port schedulers created with the port scheduler policy. Any orphan queues and schedulers on a port associated with the port scheduler policy will revert to default orphan behavior.

Parameters

level *priority-level*

Defines the port priority the orphan queues and schedulers will use to receive bandwidth for their above-CIR offered-load.

Values 1 to 8 (8 is the highest priority)

Default 1

weight *weight*

Defines the weight the orphan queues and schedulers will use in the above-CIR port priority level (defined by the level parameter).

Values 1 to 100

Default 1

cir-level *cir-level*

Defines the port priority the orphan queues and schedulers will use to receive bandwidth for their within-CIR offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the orphan queues and schedulers do not receive bandwidth during the port scheduler's within-CIR pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 1 to 8 (8 is the highest level)

cir-weight *cir-weight*

Defines the weight the orphan queues and schedulers will use in the within-CIR port priority level (defined by the cir-level parameter). When the cir-weight parameter is set to a value of 0 (the default value), the orphan queues and schedulers do not receive bandwidth

during the port scheduler's within-CIR pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 1 to 100 (100 is the highest weight)

Platforms

7705 SAR Gen 2

20.26 ospf

ospf

Syntax

ospf [*router-id*]

no ospf

Context

[Tree] (config>service>vprn ospf)

Full Context

configure service vprn ospf

Description

This command enables access to the context to enable an OSPF protocol instance.

OSPF instances are **shutdown** when created, so that all parameters can be configured prior to the instance being enabled.

The **no** form of this command deletes the OSPF protocol instance removing all associated configuration parameters.

Default

no ospf

Parameters

router-id

Specifies the OSPF router ID to be used with the associated OSPF instance. The *router-id* must be given a dot decimal notation format.

Values a.b.c.d

Platforms

7705 SAR Gen 2

ospf

Syntax

ospf *ospf-instance* [*router-id*]

[**no**] **ospf** *ospf-instance*

Context

[Tree] (config>router ospf)

Full Context

configure router ospf

Description

This command creates an OSPF routing instance and then enters the associated context to configure the associated protocol parameters.

Additionally, the router ID can be specified as another parameter of the OSPF command. This parameter is required for all non-base OSPF instances.

The default value for the base instance is inherited from the configuration in the **config**>**router** context. When that is not configured, the following apply:

1. the system uses the system interface address (which is also the loopback address)
2. if a system interface address is not configured, it uses the last 32 bits of the chassis MAC address

This is a required command when configuring multiple instances and the instance being configured is not the base instance. When configuring multiple instances of OSPF, there is a risk of loops because networks are advertised by multiple domains configured with multiple interconnections to one another. To prevent this from happening, all routers in a domain should be configured with the same domain ID. Each domain (OSPF-instance) should be assigned a specific bit value in the 32-bit tag mask.

The default value for non-base instances is 0.0.0.0 and is invalid; in this case, the instance of OSPF will not start. When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.

Issue the shutdown and no shutdown commands for the instance for the new router ID to be used, or reboot the entire router.

OSPF instances are **shutdown** when created, so that all parameters can be configured prior to the instance being enabled.

The **no** form of this command to reverts to the default value.

Default

no ospf

Parameters

ospf-instance

Specifies a unique integer that identifies a specific instance of a version of the OSPF protocol running in the router instance specified by the router ID.

Values 1 to 31

router-id

Specifies the OSPF router ID to be used with the associated OSPF instance. This IP address must be given a dot decimal notation format.

Platforms

7705 SAR Gen 2

ospf

Syntax

ospf [*ospf-instance*]

no ospf [*ospf-instance*]

Context

[Tree] (debug>router ospf)

Full Context

debug router ospf

Description

Indicates the OSPF instance for debugging purposes.

Parameters

ospf-instance

Debugs the specified OSPF instance.

Values 0 to 31

Platforms

7705 SAR Gen 2

20.27 ospf-dynamic-hostnames

ospf-dynamic-hostnames

Syntax

[**no**] **ospf-dynamic-hostnames**

Context

[Tree] (config>system ospf-dynamic-hostnames)

Full Context

configure system ospf-dynamic-hostnames

Description

This command enables OSPF dynamic hostnames.

The router receiving the new Dynamic Hostname within the OSPF Router Information (RI) LSA is instructed to process the received dynamic hostname information.

The **no** form of this command disables OSPF dynamic hostnames.

Default

no ospf-dynamic-hostnames

Platforms

7705 SAR Gen 2

20.28 ospf3

ospf3

Syntax

ospf3 [*instance-id*] [*router-id*]

[**no**] **ospf3** *instance-id*

Context

[Tree] (config>service>vprn ospf3)

Full Context

configure service vprn ospf3

Description

This command creates an OSPFv3 routing instance and then enters the associated context to configure associated protocol parameters.

OSPF instances are **shutdown** when created, so that all parameters can be configured before the instance is enabled.

The **no** form of this command deletes the OSPFv3 protocol instance, removing all associated configuration parameters.

Default

no ospf3

Parameters

instance-id

Specifies the instance ID for the OSPFv3 instance being created or modified. The instance ID must match the specified range based on the address family.

Values 0 to 31: IPv6 unicast

64 to 95: IPv4 unicast

router-id

Specifies the IP address.

Platforms

7705 SAR Gen 2

ospf3

Syntax

ospf3 [*ospf-instance*] [*router-id*]

[**no**] **ospf3** *instance-id*

Context

[Tree] (config>router ospf3)

Full Context

configure router ospf3

Description

This command creates an OSPFv3 routing instance and then enters the associated context to configure associated protocol parameters.

OSPFv3 instances are **shutdown** when created, so that all parameters can be configured prior to the instance being enabled.

The **no** form of this command deletes the OSPFv3 protocol instance, removing all associated configuration parameters.

Parameters

ospf-instance

Specifies the instance ID for the OSPFv3 instance being created or modified. The instance ID must match the specified range based on the address family.

Values 0 to 31: IPV6 unicast

64 to 95: IPV4 unicast

router-id

Specifies the OSPF router ID to be used with the associated OSPF instance. This IP address must be given a dot decimal notation format.

Platforms

7705 SAR Gen 2

ospf3

Syntax

ospf3 [*ospf-instance*]

no ospf3 [*ospf-instance*]

Context

[Tree] (debug>router ospf3)

Full Context

debug router ospf3

Description

Indicates the OSPF3 instance for debugging purposes.

Parameters

ospf-instance

Debugs the specified OSPF3 instance.

Values 0 to 31 | 64 to 95

0 to 31 — IPv6-unicast address-family

64 to 95 — IPv4-unicast address-family

Platforms

7705 SAR Gen 2

20.29 other-stateful-configuration

other-stateful-configuration

Syntax

[**no**] **other-stateful-configuration**

Context

[Tree] (config>service>vprn>router-advert>if other-stateful-configuration)

Full Context

configure service vprn router-advertisement interface other-stateful-configuration

Description

This command sets the "other configuration" flag. This flag indicates that DHCPv6 is available for auto-configuration of other (non-address) information such as DNS-related information or information on other servers in the network. See RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) for IPv6*.

The **no** form of this command removes the flag.

Default

no other-stateful-configuration

Platforms

7705 SAR Gen 2

other-stateful-configuration

Syntax

[**no**] **other-stateful-configuration**

Context

[Tree] (config>router>router-advert>if other-stateful-configuration)

Full Context

configure router router-advertisement interface other-stateful-configuration

Description

This command sets the "Other configuration" flag. This flag indicates that DHCPv6lite is available for autoconfiguration of other (non-address) information such as DNS-related information or information about other servers in the network. See RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) for IPv6*.

Default

no other-stateful-configuration

Platforms

7705 SAR Gen 2

20.30 out-profile-octets-discarded-count

out-profile-octets-discarded-count

Syntax

[**no**] **out-profile-octets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters out-profile-octets-discarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>e-counters out-profile-octets-discarded-count)

[Tree] (config>log>acct-policy>cr>queue>e-counters out-profile-octets-discarded-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters out-profile-octets-discarded-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters out-profile-octets-discarded-count

configure log accounting-policy custom-record ref-queue e-counters out-profile-octets-discarded-count

configure log accounting-policy custom-record queue e-counters out-profile-octets-discarded-count

configure log accounting-policy custom-record policer e-counters out-profile-octets-discarded-count

Description

This command includes the out of profile packets discarded count.

The **no** form of this command excludes the out of profile packets discarded count.

Default

no out-profile-octets-discarded-count

Platforms

7705 SAR Gen 2

out-profile-octets-discarded-count

Syntax

[**no**] **out-profile-octets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters out-profile-octets-discarded-count)

[Tree] (config>log>acct-policy>cr>policer>i-counters out-profile-octets-discarded-count)

Full Context

configure log accounting-policy custom-record ref-policer i-counters out-profile-octets-discarded-count

configure log accounting-policy custom-record policer i-counters out-profile-octets-discarded-count

Description

This command includes the out of profile octets discarded count.

The **no** form of this command excludes the out of profile octets discarded count.

Default

no out-profile-octets-discarded-count

Platforms

7705 SAR Gen 2

20.31 out-profile-octets-forwarded-count

out-profile-octets-forwarded-count

Syntax

[**no**] **out-profile-octets-forwarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters out-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters out-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>e-counters out-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>queue>e-counters out-profile-octets-forwarded-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters out-profile-octets-forwarded-count

configure log accounting-policy custom-record policer e-counters out-profile-octets-forwarded-count

configure log accounting-policy custom-record ref-queue e-counters out-profile-octets-forwarded-count

configure log accounting-policy custom-record queue e-counters out-profile-octets-forwarded-count

Description

This command includes the out of profile octets forwarded count.

The **no** form of this command excludes the out of profile octets forwarded count.

Default

no out-profile-octets-forwarded-count

Platforms

7705 SAR Gen 2

out-profile-octets-forwarded-count

Syntax

[**no**] **out-profile-octets-forwarded-count**

Context

[Tree] (config>log>acct-policy>cr>queue>i-counters out-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters out-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>policer>i-counters out-profile-octets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters out-profile-octets-forwarded-count)

Full Context

configure log accounting-policy custom-record queue i-counters out-profile-octets-forwarded-count

configure log accounting-policy custom-record ref-policer i-counters out-profile-octets-forwarded-count

configure log accounting-policy custom-record policer i-counters out-profile-octets-forwarded-count

configure log accounting-policy custom-record ref-queue i-counters out-profile-octets-forwarded-count

Description

This command includes the out of profile octets forwarded count.

The **no** form of this command excludes the out of profile octets forwarded count.

Default

no out-profile-octets-forwarded-count

Platforms

7705 SAR Gen 2

20.32 out-profile-octets-offered-count

out-profile-octets-offered-count

Syntax

[**no**] **out-profile-octets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>policer>e-counters out-profile-octets-offered-count)

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters out-profile-octets-offered-count)

Full Context

configure log accounting-policy custom-record policer e-counters out-profile-octets-offered-count

configure log accounting-policy custom-record ref-policer e-counters out-profile-octets-offered-count

Description

This command includes the out of profile octets offered count.

The **no** form of this command excludes the out of profile octets offered count.

Default

no out-profile-octets-offered-count

Platforms

7705 SAR Gen 2

out-profile-octets-offered-count

Syntax

[**no**] **out-profile-octets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters out-profile-octets-offered-count)

[Tree] (config>log>acct-policy>cr>policer>i-counters out-profile-octets-offered-count)

Full Context

configure log accounting-policy custom-record ref-policer i-counters out-profile-octets-offered-count

configure log accounting-policy custom-record policer i-counters out-profile-octets-offered-count

Description

This command includes the out of profile octets offered count.

The **no** form of this command excludes the out of profile octets offered count.

Default

no out-profile-octets-offered-count

Platforms

7705 SAR Gen 2

20.33 out-profile-packets-discarded-count

out-profile-packets-discarded-count

Syntax

[**no**] **out-profile-packets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>queue>e-counters out-profile-packets-discarded-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters out-profile-packets-discarded-count)

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters out-profile-packets-discarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>e-counters out-profile-packets-discarded-count)

Full Context

configure log accounting-policy custom-record queue e-counters out-profile-packets-discarded-count

configure log accounting-policy custom-record policer e-counters out-profile-packets-discarded-count

configure log accounting-policy custom-record ref-policer e-counters out-profile-packets-discarded-count

configure log accounting-policy custom-record ref-queue e-counters out-profile-packets-discarded-count

Description

This command includes the out of profile packets discarded count.

The **no** form of this command excludes the out of profile packets discarded count.

Default

no out-profile-packets-discarded-count

Platforms

7705 SAR Gen 2

out-profile-packets-discarded-count

Syntax

[**no**] **out-profile-packets-discarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters out-profile-packets-discarded-count)

[Tree] (config>log>acct-policy>cr>policer>i-counters out-profile-packets-discarded-count)

Full Context

configure log accounting-policy custom-record ref-policer i-counters out-profile-packets-discarded-count

configure log accounting-policy custom-record policer i-counters out-profile-packets-discarded-count

Description

This command includes the out of profile packets discarded count.

The **no** form of this command excludes the out of profile packets discarded count.

Default

no out-profile-packets-discarded-count

Platforms

7705 SAR Gen 2

20.34 out-profile-packets-forwarded-count

out-profile-packets-forwarded-count

Syntax

[**no**] **out-profile-packets-forwarded-count**

Context

[Tree] (config>log>acct-policy>cr>ref-queue>e-counters out-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters out-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters out-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>queue>e-counters out-profile-packets-forwarded-count)

Full Context

configure log accounting-policy custom-record ref-queue e-counters out-profile-packets-forwarded-count

configure log accounting-policy custom-record policer e-counters out-profile-packets-forwarded-count

configure log accounting-policy custom-record ref-policer e-counters out-profile-packets-forwarded-count

configure log accounting-policy custom-record queue e-counters out-profile-packets-forwarded-count

Description

This command includes the out of profile packets forwarded count.

The **no** form of this command excludes the out of profile packets forwarded count.

Platforms

7705 SAR Gen 2

out-profile-packets-forwarded-count

Syntax

[**no**] **out-profile-packets-forwarded-count**

Context

[Tree] (config>log>acct-policy>cr>queue>i-counters out-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters out-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>policer>i-counters out-profile-packets-forwarded-count)

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters out-profile-packets-forwarded-count)

Full Context

configure log accounting-policy custom-record queue i-counters out-profile-packets-forwarded-count

configure log accounting-policy custom-record ref-policer i-counters out-profile-packets-forwarded-count

configure log accounting-policy custom-record policer i-counters out-profile-packets-forwarded-count

configure log accounting-policy custom-record ref-queue i-counters out-profile-packets-forwarded-count

Description

This command includes the out of profile packets forwarded count.

The **no** form of this command excludes the out of profile packets forwarded count.

Default

no out-profile-packets-forwarded-count

Platforms

7705 SAR Gen 2

20.35 out-profile-packets-offered-count

out-profile-packets-offered-count

Syntax

[**no**] **out-profile-packets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>policer>e-counters out-profile-packets-offered-count)

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters out-profile-packets-offered-count)

Full Context

configure log accounting-policy custom-record policer e-counters out-profile-packets-offered-count

configure log accounting-policy custom-record ref-policer e-counters out-profile-packets-offered-count

Description

This command includes the out of profile packets offered count.

The **no** form of this command excludes the out of profile packets offered count.

Default

no out-profile-packets-offered-count

Platforms

7705 SAR Gen 2

out-profile-packets-offered-count

Syntax

[**no**] **out-profile-packets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>policer>i-counters out-profile-packets-offered-count)

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters out-profile-packets-offered-count)

Full Context

configure log accounting-policy custom-record policer i-counters out-profile-packets-offered-count

configure log accounting-policy custom-record ref-policer i-counters out-profile-packets-offered-count

Description

This command includes the out of profile packets offered count.

The **no** form of this command excludes the out of profile packets offered count.

Default

no out-profile-packets-offered-count

Platforms

7705 SAR Gen 2

20.36 out-remark

out-remark

Syntax

out-remark {**dscp** *dscp-name* | **prec** *ip-prec-value*}

no out-remark

Context

[Tree] (config>qos>sap-ingress>fc out-remark)

Full Context

configure qos sap-ingress fc out-remark

Description

This command is used in a SAP ingress QoS policy to define an explicit out-of-profile remark action for a forwarding class or subclass. While the SAP ingress QoS policy may be applied to any SAP, the remarking functions are only enforced when the SAP is associated with an IP or subscriber interface (in an IES or VPRN). When the policy is applied to a Layer 2 SAP (for example, Epipe or VPLS), the remarking definitions are silently ignored.

In the case where the policy is applied to a Layer 3 SAP, the out-of-profile remarking definition will be applied to packets that have been classified to the forwarding class or subclass. It is possible for a packet to match a classification command that maps the packet to a particular forwarding class or subclass, only to have a more explicit (higher priority match) override the association. Only the highest priority match forwarding class or subclass association will drive the out-of-profile marking.

The out-remark command is only applicable to ingress IP routed packets that are considered out-of-profile. The profile of a SAP ingress packet is affected by either the explicit in-profile/out-of-profile definitions or the ingress policing function applied to the packet. Table 73: Out-remark Command Effect describes the effect of the out-remark command on received SAP ingress packets. Within the out-of-profile IP packet's ToS field, either the six DSCP bits or the three precedence bits are remarked.

Table 73: Out-remark Command Effect

SAP Ingress Packet State	out-remark Command Effect
Non-Routed, Policed In-Profile	No Effect (non-routed packet)
Non-Routed, Policed Out-of-Profile	No Effect (non-routed packet)
Non-Routed, Explicit In-Profile	No Effect (non-routed packet)
Non-Routed, Explicit Out-of-Profile	No Effect (non-routed packet)
IP Routed, Policed In-Profile	No Effect (in-profile packet)
IP Routed, Policed Out-of-Profile	out-remark value applied to IP header ToS field
IP Routed, Explicit In-Profile	No Effect (in-of-profile packet)
IP Routed, Explicit Out-of-Profile	out-remark value applied to IP header ToS field

A packet that is explicitly remarked at ingress will not be affected by any egress remarking decision. Explicit ingress remarking has highest priority.

An explicit dscp name or precedence value must be specified for out-of-profile remarking to be applied.

The **no** form of this command disables ingress remarking of out-of-profile packets classified to the forwarding class or subclass.

Default

no out-remark

Parameters

dscp *dscp-name*

Specifies that the matching packet's DSCP bits should be overridden with the value represented by *dscp-name*.

The *dscp-name* parameter is a 6-bit value. It must be one of the predefined DSCP names defined on the system.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, e f, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

prec *ip-prec-value*

Specifies that the matching packet's precedence bits should be overridden with the value represented by *ip-prec-value*.

The value specified by *ip-prec-value* is used to overwrite the precedence bits within a matching routed packets IP header ToS field.

Values 0 to 7

Platforms

7705 SAR Gen 2

20.37 outband

outband

Syntax

outband *service-id*

no outband

Context

[Tree] (config>system>security>vprn-aaa-server outband)

Full Context

configure system security vprn-aaa-server outband

Description

This command configures TACACS+ and RADIUS servers in a VPRN to be used for AAA by that VPRN and by sessions on the console or out-of-band (OOB) Ethernet ports.

The **no** form of this command disables the use of servers in out-of-band management.

Default

no outband

Parameters

service-id

Specifies the VPRN server for AAA to use for OOB sessions.

Values *service-id*: 1 to 2147483648

svc-name: 64 characters maximum

Platforms

7705 SAR Gen 2

20.38 outbound-max-sessions

outbound-max-sessions

Syntax

outbound-max-sessions *number-of-sessions*

no outbound-max-sessions

Context

[Tree] (config>system>login-control>ssh outbound-max-sessions)

[Tree] (config>system>login-control>telnet outbound-max-sessions)

Full Context

configure system login-control ssh outbound-max-sessions

configure system login-control telnet outbound-max-sessions

Description

This parameter limits the number of outbound Telnet and SSH sessions. A maximum of 15 Telnet and SSH connections can be established from the router. The local serial port cannot be disabled.

The **no** form of this command reverts to the default value.

Default

outbound-max-sessions 5

Parameters

value

Specifies the maximum number of concurrent outbound Telnet sessions, expressed as an integer.

Values 0 to 15

Platforms

7705 SAR Gen 2

20.39 outbound-route-filtering

outbound-route-filtering

Syntax

[**no**] **outbound-route-filtering**

Context

[Tree] (config>router>bgp>group>neighbor outbound-route-filtering)

[Tree] (config>router>bgp outbound-route-filtering)

[Tree] (config>router>bgp>group outbound-route-filtering)

Full Context

configure router bgp group neighbor outbound-route-filtering

configure router bgp outbound-route-filtering

configure router bgp group outbound-route-filtering

Description

This command opens the configuration tree for sending or accepting BGP filter lists from peers (outbound route filtering).

Default

no outbound-route-filtering

Platforms

7705 SAR Gen 2

outbound-route-filtering

Syntax

[**no**] **outbound-route-filtering**

Context

[Tree] (debug>router>bgp outbound-route-filtering)

Full Context

debug router bgp outbound-route-filtering

Description

This command enables debugging for all BGP outbound route filtering (ORF) packets. ORF is used to inform a neighbor of targets (using target-list) that it is willing to receive.

Platforms

7705 SAR Gen 2

20.40 outer-tag

outer-tag

Syntax

outer-tag *value* [*vid-mask*]

no outer-tag

Context

[Tree] (config>qos>sap-ingress>mac-criteria>entry>match outer-tag)

Full Context

configure qos sap-ingress mac-criteria entry match outer-tag

Description

This command configures the matching of the first tag that is carried transparently through the service. Service delimiting tags are stripped from the frame and the outer tag on ingress is the first tag after any service delimiting tags. The outer tag is the first tag before any service delimiting tags on egress. This allows matching VLAN tags for explicit filtering or QoS setting when using default or null encapsulations.

On dot1Q SAPs, the outer tag is the only tag that can be matched. On dot1Q SAPs with exact match (sap 2/1/1:50), the outer tag will be populated with the next tag that is carried transparently through the service or 0 if there is no additional VLAN tags on the frame.

On QinQ SAPs that strip a single service delimiting tag, the outer tag will contain the next tag (which is still the first tag carried transparently through the service.) On SAPs with two service delimiting tags (two tags stripped), the **outer-tag** will contain 0 even if there are more than two tags on the frame.

The optional *vid_mask* is defaulted to 4095 (exact match) but may be specified to allow pattern matching. The masking operation is ((value & vid-mask) = = (tag & vid-mask)). A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6.

For QoS, the VID type cannot be specified on the default QoS policy.

The default vid-mask is set to 4095 for exact match.

Platforms

7705 SAR Gen 2

20.41 output-authorization

output-authorization

Syntax

output-authorization

Context

[Tree] (config>system>security>management-interface output-authorization)

Full Context

configure system security management-interface output-authorization

Description

This command configures the authorization of the configuration and state output in model-driven interfaces and telemetry. When enabled, commands that display configuration or state output authorize every element in the output. If a remote AAA server is configured, this can cause delays in displaying output while it is authorized. If a large amount of output is displayed, for example, when displaying the system configuration, the remote AAA server receives a large number of authorization requests.

Input to edit the configuration is not affected by this command, and is always authorized.

Platforms

7705 SAR Gen 2

20.42 outside

outside

Syntax

outside

Context

[Tree] (config>service>vprn>nat outside)

[Tree] (config>router>nat outside)

Full Context

configure service vprn nat outside

configure router nat outside

Description

Commands in this context configure the outside NAT instance.

Platforms

7705 SAR Gen 2

20.43 outside-range

outside-range

Syntax

outside-range *outside-ip-address*

no outside-range

Context

[Tree] (config>service>vprn>nat>inside>deterministic>address-map outside-range)

[Tree] (config>router>nat>inside>deterministic>address-map outside-range)

Full Context

configure service vprn nat inside deterministic address-map outside-range

configure router nat inside deterministic address-map outside-range

Description

This command configures the mapping of the inside address range of deterministic NAT subscribers to the first IP address in the outside IP address range in the NAT pool.

The last outside IP address is determined by the number of subscribers mapped to an outside IP address via the **configure router nat outside pool subscriber-limit** and **configure service vprn nat outside pool subscriber-limit** commands.

The **no** form of this command removes the configuration.

Default

no outside-range

Parameters

outside-ip-address

Specifies an outside IPv4 address.

Values a.b.c.d

Platforms

7705 SAR Gen 2

20.44 overlapping-reverse-route

overlapping-reverse-route

Syntax

[**no**] **overlapping-reverse-route**

Context

[Tree] (config>service>vprn>ipsec overlapping-reverse-route)

Full Context

configure service vprn ipsec overlapping-reverse-route

Description

This command configures the router to accept overlapping DL2L tunnel reverse routes from different tunnels and install the routes based on the preference, metric, or ECMP configuration.

This command is mutually exclusive with the **configure service vprn ipsec allow-reverse-route-override-type** command.

The **no** form of this command disables the acceptance of overlapping reverse routes. The router handles the overlapping route according to the **allow-reverse-route-override-type** command configuration. See the *7705 SAR Gen 2 Multiservice ISA and ESA Guide* for more information.

Default

no overlapping-reverse-route

Platforms

7705 SAR Gen 2

20.45 overload

overload

Syntax

overload [**timeout** *seconds*] [**max-metric**]

no overload

Context

[Tree] (config>service>vprn>isis overload)

Full Context

configure service vprn isis overload

Description

This command administratively sets the IS-IS router to operate in the overload state for a specific time period, in seconds, or indefinitely.

During normal operation, the router may be forced to enter an overload state due to a lack of resources. When in the overload state, the router is only used if the destination is reachable by the router and will not be used for other transit traffic.

If a time period is specified, the overload state persists for the configured length of time. If no time is specified, the overload state operation is maintained indefinitely.

The **overload** command can be useful in circumstances where the router is overloaded or used prior to executing a **shutdown** command to divert traffic around the router.

The **max-metric** parameter can be set to advertise transit links with the maximum metric of 0xffffffe (wide metrics) or 0x3f (regular metrics), instead of setting the overload bit when placing the router in overload.

The **no** form of this command causes the router to exit the overload state.

Default

no overload

Parameters

seconds

Specifies the time, in seconds, that this router must operate in overload state.

Values 60 to 1800

Default infinity (overload state maintained indefinitely)

max-metric

Set the maximum metric instead of overload.

Platforms

7705 SAR Gen 2

overload

Syntax

overload [**timeout** *seconds*]

no overload

Context

[Tree] (config>service>vprn>ospf overload)

[Tree] (config>service>vprn>ospf3 overload)

Full Context

configure service vprn ospf overload

configure service vprn ospf3 overload

Description

This command changes the overload state of the local router so that it appears to be overloaded. When overload is enabled, the router can participate in OSPF routing, but is not used for transit traffic. Traffic destined for directly attached interfaces continue to reach the router.

To put the IGP in an overload state, enter a **timeout** value. The IGP enters the overload state until the **timeout** timer expires or a **no overload** command is executed.

If the **overload** command is performed during the execution of an **overload-on-boot** command, the **overload** command takes precedence. This could occur as a result of a saved configuration file in which both parameters are saved. When the file is saved by the system, the **overload-on-boot** command is saved after the **overload** command.

Use the **no** form of this command to return to the default. When the **no overload** command is executed, the overload state is terminated, regardless the reason the protocol entered overload state.

Default

no overload

Parameters

timeout seconds

Specifies the number of seconds to reset overloading.

Values 60 to1800

Default 60

Platforms

7705 SAR Gen 2

overload

Syntax

overload [**timeout** *seconds*] [**max-metric**]

no overload

Context

[Tree] (config>router>isis overload)

Full Context

configure router isis overload

Description

This command administratively sets the IS-IS router to operate in the overload state for a specific time period, in seconds, or indefinitely.

During normal operation, the router may be forced to enter an overload state because of a lack of resources. When in the overload state, the router is only used if the destination is reachable by the router and is not used for other transit traffic.

If a time period is specified, the overload state persists for the configured length of time. If no time is specified, the overload state operation is maintained indefinitely.

The overload command is cleared from the configuration after a reboot if **overload-on-boot** is configured with or without a timeout value. To keep the IS-IS router in the overload state indefinitely after rebooting, configure **overload-on-boot** with no timeout value or configure the **overload** command with **no overload-on-boot** command.

The **overload** command can be useful in circumstances where the router is overloaded or used before executing a **shutdown** command to divert traffic around the router.

The **max-metric** parameter can be set to advertise transit links with the maximum metric of 0xfffffffe (wide metrics) or 0x3f (regular metrics), instead of setting the overload bit when placing the router in overload.

The **no** form of this command causes the router to exit the overload state.

Default

no overload

Parameters

seconds

Specifies the time, in seconds, that this router must operate in overload state.

Default infinity (overload state maintained indefinitely)

Values 60 to 1800

max-metric

Sets the maximum metric instead of overload.

Platforms

7705 SAR Gen 2

overload

Syntax

overload [**timeout** *seconds*]

no overload

Context

[Tree] (config>router>ospf3 overload)

[Tree] (config>router>ospf overload)

Full Context

configure router ospf3 overload

configure router ospf overload

Description

This command changes the overload state of the local router so that it appears to be overloaded. When overload is enabled, the router can participate in OSPF routing, but is not used for transit traffic. Traffic destined to directly attached interfaces continues to reach the router.

To put the IGP in an overload state enter a timeout value. The IGP will enter the overload state until the timeout timer expires or a **no overload** command is executed.

If the **overload** command is encountered during the execution of an **overload-on-boot** command then this command takes precedence. This could occur as a result of a saved configuration file where both parameters are saved. When the file is saved by the system the **overload-on-boot** command is saved after the **overload** command. **However**, when **overload-on-boot** is configured under OSPF with no timeout value configured, the router will remain in overload state indefinitely after a reboot.

The **no** form of this command reverts to the default. When the **no overload** command is executed, the overload state is terminated regardless of the reason the protocol entered overload state.

Default

no overload

Parameters

timeout seconds

Specifies the number of seconds to reset overloading.

Values 1 to 1800 in the following context.

```
configure router ospf
```

60 to 1800 in the following context.

```
configure router ospf3
```

Platforms

7705 SAR Gen 2

20.46 overload-export-external

overload-export-external

Syntax

[**no**] **overload-export-external**

Context

[Tree] (config>service>vprn>isis overload-export-external)

Full Context

configure service vprn isis overload-export-external

Description

This command enables external routes that are exported with an IS-IS export policy to continue to be advertised when the router is in overload.

The **no** form of this command causes external routes to be withdrawn when the router is in overload.

Default

no overload-export-external

Platforms

7705 SAR Gen 2

overload-export-external

Syntax

[**no**] **overload-export-external**

Context

[Tree] (config>router>isis overload-export-external)

Full Context

configure router isis overload-export-external

Description

This command enables external routes that are exported with an IS-IS export policy to continue to be advertised when the router is in overload.

The **no** form of this command causes external routes to be withdrawn when the router is in overload.

Default

no overload-export-external

Platforms

7705 SAR Gen 2

20.47 overload-export-interlevel

overload-export-interlevel

Syntax

[**no**] **overload-export-interlevel**

Context

[Tree] (config>service>vprn>isis overload-export-interlevel)

Full Context

configure service vprn isis overload-export-interlevel

Description

This command enables inter-level routes that are exported with an IS-IS export policy to continue to be advertised when the router is in overload.

The **no** form of this command causes inter-level routes to be withdrawn when the router is in overload.

Default

no overload-export-interlevel

Platforms

7705 SAR Gen 2

overload-export-interlevel

Syntax

[**no**] **overload-export-interlevel**

Context

[Tree] (config>router>isis overload-export-interlevel)

Full Context

configure router isis overload-export-interlevel

Description

This command enables inter-level routes that are exported with an IS-IS export policy to continue to be advertised when the router is in overload.

The **no** form of this command causes inter-level routes to be withdrawn when the router is in overload.

Default

no overload-export-interlevel

Platforms

7705 SAR Gen 2

20.48 overload-fib-error-notify-only

overload-fib-error-notify-only

Syntax

overload-fib-error-notify-only [**retry** *seconds*]

no overload-fib-error-notify-only

Context

[Tree] (config>router>isis overload-fib-error-notify-only)

[Tree] (config>service>vprn>isis overload-fib-error-notify-only)

Full Context

configure router isis overload-fib-error-notify-only

configure service vprn isis overload-fib-error-notify-only

Description

This command configures the IS-IS router to send a notification when an overload condition occurs when programming the FIB, instead of advertising the overload condition of the router in the IS-IS LSP.

Note: Nokia recommends being careful using this command. When you configure the router not to advertise the IS-IS overload state in the IS-IS LSP, other routers are not instructed to take the overloaded router out of the IS-IS forwarding topology and this will cause suboptimal forwarding and non-deterministic behavior on the overloaded router. To avoid changing the default IS-IS overflow behavior, leave this command disabled.

When this command is configured, the IS-IS router enters a suboptimal state where it only sends a notification trap; the router can still be used by transit traffic in this state. The IS-IS router tracks the segment routing prefix SIDs where FIB programming failed. With the **retry** parameter configured, the router retries programming the segment routing prefix SIDs in the FIB using this tracked information.

When this command is not configured, during normal operation, the system may force the router to enter an overload state because of a lack of FIB resources. In this state, the router is used to terminate traffic and is not used to transit traffic.

The removal of the **overload-fib-error-notify-only** command configuration causes the system to program the failed entries in the FIB by triggering an immediate SPF.

The **no** form of this command causes the router to enter the full overload state.

Default

no overload-fib-error-notify-only

Parameters

seconds

Specifies the time, in seconds, this router uses to retry programming the failed entries in the FIB when **overload-fib-error-notify-only** is configured. The **overload-fib-error-notify-only** command must be configured to use the retry timer.

Values 10 to 1800

Default 10

Platforms

7705 SAR Gen 2

20.49 overload-include-ext-1

overload-include-ext-1

Syntax

[**no**] **overload-include-ext-1**

Context

[Tree] (config>service>vprn>ospf overload-include-ext-1)

[Tree] (config>service>vprn>ospf3 overload-include-ext-1)

Full Context

configure service vprn ospf overload-include-ext-1

configure service vprn ospf3 overload-include-ext-1

Description

This command controls whether routes should be readvertised with a maximum metric value when the system goes into overload state for any reason. When this command is enabled and the router is in overload, all external type-1 routes are advertised with the maximum metric.

The **no** form of this command reverts to the default value.

Default

no overload-include-ext-1

Platforms

7705 SAR Gen 2

overload-include-ext-1

Syntax

[**no**] **overload-include-ext-1**

Context

[Tree] (config>router>ospf3 overload-include-ext-1)

[Tree] (config>router>ospf overload-include-ext-1)

Full Context

configure router ospf3 overload-include-ext-1

configure router ospf overload-include-ext-1

Description

This command controls whether external type-1 routes should be readvertised with a maximum metric value when the system goes into overload state for any reason. When this command is enabled and the router is in overload, all external type-1 routes are advertised with the maximum metric.

The **no** form of this command reverts to the default value.

Default

no overload-include-ext-1

Platforms

7705 SAR Gen 2

20.50 overload-include-ext-2

overload-include-ext-2

Syntax

[**no**] **overload-include-ext-2**

Context

[Tree] (config>service>vprn>ospf overload-include-ext-2)

[Tree] (config>service>vprn>ospf3 overload-include-ext-2)

Full Context

configure service vprn ospf overload-include-ext-2

configure service vprn ospf3 overload-include-ext-2

Description

This command controls whether external type-2 routes should be readvertised with a maximum metric value when the system goes into overload state for any reason. When this command is enabled and the router is in overload, all external type-2 routes is advertised with the maximum metric.

The **no** form of this command reverts to the default value.

Default

no overload-include-ext-2

Platforms

7705 SAR Gen 2

overload-include-ext-2

Syntax

[**no**] **overload-include-ext-2**

Context

[Tree] (config>router>ospf overload-include-ext-2)

[Tree] (config>router>ospf3 overload-include-ext-2)

Full Context

configure router ospf overload-include-ext-2

configure router ospf3 overload-include-ext-2

Description

This command controls whether external type-2 routes should be readvertised with a maximum metric value when the system goes into overload state for any reason. When this command is enabled and the router is in overload, all external type-2 routes are advertised with the maximum metric.

The **no** form of this command reverts to the default value.

Default

no overload-include-ext-2

Platforms

7705 SAR Gen 2

20.51 overload-include-locators

overload-include-locators

Syntax

overload-include-locators

no overload-include-locators

Context

[Tree] (config>router>isis overload-include-locators)

Full Context

configure router isis overload-include-locators

Description

This command configures the router to include SRv6 locators when advertising links and prefixes with **max-metric** if the IS-IS instance goes into overload because of resource depletion or manual configuration.

The **no** form of this command causes the router to stop including the SRv6 locators when advertising links and prefixes with **max-metric**.

Default

no overload-include-locators

Platforms

7705 SAR Gen 2

20.52 overload-include-stub

overload-include-stub

Syntax

[**no**] **overload-include-stub**

Context

[Tree] (config>service>vprn>ospf overload-include-stub)

[Tree] (config>service>vprn>ospf3 overload-include-stub)

Full Context

configure service vprn ospf overload-include-stub

configure service vprn ospf3 overload-include-stub

Description

This command controls whether the OSPF stub networks should be advertised with a maximum metric value when the system goes into overload state for any reason. When enabled, the system uses the maximum metric value. When this command is enabled and the router is in overload, all stub interfaces, including loopback and system interfaces, are advertised at the maximum metric.

The **no** form of this command reverts to the default value.

Default

no overload-include-stub

Platforms

7705 SAR Gen 2

overload-include-stub

Syntax

[**no**] **overload-include-stub**

Context

[Tree] (config>router>ospf3 overload-include-stub)

[Tree] (config>router>ospf overload-include-stub)

Full Context

configure router ospf3 overload-include-stub

configure router ospf overload-include-stub

Description

This command controls whether the OSPF stub networks should be advertised with a maximum metric value when the system goes into overload state for any reason. When enabled, the system uses the maximum metric value. When this command is enabled and the router is in overload, all stub interfaces, including loopback and system interfaces, are advertised at the maximum metric.

The **no** form of this command reverts to the default value.

Default

no overload-include-stub

Platforms

7705 SAR Gen 2

20.53 overload-on-boot

overload-on-boot

Syntax

overload-on-boot [**timeout** *seconds*] [**max-metric**]

no overload-on-boot

Context

[Tree] (config>service>vprn>isis overload-on-boot)

Full Context

configure service vprn isis overload-on-boot

Description

When the router is in an overload state, it is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur:

- The timeout timer expires.
- A manual override of the current overload state is entered with the **config>router>isis>no overload** command.

  The **no overload** command does not affect the **overload-on-boot** function.

If no timeout is specified, IS-IS will go into overload indefinitely after a reboot. After the reboot, the IS-IS status will display a permanent overload state:

- L1 LSDB Overload : Manual on boot (Indefinitely in overload)
- L2 LSDB Overload : Manual on boot (Indefinitely in overload)

This state can be cleared with the **config>router>isis>no overload** command.

When specifying a timeout value, IS-IS will go into overload for the configured timeout after a reboot. After the reboot, the IS-IS status will display the remaining time the system stays in overload:

- L1 LSDB Overload : Manual on boot (Overload Time Left : 17)
- L2 LSDB Overload : Manual on boot (Overload Time Left : 17)

The overload state can be cleared before the timeout expires with the **config>router>isis>no overload** command.

The **no** form of this command removes the overload-on-boot functionality from the configuration.

Use the **show router isis status** command to display the administrative and operational state as well as all timers.

Default

no overload-on-boot

Parameters

timeout *seconds*

Configure the timeout timer for overload-on-boot in seconds.

Values 60 to 1800

max-metric

Set the maximum metric instead of overload.

Platforms

7705 SAR Gen 2

overload-on-boot

Syntax

overload-on-boot [**timeout** *seconds*]

no overload

Context

[Tree] (config>service>vprn>ospf3 overload-on-boot)

[Tree] (config>service>vprn>ospf overload-on-boot)

Full Context

configure service vprn ospf3 overload-on-boot

configure service vprn ospf overload-on-boot

Description

When the router is in an overload state, it is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur:

- The timeout timer expires.
- A manual override of the current overload state is entered with the **no overload** command.

The **no overload** command does not affect the **overload-on-boot** function.

The **no** form of this command removes the overload-on-boot functionality from the configuration.

Default

no overload-on-boot

Parameters

timeout seconds

Specifies the number of seconds to reset overloading.

Values 60 to1800

Default 60

Platforms

7705 SAR Gen 2

overload-on-boot

Syntax

overload-on-boot [**timeout** *seconds*] [**max-metric**]

no overload-on-boot

Context

[Tree] (config>router>isis overload-on-boot)

Full Context

configure router isis overload-on-boot

Description

When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur:

1. The timeout timer expires.
2. A manual override of the current overload state is entered with the **config>router>isis>no overload** command.

The **no overload** command does not affect the **overload-on-boot** function.

If no timeout is specified, IS-IS will go into overload indefinitely after a reboot. After the reboot, the IS-IS status will display a permanent overload state:

- L1 LSDB Overload : Manual on boot (Indefinitely in overload)
- L2 LSDB Overload : Manual on boot (Indefinitely in overload)

This state can be cleared with the **config>router>isis>no overload** command.

When specifying a timeout value, IS-IS will go into overload for the configured timeout after a reboot. After the reboot, the IS-IS status will display the remaining time the system stays in overload:

- L1 LSDB Overload : Manual on boot (Overload Time Left : 17)
- L2 LSDB Overload : Manual on boot (Overload Time Left : 17)

The overload state can be cleared before the timeout expires with the **config>router>isis>no overload** command.

The **no** form of this command removes the overload-on-boot functionality from the configuration.

Use the show router isis status command to display the administrative and operational state as well as all timers.

Default

no overload-on-boot

Parameters

seconds

Specifies the timeout timer for overload-on-boot, in seconds.

Values 60 to 1800

max-metric

Sets the maximum metric instead of overload.

Platforms

7705 SAR Gen 2

overload-on-boot

Syntax

overload-on-boot [**timeout** *seconds*]

no overload

Context

[Tree] (config>router>ospf3 overload-on-boot)

[Tree] (config>router>ospf overload-on-boot)

Full Context

configure router ospf3 overload-on-boot

configure router ospf overload-on-boot

Description

When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur:

- the timeout timer expires
- a manual override of the current overload state is entered with the **no overload** command

The **no overload** command does not affect the **overload-on-boot** function.

The **no** form of this command removes the overload-on-boot functionality from the configuration.

The default timeout value is 60 seconds, which means after 60 seconds overload status the SR will recover (change back to non-overload status). However, when **overload-on-boot** is configured under OSPF with no timeout value the router will remain in overload state indefinitely after a reboot.

Default

no overload-on-boot

Parameters

timeout seconds

Specifies the number of seconds to reset overloading.

Values 1 to 1800 in the **config**>**router**> **ospf** context
60 to 1800 in the **config**>**router**>**ospf3** context

Platforms

7705 SAR Gen 2

20.54 override

override

Syntax

[**no**] **override**

Context

[Tree] (config>router>pim>rp>static>address override)

[Tree] (config>router>pim>rp>ipv6>static>address override)

Full Context

configure router pim rp static address override

configure router pim rp ipv6 static address override

Description

This command changes the precedence of static RP over dynamically-learned Rendezvous Points (RPs).

When enabled, the static group-to-RP mappings take precedence over the dynamically learned mappings.

The **no** form of this command reverts to the default.

Default

no override

Platforms

7705 SAR Gen 2

20.55 override-bmi

override-bmi

Syntax

override-bmi *value*

no override-bmi

Context

[Tree] (config>router>isis>segm-rtng>msd override-bmi)

Full Context

configure router isis segment-routing maximum-sid-depth override-bmi

Description

This command provides the ability to override the announced MSD node Base MPLS Imposition (BMI). The MSD-BMI value announced by a router can be used by recipients to understand the number of MPLS labels that can be imposed inclusive of all service, transport, or special labels.

When **override-bmi** is not configured, the router announces the node maximum supported BMI assuming the most simple services and Layer 2 encapsulation.

The **no** form of this command reverts to the default.

Default

no override-bmi

Parameters

values

Specifies the override BMI.

Values 0 to 12

Platforms

7705 SAR Gen 2

20.56 override-erld

override-erld

Syntax

override-erld *value*

no override-erld

Context

[Tree] (config>router>isis>segm-rtng>msd override-erld)

Full Context

configure router isis segment-routing maximum-sid-depth override-erld

Description

This command provides the ability to override the announced MSD node Entropy Readable Label Depth (ERLD). It is useful for ingress LSRs to know each intermediate LSR's capability of reading the maximum label stack depth and performing EL-based load balancing.

When **override-erld** is not configured, then the router announces the node maximum supported ERLD assuming the most simple Layer 2 encapsulation.

The **no** form of this command reverts to the default.

Default

no override-erld

Parameters

values

Specifies the override ERLD.

Values 0 to 15

Platforms

7705 SAR Gen 2

20.57 own-auth-method

own-auth-method

Syntax

own-auth-method {**psk** | **cert** | **eap-only**}

no own-auth-method

Context

[Tree] (config>ipsec>ike-policy own-auth-method)

Full Context

configure ipsec ike-policy own-auth-method

Description

This command configures the authentication method used with this IKE policy on its own side.

Default

no own-auth-method

Platforms

7705 SAR Gen 2

21 p Commands – Part I

21.1 p2p-active-path-fast-retry

p2p-active-path-fast-retry

Syntax

p2p-active-path-fast-retry *seconds*

no p2p-active-path-fast-retry

Context

[Tree] (config>router>mpls p2p-active-path-fast-retry)

Full Context

configure router mpls p2p-active-path-fast-retry

Description

This command configures a global parameter to allow the user to apply a shorter retry timer for the first try after an active LSP path went down due to a local failure or the receipt of a ResvTear. This timer is used only in the first try. Subsequent retries will continue to be governed by the existing LSP level retry-timer.

The **no** form of this command disables the timer.

Default

no p2p-active-path-fast-retry

Parameters

seconds

Specifies the length of time for retry timer, in seconds

Values 1 to 10 seconds

Platforms

7705 SAR Gen 2

21.2 p2p-merge-point-abort-timer

p2p-merge-point-abort-timer

Syntax

p2p-merge-point-abort-timer *seconds*

no p2p-merge-point-abort-timer

Context

[Tree] (config>router>rsvp p2p-merge-point-abort-timer)

Full Context

configure router rsvp p2p-merge-point-abort-timer

Description

This command configures a timer to abort Merge-Point (MP) node procedures for a P2P LSP path. When a value higher than zero is configured for this timer, it will enter into effect anytime this node activates Merge-Point procedures for one or more P2P LSP paths. As soon an ingress interface goes operationally down, the Merge-Point node starts the abort timer. Upon expiry of the timer, MPLS will clean up all P2P LSP paths which ILM is on the failed interface and which have not already received a Path refresh over the bypass LSP.

The **no** form of this command disables the timer.

Default

no p2p-merge-point-abort-timer

Parameters

seconds

Specifies the length of the abort timer in seconds

Values 1 to 65535

Platforms

7705 SAR Gen 2

21.3 packet

packet

Syntax

packet [**query** | **v1-report** | **v2-report** | **v3-report** | **v2-leave**] [*ip-int-name* | *ip-address*] [**mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}]

packet [**query** | **v1-report** | **v2-report** | **v3-report** | **v2-leave**] [**mode** { **dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}] **group-interface** *ip-int-name*

packet [**query** | **v1-report** | **v2-report** | **v3-report** | **v2-leave**] **host** *ip-address* [**mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}]

no packet [**query** | **v1-report** | **v2-report** | **v3-report** | **v2-leave**] [*ip-int-name* | *ip-address*]

no packet [**query** | **v1-report** | **v2-report** | **v3-report** | **v2-leave**] **group-interface** *ip-int-name*

no packet [**query** | **v1-report** | **v2-report** | **v3-report** | **v2-leave**] **host** *ip-address*

Context

[Tree] (debug>router>igmp packet)

Full Context

debug router igmp packet

Description

This command enables/disables debugging for IGMP packets.

Parameters

query

Specifies to log the IGMP group- and source-specific queries transmitted and received on this interface.

v1-report

Specifies to debug IGMP V1 reports transmitted and received on this interface.

v2-report

Specifies to debug IGMP V2 reports transmitted and received on this interface.

v3-report

Specifies to debug IGMP V3 reports transmitted and received on this interface.

v2-leave

Specifies to debug the IGMP Leaves transmitted and received on this interface.

ip-int-name

Debugs the information associated with the specified IP interface name.

Values IP interface address

ip-address

Debugs the information associated with the specified IP address.

Platforms

7705 SAR Gen 2

packet

Syntax

packet [**detail**]

no packet

Context

[Tree] (debug>router>ldp>if packet)

[Tree] (debug>router>ldp>peer packet)

Full Context

debug router ldp interface packet

debug router ldp peer packet

Description

This command enables debugging for specific LDP packets.

The **no** form of the command disables the debugging output.

Parameters

detail

Displays detailed information.

Platforms

7705 SAR Gen 2

packet

Syntax

[**no**] **packet**

Context

[Tree] (debug>router>rsvp packet)

Full Context

debug router rsvp packet

Description

Commands in this context debug packets.

Platforms

7705 SAR Gen 2

packet

Syntax

packet [**hello** | **register** | **register-stop** | **jp** | **bsr** | **assert** | **crp** | **mdt-tlv** | **auto-rp-announcement** | **auto-rp-mapping** | **graft** | **graft-ack**] [*ip-int-name* | *mt-int-name* | *int-ip-address* | *mpls-if-name*] [**family** {**ipv4** | **ipv6**}] [**send** | **receive**]

no packet

Context

[Tree] (debug>router>pim packet)

Full Context

debug router pim packet

Description

This command enables debugging for PIM packets.

The **no** form of this command disables debugging for PIM packets.

Parameters

hello | register | register-stop | jp | bsr | assert | crp | mdt-tlv | auto-rp-announcement | auto-rp-mapping | graft | graft-ack

Specifies PIM packet types.

ip-int-name

Debugs the information associated with the specified IP interface name, up to 32 characters.

mt-int-name

Debugs the information associated with the specified VPRN ID and group address.

Values *vprn-id*-mt-*grp-ip-address*

int-ip-address

Debugs the information associated with the specified IP address.

ipv4

Specifies to display IPv4 packets.

ipv6

Specifies to display IPv6 packets.

mpls-if-name

Debugs the information associated with the specified MPLS interface.

Values mpls-if-*index*

receive

Specifies to display received packets.

send

Specifies to display sent packets.

family

Debugs database packet information.

Values ipv4, ipv6

Platforms

7705 SAR Gen 2

packet

Syntax

packet [{*ip-int-name* | *ip-address*}] [**headers**] [*protocol-id*]

no packet [{*ip-int-name* | *ip-address*}]

Context

[Tree] (debug>router>ip packet)

Full Context

debug router ip packet

Description

This command enables debugging for IP packets.

Parameters

ip-int-name

Only displays the interface information associated with the specified IP interface name.

Values 32 characters maximum

ip-address

Only displays the interface information associated with the specified IP address.

headers

Only displays information associated with the packet header.

protocol-id

Specifies the decimal value representing the IP protocol to debug. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17). The **no** form the command removes the protocol from the criteria.

Values 0 to 255 (values can be expressed in decimal, hexadecimal, or binary)

Platforms

7705 SAR Gen 2

packet

Syntax

[**no**] **packet**

Context

[Tree] (debug>router>rpki-session packet)

Full Context

debug router rpki-session packet

Description

This command enables debugging for specific RPKI packets.

The **no** form of this command disables debugging for specific RPKI packets.

Platforms

7705 SAR Gen 2

packet

Syntax

packet [*packet-type*] [*ip-int-name* | *ip-address*] [**detail**]

Context

[Tree] (debug>router>isis packet)

Full Context

debug router isis packet

Description

This command enables debugging for IS-IS packets.

The **no** form of the command disables debugging.

Parameters

ip-address

When specified, only packets with the specified interface address are debugged.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

ip-int-name

When specified, only packets with the specified interface name are debugged.

packet-type

When specified, only packets of the specified type are debugged.

Values ptop-hello | l1-hello | l2-hello | l1-psnp | l2-psnp | l1-csnp | l2-csnp | l1-lsp | l2-lsp

detail

All output is displayed in the detailed format.

Platforms

7705 SAR Gen 2

packet

Syntax

packet [*packet-type*] [*interface-name*] [**ingress** | **egress**] [**detail**]

packet [*packet-type*] [*interface-name*] [**ingress** | **egress** | **drop**] [**detail**]

no packet

Context

[Tree] (debug>router>ospf packet)

[Tree] (debug>router>ospf3 packet)

Full Context

debug router ospf packet

debug router ospf3 packet

Description

This command enables debugging for OSPF packets.

Parameters

packet-type

Specifies the OSPF packet type to debug.

Values hello, dbdescr, lsrequest, lsupdate, lsack

interface-name

Specifies the interface to debug, up to 32 characters.

ingress

Specifies to display ingress packets.

egress

Specifies to display egress packets.

drop

Specifies to display dropped packets.

Platforms

7705 SAR Gen 2

packet

Syntax

packet *packet-type* [**detail**]

no packet *packet-type*

Context

[Tree] (debug>router>pcep>pcc packet)

[Tree] (debug>router>pcep>pcc>conn packet)

Full Context

debug router pcep pcc packet

debug router pcep pcc connection packet

Description

This command enables debugging for PCEP PCC or connection packets.

The **no** form of this command disables debugging.

Parameters

packet-type

Specifies only packets of the specified type are debugged.

Values open | request | reply | notify | error | close | report | update | keepalive | pce-initiated

detail

Keyword used to specify detailed output.

Platforms

7705 SAR Gen 2

21.4 packet-byte-offset

packet-byte-offset

Syntax

packet-byte-offset {**add** *add-bytes* | **subtract** *sub-bytes*}

no packet-byte-offset

Context

[Tree] (config>card>fp>ingress>access>qgrp>policer-over>plcr packet-byte-offset)

[Tree] (config>card>fp>ingress>network>qgrp>policer-over>plcr packet-byte-offset)

Full Context

configure card fp ingress access queue-group policer-override policer packet-byte-offset

configure card fp ingress network queue-group policer-override policer packet-byte-offset

Description

This command modifies the size of each packet handled by the policer by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed. The **packet-byte-offset** command is meant to be an arbitrary mechanism the can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the policing metering and profiling throughput is affected by the offset as well as the stats associated with the policer.

When child policers are adding to or subtracting from the size of each packet, the parent policer's **min-thresh-separation** value should also need to be modified by the same amount.

The policer's **packet-byte-offset** defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command removes per packet size modifications from the policer.

Parameters

add-bytes

The **add** keyword is mutually exclusive to the **subtract** keyword. Either **add** or **subtract** must be specified. When **add** is defined the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate

metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 1 to 31

sub-bytes

The **subtract** keyword is mutually exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When b is defined the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet. Note that the minimum resulting packet size used by the system is 1 byte.

Values 0 to 32

Platforms

7705 SAR Gen 2

packet-byte-offset

Syntax

packet-byte-offset {**add** *add-bytes* | **subtract** *sub-bytes*}

no packet-byte-offset

Context

[Tree] (config>service>epipe>sap>ingress>policer-over>plcr packet-byte-offset)

[Tree] (config>service>epipe>sap>egress>policer-over>plcr packet-byte-offset)

Full Context

configure service epipe sap ingress policer-override policer packet-byte-offset

configure service epipe sap egress policer-override policer packet-byte-offset

Description

This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured packet-byte-offset parameter for the specified policer-id. Packet byte offset settings are not included in the applied rate when (queue) frame based accounting is configured; however, the offsets are applied to the statistics.

The **no** packet-byte-offset command is used to restore the policer's packet-byte-offset setting to the policy defined value.

Default

no packet-byte-offset

Parameters

add-bytes

Specifies the number of bytes that are added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 1 to 31

sub-bytes

Specifies the number of bytes that are subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

Values 1 to 64

Platforms

7705 SAR Gen 2

packet-byte-offset

Syntax

packet-byte-offset {**add** *add-bytes* | **subtract** *sub-bytes*}

no packet-byte-offset

Context

[Tree] (config>service>vpls>sap>egress>policer-override>plcr packet-byte-offset)

[Tree] (config>service>vpls>sap>ingress>policer-override>plcr packet-byte-offset)

Full Context

configure service vpls sap egress policer-override policer packet-byte-offset

configure service vpls sap ingress policer-override policer packet-byte-offset

Description

This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured packet-byte-offset parameter for the specified policer-id. Packet byte offset settings are not included in the applied rate when (queue) frame based accounting is configured, however the offsets are applied to the statistics.

The **no** form of this command restores the policer's packet-byte-offset setting to the policy defined value.

Default

no packet-byte-offset

Parameters

add-bytes

The add keyword is mutually exclusive to the subtract keyword. Either add or subtract must be specified. When add is defined the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 0 to 31

sub-bytes

The subtract keyword is mutually exclusive to the add keyword. Either add or subtract must be specified. When subtract is defined the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

Values 1 to 64

Platforms

7705 SAR Gen 2

packet-byte-offset

Syntax

packet-byte-offset {**add** *add-bytes* | **subtract** *sub-bytes*}

no packet-byte-offset

Context

[Tree] (config>service>ies>if>sap>egress>policer-override>plcr packet-byte-offset)

[Tree] (config>service>ies>if>sap>ingress>policer-override>plcr packet-byte-offset)

Full Context

configure service ies interface sap egress policer-override policer packet-byte-offset

configure service ies interface sap ingress policer-override policer packet-byte-offset

Description

This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured packet-byte-offset parameter for the specified policer-id. Packet byte offset settings are not included in the applied rate when (queue) frame based accounting is configured, however the offsets are applied to the statistics.

The **no** form of this command restores the policer's packet-byte-offset setting to the policy defined value.

Default

no packet-byte-offset

Parameters

add *add-bytes*

The **add** keyword is mutually exclusive to the **subtract** keyword. Either **add** or **subtract** must be specified. When **add** is defined, the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 0 to 31

subtract *sub-bytes*

The **subtract** keyword is mutually exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When **subtract** is defined, the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

Values 1 to 64

Platforms

7705 SAR Gen 2

packet-byte-offset

Syntax

packet-byte-offset add *add-bytes*

packet-byte-offset subtract *sub-bytes*

no packet-byte-offset

Context

[Tree] (config>service>vprn>if>sap>ingress>policer-override>plcr packet-byte-offset)

[Tree] (config>service>vprn>if>sap>egress>policer-override>plcr packet-byte-offset)

Full Context

configure service vprn interface sap ingress policer-override policer packet-byte-offset

configure service vprn interface sap egress policer-override policer packet-byte-offset

Description

This command, within the SAP ingress and egress policer-overrides contexts, is used to override the sap-ingress and sap-egress QoS policy configured packet-byte-offset parameter for the specified policer-id.

Packet byte offset settings are not included in the applied rate when (queue) frame based accounting is configured, however the offsets are applied to the statistics.

The **no** form of this command restores the policer's packet-byte-offset setting to the policy defined value.

Default

no packet-byte-offset

Parameters

add *add-bytes*

The **add** keyword is mutually exclusive to the **subtract** keyword. Either **add** or **subtract** must be specified. When **add** is defined, the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 0 to 31

subtract *sub-bytes*

The **subtract** keyword is mutually exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When **subtract** is defined, the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

Values 1 to 64

Platforms

7705 SAR Gen 2

packet-byte-offset

Syntax

packet-byte-offset {**add** *add-bytes* | **subtract** *sub- bytes*}

no packet-byte-offset

Context

[Tree] (config>qos>sap-ingress>policer packet-byte-offset)

[Tree] (config>qos>sap-egress>policer packet-byte-offset)

Full Context

configure qos sap-ingress policer packet-byte-offset

configure qos sap-egress policer packet-byte-offset

Description

This command is used to modify the size of each packet handled by the policer by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed. The **packet-byte-offset** command is meant to be an arbitrary mechanism that can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the policing metering and profiling throughput is affected by the offset as well as the stats associated with the policer.

When child policers are adding to or subtracting from the size of each packet, the parent policer's **min-thresh-separation** value should also need to be modified by the same amount.

The policer's **packet-byte-offset** defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied. Packet byte offset settings are not included in the applied rate when (queue) frame-based accounting is configured and the policer is managed by HQoS; however, the offsets are applied to the statistics.

The **no** form of this command is used to remove per packet size modifications from the policer.

Parameters

add *add-bytes*

The **add** keyword is mutually exclusive to the **subtract** keyword. Either **add** or **subtract** must be specified. When **add** is defined, the corresponding bytes parameter specifies the number of bytes that is added to the size of each packet associated with the policer for rate metering, profiling, and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 0 to 31

subtract *sub-bytes*

The **subtract** keyword is mutually exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When **subtract** is defined, the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling, and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet. The minimum resulting packet size used by the system is 1 byte.

Values 1 to 64

Platforms

7705 SAR Gen 2

packet-byte-offset

Syntax

packet-byte-offset {**add** *add-bytes* | **subtract** *sub-bytes*}

no packet-byte-offset

Context

[Tree] (config>qos>sap-ingress>queue packet-byte-offset)

Full Context

configure qos sap-ingress queue packet-byte-offset

Description

This command modifies the size of each packet handled by the queue by adding or subtracting the specified number of bytes. The actual packet size is not modified, only the size used to determine the ingress scheduling and profiling is changed. The **packet-byte-offset** command is an arbitrary mechanism that can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the scheduling and profiling throughput is affected by the offset as well as the statistics (accounting) associated with the queue. The **packet-byte-offset** does not apply to drop statistics, received valid statistics, or the offered managed and unmanaged statistics used by Ingress Multicast Path Management.

The **no** form of this command removes per-packet size modifications from the queue.

Parameters

add-bytes

Specifies the number of bytes added to the size of each packet associated with the queue for scheduling, profiling, and accounting purposes. From the queue's perspective, the packet size is increased by the amount specified.

Values 0 to 30, in increments of 2

sub-bytes

Specifies the number of bytes subtracted from the size of each packet associated with the queue for scheduling, profiling, and accounting purposes. From the queue's perspective, the packet size is reduced by the amount specified. The minimum resulting packet size used by the system is 1 byte.

Values 0 to 64, in increments of 2

Platforms

7705 SAR Gen 2

packet-byte-offset

Syntax

packet-byte-offset {**add** *add-bytes* | **subtract** *sub-bytes*}

no packet-byte-offset

Context

[Tree] (config>qos>sap-egress>queue packet-byte-offset)

Full Context

configure qos sap-egress queue packet-byte-offset

Description

This command is used to modify the size of each packet handled by the queue by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed.

The packet-byte-offset command is meant to be an arbitrary mechanism the can be used to either add downstream frame encapsulation or remove portions of packet headers.

When a packet-byte-offset value is applied to a queue instance, it adjusts the immediate packet size. This means that the queue rates, in other words, operational PIR and CIR, and queue bucket updates use the adjusted packet size. In addition, the queue statistics also reflects the adjusted packet size. Scheduler policy rates, which are data rates, uses the adjusted packet size.

The port scheduler max-rate and the priority level rates and weights, if a Weighted Scheduler Group is used, are always on-the-wire rates and uses the actual frame size. The same applies for the agg-rate-limit on a SAP, a subscriber, or a multiservice Site (MSS) when the queue is port-parented.

When the user enables frame-based-accounting in a scheduler policy or queue-frame-based-accounting with agg-rate-limit in a port scheduler policy, the queue rate will be capped to a user-configured on-the-wire rate and the packet-byte-offset is not included. However, the offsets are applied to the statistics.

The **no** form of this command is used to remove per packet size modifications from the queue.

Parameters

add-bytes

The **add** keyword is mutually exclusive to the **subtract** keyword. Either parameter must be specified. When **add** is defined, the corresponding bytes parameter specifies the number of bytes that is added to the size of each packet associated with the queue for scheduling and accounting purposes.

Values 0 to 32

sub-bytes

The **subtract** keyword is mutually exclusive to the **add** keyword. Either parameter must be specified. When subtract is defined, the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the queue for scheduling and accounting purposes. The minimum resulting packet size used by the system is 1 byte.

Values 0 to 64

Platforms

7705 SAR Gen 2

packet-byte-offset

Syntax

packet-byte-offset {**add** *add-bytes* | **subtract** *sub-bytes*}

no packet-byte-offset

Context

[Tree] (config>qos>qgrps>ing>qgrp>policer packet-byte-offset)

Full Context

configure qos queue-group-templates ingress queue-group policer packet-byte-offset

Description

This command configures a packet byte offset for the QoS ingress queue-group policer.

Default

no packet-byte-offset

Parameters

add-bytes

Specifies the number of bytes to add as the offset amount.

Values 0 to 31

sub-bytes

Specifies the number of bytes to add as the offset amount.

Values 1 to 32

Platforms

7705 SAR Gen 2

packet-byte-offset

Syntax

packet-byte-offset {**add** *add-bytes* | **subtract** *sub-bytes*}

no packet-byte-offset

Context

[Tree] (config>qos>qgrps>ing>qgrp>queue packet-byte-offset)

Full Context

configure qos queue-group-templates ingress queue-group queue packet-byte-offset

Description

This command is used to modify the size of each packet handled by the queue by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the ingress scheduling and profiling is changed. The packet-byte-offset command is meant to be an arbitrary mechanism that can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the scheduling and profiling throughput is affected by the offset as well as the

stats (accounting) associated with the queue. The packet-byte-offset does not apply to drop statistics, received valid statistics, or the offered managed and unmanaged statistics used by Ingress Multicast Path Management.

The **no** form of this command is used to remove per packet size modifications from the queue.

Parameters

add-bytes

The add keyword is mutually exclusive to the subtract keyword. Either add or subtract must be specified. When add is defined, the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the queue for scheduling, profiling and accounting purposes. From the queue's perspective, the packet size is increased by the amount being added to the size of each packet.

Values 0 to 30, in steps of 2

sub-bytes

The subtract keyword is mutually exclusive to the add keyword. Either add or subtract must be specified. When subtract is defined, the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the queue for scheduling, profiling and accounting purposes. From the queue's perspective, the packet size is reduced by the amount being subtracted from the size of each packet. The minimum resulting packet size used by the system is 1 byte.

Values 0 to 64, in steps of 2

Platforms

7705 SAR Gen 2

packet-byte-offset

Syntax

packet-byte-offset {**add add-***bytes* | **subtract** *sub-bytes*}

no packet-byte-offset

Context

[Tree] (config>qos>qgrps>egr>qgrp>policer packet-byte-offset)

[Tree] (config>qos>qgrps>egr>qgrp>queue packet-byte-offset)

Full Context

configure qos queue-group-templates egress queue-group policer packet-byte-offset

configure qos queue-group-templates egress queue-group queue packet-byte-offset

Description

This command is used to modify the size of each packet handled by the queue by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed.

The packet-byte-offset command is meant to be an arbitrary mechanism that can be used to either add downstream frame encapsulation or remove portions of packet headers.

When a packet-byte-offset value is applied to a queue or policer instance, it adjusts the immediate packet size. This means that the queue rates (in other words, operational PIR and CIR) and policer or queue bucket updates use the adjusted packet size. In addition, the statistics also reflect the adjusted packet size. Scheduler policy rates, which are data rates, will use the adjusted packet size.

The port scheduler **max-rate** and the priority level rates and weights, if a Weighted Scheduler Group is used, are always on-the-wire rates and uses the actual frame size. The same applies for the agg-rate-limit on a SAP, a subscriber, or a Multiservice Site (MSS) when the queue is port-parented.

When the user enables **frame-based-accounting** in a scheduler policy or **queue-frame-based-accounting** with agg-rate-limit in a port scheduler policy, the policer or queue rate is capped to a user-configured on-the-wire rate and the packet-byte-offset is not included; however, the offsets are applied to the statistics.

The **no** form of this command is used to remove per packet size modifications from the queue.

Parameters

add-bytes

Specifies that the corresponding bytes parameter specifies the number of bytes that is added to the size of each packet associated with the queue for scheduling and accounting purposes.

Values 0 to 32

sub-bytes

Specifies that the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the queue for scheduling and accounting purposes. The minimum resulting packet size used by the system is 1 byte.

Values 0 to 64

Platforms

7705 SAR Gen 2

21.5 packet-size

packet-size

Syntax

packet-size *bytes*

no packet-size

Context

[Tree] (config>system>snmp packet-size)

Full Context

configure system snmp packet-size

Description

This command configures the maximum SNMP packet size generated by this node.

The **no** form of this command restores the default value.

Default

packet-size 1500

Parameters

bytes

Specifies the SNMP packet size in bytes.

Values 484 to 9216

Platforms

7705 SAR Gen 2

21.6 packet-too-big

packet-too-big

Syntax

packet-too-big [*number seconds*]

no packet-too-big

Context

[Tree] (config>service>ies>if>ipv6>icmp6 packet-too-big)

Full Context

configure service ies interface ipv6 icmp6 packet-too-big

Description

This command specifies whether packet-too-big ICMP messages should be sent. When enabled, ICMPv6 packet-too-big messages are generated by this interface.

The **no** form of this command disables the sending of ICMPv6 packet-too-big messages.

Default

packet-too-big 100 10

Parameters

number

Specifies the number of ICMP messages that are too large to send in the time frame specified by the *seconds* parameter.

Values 10 to 1000

Default 100

seconds

Specifies the time frame, in seconds, that is used to limit the number of "packet-too-big" ICMP messages issued.

Values 1 to 60

Default 10

Platforms

7705 SAR Gen 2

packet-too-big

Syntax

packet-too-big [*number seconds*]

no packet-too-big

Context

[Tree] (config>service>vprn>if>ipv6>icmp6 packet-too-big)

Full Context

configure service vprn interface ipv6 icmp6 packet-too-big

Description

This command configures the rate for Internet Control Message Protocol version 6 (ICMPv6) packet-too-big messages.

Parameters

number

Specifies the number of packet-too-big messages to send in the time frame specified by the *seconds* parameter.

Values 10 to 1000

Default 100

seconds

Specifies the time frame, in seconds, that is used to limit the number of packet-too-big messages issued.

Values 1 to 60

Default 10

Platforms

7705 SAR Gen 2

packet-too-big

Syntax

packet-too-big

packet-too-big number [**10..1000**] **seconds** [**1..60**]

no packet-too-big

Context

[Tree] (config>service>vprn>if>sap>ip-tunnel>icmp6-gen packet-too-big)

Full Context

configure service vprn interface sap ip-tunnel icmp6-generation packet-too-big

Description

This command enables the system to send ICMPv6 PTB (Packet Too Big) messages on the private side and optionally specifies the rate.

With this command configured, the system sends PTB back if it received an IPv6 packet on the private side that is bigger than 1280 bytes and also exceeds the private MTU of the tunnel.

The **ip-mtu** command (under **ipsec-tunnel** or **tunnel-template**) specifies the private MTU for the ipsec-tunnel or dynamic tunnel.

The **no** form of this command reverts **interval** and **message-count** values to their default values.

Platforms

7705 SAR Gen 2

packet-too-big

Syntax

packet-too-big [*number seconds*]

no packet-too-big

Context

[Tree] (config>router>if>ipv6>icmp6 packet-too-big)

Full Context

configure router interface ipv6 icmp6 packet-too-big

Description

This command configures the rate for ICMPv6 packet-too-big messages.

Parameters

number

Limits the number of packet-too-big messages issued per time frame specified in the *seconds* parameter.

Values 10 to 1000

seconds

Determines the time frame, in seconds, that is used to limit the number of packet-too-big messages issued per time frame.

Values 1 to 60

Platforms

7705 SAR Gen 2

21.7 packet-type

packet-type

Syntax

packet-type [**authentication**] [**accounting**] [**coa**]

no packet-type

Context

[Tree] (debug>router>radius packet-type)

Full Context

debug router radius packet-type

Description

This command specifies the RADIUS packet type filter of command **debug router radius**.

Default

authentication accounting coa

Parameters

authentication

Specifies the RADIUS authentication packet.

accounting

Specifies the RADIUS accounting packet.

coa

Specifies the RADIUS change of authorization packet.

Platforms

7705 SAR Gen 2

21.8 packets

packets

Syntax

[**no**] **packets**

[**no**] **packets interface** *ip-int-name* [**vrid** *virtual-router-id*]

[**no**] **packets interface** *ip-int-name* **vrid** *virtual-router-id* **ipv6**

Context

[Tree] (debug>router>vrrp packets)

Full Context

debug router vrrp packets

Description

This command enables or disables debugging for VRRP packets.

Parameters

ip-int-name

Specifies the interface name, up to 32 characters.

virtual-router-id

Specifies the router ID.

Values 1 to 255

ipv6

Debugs the specified IPv6 VRRP interface.

Platforms

7705 SAR Gen 2

packets

Syntax

packets [**neighbor** *ip-address* | **group** *name*]

no packets

Context

[Tree] (debug>router>bgp packets)

Full Context

debug router bgp packets

Description

This command decodes and logs all sent and received BGP packets in the debug log.

The **no** form of this command disables debugging.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x [-interface] (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H
- d: [0 to 255]D
- interface: up to 32 characters for link local addresses

group *name*

Debugs only events affecting the specified peer group name, up to 64 characters, and associated neighbors.

Platforms

7705 SAR Gen 2

packets

Syntax

[**no**] **packets** [**neighbor** *ip-int-name* | *ip-addr*]

Context

[Tree] (debug>router>rip packets)

Full Context

debug router rip packets

Description

This command enables debugging for RIP packets.

Parameters

ip-int-name* | *ip-address

Debugs the RIP packets sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

packets

Syntax

[**no**] **packets** [**neighbor** *ip-int-name* | *ipv6-address*]

Context

[Tree] (debug>router>ripng packets)

Full Context

debug router ripng packets

Description

This command enables debugging for RIPng packets.

Parameters

ip-int-name*| *ipv6-address

Debugs the RIPng packets sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

21.9 pad-size

pad-size

Syntax

pad-size *octets*

no pad-size

Context

[Tree] (config>oam-pm>session>ip>twamp-light pad-size)

Full Context

configure oam-pm session ip twamp-light pad-size

Description

This command defines the amount by which the TWAMP Light packet is padded. TWAMP session controller packets are 27 bytes smaller than TWAMP session reflector packets. If symmetrical packet sizes in the forward and backward direction are required, the pad size must be configured to a minimum of 27 bytes.

The **no** form of this command removes all padding.

Default

pad-size 0

Parameters

octets

Specifies the value, in octets, to pad the TWAMP Light packet.

Values 0 to 2000

Default 0

Platforms

7705 SAR Gen 2

21.10 pad-tlv-size

pad-tlv-size

Syntax

pad-tlv-size *octets* [**create**]

no pad-tlv-size

Context

[Tree] (config>oam-pm>session>ip>twamp-light pad-tlv-size)

Full Context

configure oam-pm session ip twamp-light pad-tlv-size

Description

This command configures the PAD TLV to be included in the STAMP test packet with a total byte count equivalent to the value of this leaf.

TWAMP Light does not support TLVs. To pad the size of the TWAMP Light test packet the user must configure the **pad-size** command. STAMP test packets (the standard form of TWAMP Light) introduces TLVs for padding. Therefore, STAMP test packets must use the pad-tlv-size value.

The **no** form of this command removes the TWAMP Light test function from the OAM-PM session.

Parameters

test-id

Specifies the value of the 4-byte local test identifier not sent in the TWAMP Light packets.

Values 0 to 2147483647

create

Creates the test.

Platforms

7705 SAR Gen 2

21.11 padding-size

padding-size

Syntax

padding-size *padding-size*

no padding-size

Context

[Tree] (config>service>vprn>static-route-entry>indirect>cpe-check padding-size)

[Tree] (config>service>vprn>static-route-entry>next-hop>cpe-check padding-size)

Full Context

configure service vprn static-route-entry indirect cpe-check padding-size

configure service vprn static-route-entry next-hop cpe-check padding-size

Description

This optional parameter specifies the amount of padding to add to the ICMP packet in bytes. The parameter is only applicable when the **cpe-check** option is used with the associated static route.

Default

padding-size 56

Parameters

padding-size

An integer value.

Values 0 to 16384 bytes

Platforms

7705 SAR Gen 2

padding-size

Syntax

padding-size *padding-size*

no padding-size

Context

[Tree] (config>router>static-route-entry>indirect>cpe-check padding-size)

[Tree] (config>router>static-route-entry>next-hop>cpe-check padding-size)

Full Context

configure router static-route-entry indirect cpe-check padding-size

configure router static-route-entry next-hop cpe-check padding-size

Description

This command specifies the amount of padding to add to the ICMP packet in bytes. The parameter is only applicable when the **cpe-check** option is used with the associated static route.

Default

padding-size 56

Parameters

padding-size

Specifies the integer value.

Values 0 to 16384 bytes

Platforms

7705 SAR Gen 2

padding-size

Syntax

padding-size *size*

no padding-size

Context

[Tree] (config>vrrp>policy>priority-event>host-unreachable padding-size)

Full Context

configure vrrp policy priority-event host-unreachable padding-size

Description

This command allows the operator to increase the size of IP packet by padding the PDU.

The **no** form of the command reverts to the default.

Default

padding-size 0

Parameters

size

Specifies amount of increase to the ICMP PDU.

Values 0 to 16384

Platforms

7705 SAR Gen 2

21.12 parallel

parallel

Syntax

parallel [**no-advertise**]

no parallel

Context

[Tree] (config>router>ospf>segm-rtng>adjacency-set parallel)

[Tree] (config>router>isis>segm-rtng>adjacency-set parallel)

Full Context

configure router ospf segment-routing adjacency-set parallel

configure router isis segment-routing adjacency-set parallel

Description

This command indicates that all members of the adjacency set must terminate on the same neighboring node. The system raises a trap if a user attempts to add an adjacency terminating on a neighboring node that differs from the existing members of the adjacency set. In addition, the system stops advertising the adjacency set in IS-IS or OSPF and locally deprograms it.

By default, parallel adjacency sets are advertised in the IGP. The **no-advertise** option prevents an adjacency set from being advertised in the IGP. It is only allowed in CLI and SNMP if the **parallel** command is configured.

The **no** form of this command indicates that the adjacency set can include adjacencies to different next hop nodes.

Default

parallel

Platforms

7705 SAR Gen 2

21.13 param-problem

param-problem

Syntax

param-problem [*number seconds*]

no param-problem

Context

[Tree] (config>service>ies>if>icmp param-problem)

[Tree] (config>service>ies>if>ipv6>icmp6 param-problem)

Full Context

configure service ies interface icmp param-problem

configure service ies interface ipv6 icmp6 param-problem

Description

This command specifies whether parameter-problem ICMP/ICMPv6 messages should be sent. When enabled, parameter-problem ICMP/ICMPv6 messages are generated by this interface.

The **no** form of this command disables the sending of parameter-problem ICMP/ICMPv6 messages.

Default

param-problem 100 10

Parameters

number

Specifies the number of parameter-problem ICMPv6 messages to send in the time frame specified by the *seconds* parameter.

Values 10 to 1000

Default 100

seconds

Specifies the time frame, in seconds, that is used to limit the number of parameter-problem ICMPv6 messages issued.

Values 1 to 60

Default 10

Platforms

7705 SAR Gen 2

param-problem

Syntax

param-problem *number seconds*

no param-problem [*number seconds*]

Context

[Tree] (config>service>vprn>if>icmp param-problem)

[Tree] (config>service>vprn>if>ipv6>icmp6 param-problem)

[Tree] (config>service>vprn>nw-if>icmp param-problem)

Full Context

configure service vprn interface icmp param-problem

configure service vprn interface ipv6 icmp6 param-problem

configure service vprn network-interface icmp param-problem

Description

This command specifies whether parameter-problem ICMP messages should be sent. When enabled, parameter-problem ICMP messages are generated by this interface. The **no** form of this command disables the sending of parameter-problem ICMP messages.

Parameters

number

Specifies the number of parameter-problem ICMP messages to send in the time frame specified by the *seconds* parameter.

Values 10 to 1000

Default 100

seconds

Specifies the time frame, in seconds, that is used to limit the number of parameter-problem ICMP messages issued.

Values 1 to 60

Default 10

Platforms

7705 SAR Gen 2

param-problem

Syntax

param-problem [*number seconds*]

no param-problem

Context

[Tree] (config>router>if>icmp param-problem)

Full Context

configure router interface icmp param-problem

Description

This command specifies whether parameter-problem ICMP messages should be sent. When enabled, parameter-problem ICMP messages are generated by this interface.

The **no** form of this command disables the sending of parameter-problem ICMP messages.

Parameters

number

Specifies the number of parameter-problem ICMP messages to send in the time frame specified by the *seconds* parameter.

Values 10 to 1000

Default 100

seconds

Specifies the time frame, in seconds, that is used to limit the number of parameter-problem ICMP messages issued.

Values 1 to 60

Default 10

Platforms

7705 SAR Gen 2

param-problem

Syntax

param-problem [*number seconds*]

no param-problem

Context

[Tree] (config>router>if>ipv6>icmp6 param-problem)

Full Context

configure router interface ipv6 icmp6 param-problem

Description

This command specifies whether parameter-problem ICMPv6 messages should be sent. When enabled, parameter-problem ICMPv6 messages are generated by this interface.

The **no** form of this command disables the sending of parameter-problem ICMPv6 messages.

Parameters

number

Specifies the number of parameter-problem ICMPv6 messages to send in the time frame specified by the *seconds* parameter.

Values 10 to 1000

Default 100

seconds

Specifies the time frame, in seconds, that is used to limit the number of parameter-problem ICMPv6 messages issued.

Values 1 to 60

Default 10

Platforms

7705 SAR Gen 2

21.14 parent

parent

Syntax

parent [**weight** *weight*] [**cir-weight** *cir-weight*]

no parent

Context

[Tree] (config>port>ethernet>access>egr>qgrp>qover>q parent)

Full Context

configure port ethernet access egress queue-group queue-overrides queue parent

Description

This command, when used in the *queue-overrides* context for a queue group queue, defines an optional **weight** and **cir-weight** for the queue treatment by the parent scheduler that further governs the available bandwidth for the queue aside from the queue PIR setting. When multiple schedulers and or queues share a child status with the parent scheduler, the weight or level parameters define how this queue contends with the other children for the parent bandwidth.

Parameters

weight

Specifies the relative weight of this queue in comparison to other child schedulers and queues while vying for bandwidth on the parent scheduler-name. Any queues or schedulers defined as weighted receive no parental bandwidth until all strict queues and schedulers on the parent have reached their maximum bandwidth or are idle. In this manner, weighted children are considered to be the lowest priority.

Values 0 to 100

Default 1

cir-weight

Specifies the weight the queue uses at the within-cir port priority level. The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 100

Platforms

7705 SAR Gen 2

parent

Syntax

parent [**weight** *weight*] [**cir-weight** *cir-weight*]

no parent

Context

[Tree] (config>port>ethernet>access>ing>qgrp>sched-override>scheduler parent)

[Tree] (config>port>ethernet>access>egr>qgrp>sched-override>scheduler parent)

Full Context

configure port ethernet access ingress queue-group scheduler-override scheduler parent

configure port ethernet access egress queue-group scheduler-override scheduler parent

Description

This command can be used to override the scheduler's parent weight and CIR weight. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the applied scheduler policy.

The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy - this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the queue group overrides. If the parent scheduler does not exist, causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non-default weightings for fostered schedulers.

The **no** form of this command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.

Default

no parent

Parameters

weight

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict level defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total distributes the available bandwidth at that level. A weight is considered to be active when the queue or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.

A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.

Values 0 to 100

cir-weight

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same cir-level defined by the cir-level parameter in the applied scheduler policy. Within the strict cir-level, all cir-weight values from active children at that level are summed and the ratio of each active child's cir-weight to the total distributes the available bandwidth at that level. A cir-weight is considered to be active when the policer, queue, or scheduler that the cir-weight pertains to has not reached the CIR and still has packets to transmit.

A 0 (zero) cir-weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.

Values 0 to 100

Platforms

7705 SAR Gen 2

parent

Syntax

parent {[**weight** *weight*] [**cir-weight** *cir-weight*]}

no parent

Context

[Tree] (config>service>epipe>sap>egress>queue-override>queue parent)

[Tree] (config>service>epipe>sap>ingress>queue-override>queue parent)

Full Context

configure service epipe sap egress queue-override queue parent

configure service epipe sap ingress queue-override queue parent

Description

This command defines an optional parent scheduler that further governs the available bandwidth given the queue aside from the queue's PIR setting. When multiple schedulers and/or queues share a child status with the parent scheduler, the **weight** or **level** parameters define how this queue contends with the other children for the parent's bandwidth.

Checks are not performed to see if a *scheduler-name* exists when the parent command is defined on the queue. Scheduler names are configured in the **config>qos>scheduler-policy>tier** *level* context. Multiple schedulers can exist with the *scheduler-name* and the association pertains to a scheduler that should exist on the egress SAP as the policy is applied and the queue created. When the queue is created on the egress SAP, the existence of the *scheduler-name* is dependent on a scheduler policy containing the *scheduler-name* being directly or indirectly applied (through a multi-service customer site) to the egress SAP. If the *scheduler-name* does not exist, the queue is placed in the orphaned operational state. The queue will accept packets but will not be bandwidth limited by a virtual scheduler or the scheduler hierarchy applied to the SAP. The orphaned state must generate a log entry and a trap message. The SAP which the queue belongs to must also depict an orphan queue status. The orphaned state of the queue is automatically cleared when the *scheduler-name* becomes available on the egress SAP.

The parent scheduler can be made unavailable due to the removal of a scheduler policy or scheduler. When an existing parent scheduler is removed or inoperative, the queue enters the orphaned state and automatically returns to normal operation when the parent scheduler is available again.

When a parent scheduler is defined without specifying weight or strict parameters, the default bandwidth access method is weight with a value of 1.

The **no** form of this command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. Once a parent association has been removed, the former child queue attempts to operate based on its configured rate parameter. Removing the parent association on the queue within the policy takes effect immediately on all queues using the SAP egress QoS policy.

Parameters

weight

These optional keywords are mutually exclusive to the **level** keyword. Specifies the relative weight of this queue in comparison to other child schedulers, policers, and queues while vying for bandwidth on the parent *scheduler-name*. Any policers, queues, or schedulers defined as weighted receive no parental bandwidth until all policers, queues, and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

All **weight** values from all weighted active policers, queues, and schedulers with a common parent scheduler are added together. Then, each individual active weight is divided by the total, deriving the percentage of remaining bandwidth provided to the policer, queue, or scheduler. A weight is considered to be active when the pertaining policer, queue, or scheduler has not reached its maximum rate and still has packets to transmit. All child policers, queues, and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all non-zero weighted policers, queues, and schedulers at that level are operating at the maximum bandwidth or are idle.

Values 0 to 100

Default 1

cir-weight

Specifies the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the policer, queue, or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 100

Platforms

7705 SAR Gen 2

parent

Syntax

parent [**weight** *weight*] [**cir-weight** *cir-weight*]

no parent

Context

[Tree] (config>service>epipe>sap>egress>sched-override>scheduler parent)

[Tree] (config>service>epipe>sap>ingress>sched-override>scheduler parent)

Full Context

configure service epipe sap egress scheduler-override scheduler parent

configure service epipe sap ingress scheduler-override scheduler parent

Description

This command can be used to override the scheduler's parent weight and cir-weight information. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site.

The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy – this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non-default weightings for fostered schedulers.

The **no** form of this command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.

Default

no parent

Parameters

weight

Specifies the relative weight of this scheduler in comparison to other child schedulers, policers, and queues at the same strict **level** defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the policer, queue, or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.

A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.

Values 0 to 100

cir-weight

Specifies the relative weight of this scheduler in comparison to other child schedulers, policers, and queues at the same *cir-level* defined by the **cir-level** parameter in the applied scheduler policy. Within the strict **cir-level**, all **cir-weight** values from active children at that level are summed and the ratio of each active child's **cir-weight** to the total is used to distribute the available bandwidth at that level. A **cir-weight** is considered to be active when the policer, queue, or scheduler that the **cir-weight** pertains to has not reached the CIR and still has packets to transmit.

A 0 (zero) **cir-weight** value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.

Values 0 to 100

Platforms

7705 SAR Gen 2

parent

Syntax

parent [**weight** *weight*] [**cir-weight** *cir-weight*]

no parent

Context

[Tree] (config>service>vpls>sap>egress>queue-override>queue parent)

[Tree] (config>service>vpls>sap>ingress>queue-override>queue parent)

Full Context

configure service vpls sap egress queue-override queue parent

configure service vpls sap ingress queue-override queue parent

Description

This command defines an optional parent scheduler that further governs the available bandwidth given the queue aside from the queue's PIR setting. When multiple schedulers and/or queues share a child status with the parent scheduler, the **weight** or **level** parameters define how this queue contends with the other children for the parent's bandwidth.

Checks are not performed to see if a *scheduler-name* exists when the parent command is defined on the queue. Scheduler names are configured in the **config>qos>scheduler-policy>tier** *level* context. Multiple schedulers can exist with the *scheduler-name* and the association pertains to a scheduler that should exist on the egress SAP as the policy is applied and the queue created. When the queue is created on the egress SAP, the existence of the *scheduler-name* is dependent on a scheduler policy containing the *scheduler-name* being directly or indirectly applied (through a multi-service customer site) to the egress SAP. If the *scheduler-name* does not exist, the queue is placed in the orphaned operational state. The queue will accept packets but will not be bandwidth limited by a virtual scheduler or the scheduler hierarchy applied to the SAP. The orphaned state must generate a log entry and a trap message. The SAP which the queue belongs to must also depict an orphan queue status. The orphaned state of the queue is automatically cleared when the *scheduler-name* becomes available on the egress SAP.

The parent scheduler can be made unavailable due to the removal of a scheduler policy or scheduler. When an existing parent scheduler is removed or inoperative, the queue enters the orphaned state mentioned above and automatically return to normal operation when the parent scheduler is available again.

When a parent scheduler is defined without specifying weight or strict parameters, the default bandwidth access method is weight with a value of 1.

The **no** form of this command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. Once a parent association has been removed, the former child queue attempts to operate based on its configured rate parameter. Removing the parent association on the queue within the policy takes effect immediately on all queues using the SAP egress QoS policy.

Parameters

weight

These optional keywords are mutually exclusive to the **level** keyword. The weight defines the relative weight of this queue in comparison to other child schedulers, policers, and queues while vying for bandwidth on the parent *scheduler-name*. Any policers, queues, or schedulers defined as weighted receive no parental bandwidth until all policers, queues, and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

All **weight** values from all weighted active policers, queues, and schedulers with a common parent scheduler are added together. Then, each individual active weight is divided by the total, deriving the percentage of remaining bandwidth provided to the policer, queue, or scheduler. A weight is considered to be active when the pertaining policer, queue, or scheduler has not reached its maximum rate and still has packets to transmit. All child policers, queues, and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all non-zero weighted policers, queues, and schedulers at that level are operating at the maximum bandwidth or are idle.

Values 0 to 100

Default 1

cir-weight

Specifies the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the policer, queue, or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 100

Platforms

7705 SAR Gen 2

parent

Syntax

parent [**weight** *weight*] [**cir-weight** *cir-weight*]

no parent

Context

[Tree] (config>service>vpls>sap>ingress>sched-override>scheduler parent)

[Tree] (config>service>vpls>sap>egress>sched-override>scheduler parent)

Full Context

configure service vpls sap ingress scheduler-override scheduler parent

configure service vpls sap egress scheduler-override scheduler parent

Description

This command can be used to override the scheduler's parent weight and cir-weight information. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site.

The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy – this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non-default weightings for fostered schedulers.

The **no** form of this command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.

Default

no parent

Parameters

weight

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict **level** defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the policer, queue, or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.

A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.

Values 0 to 100

cir-weight

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same *cir-level* defined by the **cir-level** parameter in the applied scheduler policy. Within the strict **cir-level**, all **cir-weight** values from active children at that level are summed and the ratio of each active child's **cir-weight** to the total is used to distribute the available bandwidth at that level. A **cir-weight** is considered to be active when the policer, queue, or scheduler that the **cir-weight** pertains to has not reached the CIR and still has packets to transmit.

A 0 (zero) **cir-weight** value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.

Values 0 to 100

Platforms

7705 SAR Gen 2

parent

Syntax

parent [**weight** *weight*] [**cir-weight** *cir-weight*]

Context

[Tree] (config>service>ies>if>sap>ingress>sched-override>scheduler parent)

[Tree] (config>service>ies>if>sap>ingress>queue-override>queue parent)

[Tree] (config>service>ies>if>sap>egress>queue-override>queue parent)

[Tree] (config>service>ies>if>sap>egress>sched-override>scheduler parent)

Full Context

configure service ies interface sap ingress scheduler-override scheduler parent

configure service ies interface sap ingress queue-override queue parent

configure service ies interface sap egress queue-override queue parent

configure service ies interface sap egress scheduler-override scheduler parent

Description

This command can be used to override the scheduler's parent weight and cir-weight information. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site.

The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy – this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non default weightings for fostered schedulers.

The **no** form of this command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.

Default

no parent

Parameters

weight *weight*

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict **level** defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the queue or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.

A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.

Values 0 to 100

cir-weight ***cir-weight***

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same *cir-level* defined by the **cir-level** parameter in the applied scheduler policy. Within the strict **cir-level**, all **cir-weight** values from active children at that level are summed and the ratio of each active child's **cir-weight** to the total is used to distribute the available bandwidth at that level. A **cir-weight** is considered to be active when the policer, queue, or scheduler that the **cir-weight** pertains to has not reached the CIR and still has packets to transmit.

A 0 (zero) **cir-weight** value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.

Values 0 to 100

Platforms

7705 SAR Gen 2

parent

Syntax

parent [**weight** *weight*] [**cir-weight cir-***weight*]

no parent

Context

[Tree] (config>service>vprn>if>sap>egress>queue-override>queue parent)

[Tree] (config>service>vprn>if>sap>ingress>queue-override>queue parent)

Full Context

configure service vprn interface sap egress queue-override queue parent

configure service vprn interface sap ingress queue-override queue parent

Description

This command can be used to override the scheduler's parent weight and cir-weight information. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site.

The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy – this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non default weightings for fostered schedulers.

The **no** form of this command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.

Default

no parent

Parameters

weight *weight*

Specifes the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict **level** defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the queue or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.

A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.

Values 0 to 100

Default 1

cir-weight *cir-weight*

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same *cir-level* defined by the **cir-level** parameter in the applied scheduler policy. Within the strict **cir-level**, all **cir-weight** values from active children at that level are summed and the ratio of each active child's **cir-weight** to the total is used to distribute the available bandwidth at that level. A **cir-weight** is considered to be active when the policer, queue, or scheduler that the **cir-weight** pertains to has not reached the CIR and still has packets to transmit.

A 0 (zero) **cir-weight** value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.

Values 0 to 100

Default 1

Platforms

7705 SAR Gen 2

parent

Syntax

parent [**weight** *weight*] [**cir-weight** *cir-weight*]

no parent

Context

[Tree] (config>service>vprn>if>sap>egress>sched-override>scheduler parent)

[Tree] (config>service>vprn>if>sap>ingress>sched-override>scheduler parent)

Full Context

configure service vprn interface sap egress scheduler-override scheduler parent

configure service vprn interface sap ingress scheduler-override scheduler parent

Description

This command can be used to override the scheduler's parent weight and cir-weight information. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site.

The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy – this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non default weightings for fostered schedulers.

The no form of this command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.

Default

no parent

Parameters

weight *weight*

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict **level** defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the policer, queue, or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.

A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.

Values 0 to 100

cir-weight *cir-weight*

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same *cir-level* defined by the **cir-level** parameter in the applied scheduler policy. Within the strict **cir-level**, all **cir-weight** values from active children at that level are summed and the ratio of each active child's **cir-weight** to the total is used to distribute the available bandwidth at that level. A **cir-weight** is considered to be active when the policer, queue, or scheduler that the **cir-weight** pertains to has not reached the CIR and still has packets to transmit.

A 0 (zero) **cir-weight** value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.

Values 0 to 100

Platforms

7705 SAR Gen 2

parent

Syntax

parent {**root** | *arbiter-name*} [**level** *priority-level*] [**weight** *weight-within-level*]

no parent

Context

[Tree] (config>qos>plcr-ctrl-plcy>tier>arbiter parent)

Full Context

configure qos policer-control-policy tier arbiter parent

Description

This command is used to define from where the tiered arbiter receives bandwidth. Both tier 1 and tier 2 arbiters default to parenting to the root arbiter. Tier 2 arbiters may be modified to parent to a tier 1 arbiter. The tier 1 arbiter parent cannot be changed.

The **no** form of this command is used to return the tiered arbiter to the default parenting behavior.

Default

parent root level 1 weight 1

Parameters

root

In tier 1, *arbiter-name* is not allowed and only **root** is accepted. When **root** is specified, the arbiter will receive all bandwidth directly from the root arbiter. This is the default parent for tiered arbiters.

arbiter-name

In tier 1, *arbiter-name* is not allowed and only **root** is accepted. The specified *arbiter-name* must exist within the policer-control-policy at tier 1 or the parent command will fail. When a tiered arbiter is acting as a parent for another tiered arbiter, the parent arbiter cannot be removed from the policy. The child arbiter will receive all bandwidth directly from its parent arbiter (that receives bandwidth from the root arbiter).

priority-level

Each child arbiter attaches to its parent on one of the parent's eight strict levels. Level 1 is the lowest and 8 is the highest. The level attribute is used to define which level the child arbiter uses on its parent. The parent distributes its available bandwidth based on strict priority starting with priority level 8 and proceeding towards level 1.

Values 1 to 8

Default 1

weight-within-level

The **weight** attribute is used to define how multiple children at the same parent strict level compete when insufficient bandwidth exists on the parent for that level. Each child's weight is divided by the sum of the active children's weights and the result is multiplied by the available bandwidth. If a child cannot receive its entire weighted fair share of bandwidth due to a defined child rate limit, the remainder of its bandwidth is distributed between the other children based on their weights.

Values 1 to 100

Default 1

Platforms

7705 SAR Gen 2

parent

Syntax

parent *arbiter-name* [**weight** *weight-within-level*] [**level** *level*]

no parent

Context

[Tree] (config>qos>sap-ingress>policer parent)

[Tree] (config>qos>sap-egress>policer parent)

Full Context

configure qos sap-ingress policer parent

configure qos sap-egress policer parent

Description

This command is used to create a child-to-parent mapping between each instance of the policer and either the **root** arbiter or a specific tiered **arbiter** on the object where the policy is applied. Defining a parent association for the policer causes the policer to compete for the parent policer's available bandwidth with other child policers mapped to the policer control hierarchy.

Policer control hierarchies may be created on SAPs or on a subscriber or multiservice site context. To create a policer control hierarchy on an ingress or egress SAP context, a **policer-control-policy** must be applied to the SAP. When applied, the system will create a parent policer that is bandwidth limited by the policy's **max-rate** value under the root arbiter. The root arbiter in the policy also provides the information used to determine the various priority-level discard-unfair and discard-all thresholds. Besides the root arbiter, the policy may also contain user-defined tiered arbiters that provide arbitrary bandwidth control for subsets of child policers that are either directly or indirectly parented by the arbiter.

When the QoS policy containing the policer with a **parent** mapping to an arbiter name exists on the SAP, the system will scan the available arbiters on the SAP. If an arbiter exists with the appropriate name, the policer to arbiter association is created. If the specified arbiter does not exist either because a **policer-control-policy** is not currently applied to the SAP or the arbiter name does not exist within the applied

policy, the policer is placed in an orphan state. Orphan policers operate as if they are not parented and are not subject to any bandwidth constraints other than their own PIR. When a policer enters the orphan state, it is flagged as operationally degraded due to the fact that it is not operating as intended and a trap is generated. Whenever a **policer-control-policy** is added to the SAP or the existing policy is modified, the SAP's policer's parenting configurations must be reevaluated. If an orphan policer becomes parented, the degraded flag is cleared, and a resulting trap is generated.

For subscribers, the policer control hierarchy is created through the **policer-control-policy** applied to the **sub-profile** used by the subscriber. A unique policer control hierarchy is created for each subscriber associated with the **sub-profile**. The QoS policy containing the policer with the parenting command comes into play through the subscriber **sla-profile**, which references the QoS policy. The combining of the **sub-profile** and the **sla-profile** at the subscriber level provides the system with the proper information to create the policer control hierarchy instance for the subscriber.

Executing the **parent** command will fail if:

- The policer’s stat-mode in the QoS policy is set to no-stats
- A stat-mode no-stats override exists on an instance of the policer on a SAP or subscriber or multiservice site context

A policer with a **parent** command applied cannot be configured with **stat-mode no-stats** in either the QoS policy or as an override on an instance of the policer.

When a policer is successfully parented to an arbiter, the **parent** commands **level** and **weight** parameters are used to determine at what priority level and at which weight in the priority level that the child policer competes with other children (policers or other arbiters) for bandwidth.

The **no** form of this command is used to remove the parent association from all instances of the policer.

Parameters

{root | *arbiter-name*}

When the **parent** command is executed, either the keyword **root** or an *arbiter-name* must be specified.

root

Specifies that the policer is intended to become a child to the **root** arbiter where an instance of the policer is created. If the **root** arbiter does not exist, the policer will be placed in the orphan state.

Default **root**

arbiter-name

Specifies that the policer is intended to become a child to one of the tiered arbiters with the given arbiter-name where an instance of the policer is created. If the specified *arbiter-name* does not exist, the policer will be placed in the orphan state.

weight *weight-within-level*

The **weight** *weight-within-level* keyword and parameter are optional when executing the **parent** command. When **weight** is not specified, a default level of 1 is used in the parent arbiter’s priority level. When **weight** is specified, the *weight-within-level* parameter must be specified as an integer value from 1 through 100.

Default 1

Platforms

7705 SAR Gen 2

parent

Syntax

parent *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]

no parent

Context

[Tree] (config>qos>sap-ingress>queue parent)

[Tree] (config>qos>sap-egress>queue parent)

Full Context

configure qos sap-ingress queue parent

configure qos sap-egress queue parent

Description

This command defines an optional parent scheduler that further governs the available bandwidth given the queue aside from the queue's PIR setting. When multiple schedulers, policers (at egress only), and/or queues share a child status with the parent scheduler, the **weight** or **level** parameters define how this queue contends with the other children for the parent's bandwidth.

Checks are not performed to see if a *scheduler-name* exists when the parent command is defined on the queue. Scheduler names are configured in the **config>qos>scheduler-policy>tier** *level* context. Multiple schedulers can exist with the *scheduler-name* and the association pertains to a scheduler that should exist on the egress SAP as the policy is applied and the queue created. When the queue is created on the egress SAP, the existence of the *scheduler-name* is dependent on a scheduler policy containing the *scheduler-name* being directly or indirectly applied (through a multiservice customer site) to the egress SAP. If the *scheduler-name* does not exist, the queue is placed in the orphaned operational state. The queue will accept packets but will not be bandwidth limited by a virtual scheduler or the scheduler hierarchy applied to the SAP. The SAP that the queue belongs to also depicts an orphan queue status. The orphaned state of the queue is automatically cleared when the *scheduler-name* becomes available on the egress SAP.

The parent scheduler can be made unavailable due to the removal of a scheduler policy or scheduler. When an existing parent scheduler is removed or inoperative, the queue enters the orphaned state and automatically returns to normal operation when the parent scheduler is available again.

When a parent scheduler is defined without specifying the weight parameter, the default is a weight of 1.

The **no** form of this command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. When a parent association has been removed, the former child queue attempts to operate based on its configured rate parameter. Removing the parent association on the queue within the policy takes effect immediately on all queues using the SAP egress QoS policy.

Parameters

scheduler-name

The defined *scheduler-name* conforms to the same input criteria as the schedulers defined within a scheduler policy. Scheduler names are configured in the config>qos>scheduler-policy>tier *level* context. There are no checks performed at the time of definition to ensure that the *scheduler-name* exists within an existing scheduler policy. For the queue to use the defined *scheduler-name*, the scheduler exists on each egress SAP that the queue is eventually created on. For the duration where *scheduler-name* does not exist on the egress SAP, the queue operates in an orphaned state.

Values Any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

weight

Specifies the relative weight of this queue in comparison to other child schedulers, policers, and queues, while vying for bandwidth on the parent *scheduler-name*. Any queues, policers, or schedulers defined as weighted receive no parental bandwidth until all policers, queues, and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

All **weight** values from all weighted active queues, policers, and schedulers with a common parent scheduler are added together. Then, each individual active weight is divided by the total, deriving the percentage of remaining bandwidth provided to the queue, policer, or scheduler. A weight is considered to be active when the pertaining queue, policer, or scheduler has not reached its maximum rate and still has packets to transmit. All child queues, policers, and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all non-zero weighted queues, policers, and schedulers at that level are operating at the maximum bandwidth or are idle.

Values 0 to 100

Default 1

level

The optional **level** parameter defines the level of hierarchy when compared to other schedulers and queues competing for bandwidth on the parent *scheduler-name*. Queues or schedulers will not receive parental bandwidth until all queues, policers, and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

Children of the parent scheduler with a lower strict priority will not receive bandwidth until all children with a higher strict priority have either reached their maximum bandwidth or are idle. Children with the same strict level are serviced in relation to their relative weights.

Values 1 to 8

Default 1

cir-weight

Specifies the weight that the queue or scheduler uses at the within-CIR port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to

100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-CIR pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 100

Default 1

cir-level

Specifies the port priority that the queue or scheduler will use to receive bandwidth for its within-CIR offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-CIR pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 8 (8 is the highest priority)

Default 0

Platforms

7705 SAR Gen 2

parent

Syntax

parent {**root** | *arbiter-name*} [**level** *level*] [**weight** *weight-within-level*]

no parent

Context

[Tree] (config>qos>qgrps>ing>qgrp>policer parent)

[Tree] (config>qos>qgrps>egr>qgrp>policer parent)

Full Context

configure qos queue-group-templates ingress queue-group policer parent

configure qos queue-group-templates egress queue-group policer parent

Description

This command is used to create a child-to-parent mapping between each instance of the policer and either the **root** arbiter or a specific tiered **arbiter** on the object where the policy is applied. Defining a parent association for the policer causes the policer to compete for the parent policer's available bandwidth with other child policers mapped to the policer control hierarchy.

Policer control hierarchies may be created on SAPs or on a subscriber or multiservice site context. To create a policer control hierarchy on an ingress or egress SAP context, a **policer-control-policy** must be applied to the SAP. When applied, the system will create a parent policer that is bandwidth limited by the policy's **max-rate** value under the root arbiter. The root arbiter in the policy also provides the information

used to determine the various priority level discard-unfair and discard-all thresholds. Besides the root arbiter, the policy may also contain user-defined tiered arbiters that provide arbitrary bandwidth control for subsets of child policers that are either directly or indirectly parented by the arbiter.

When the QoS policy containing the policer with a **parent** mapping to an arbiter name exists on the SAP, the system will scan the available arbiters on the SAP. If an arbiter exists with the appropriate name, the policer to arbiter association is created. If the specified arbiter does not exist either because a **policer-control-policy** is not currently applied to the SAP or the arbiter name does not exist within the applied policy, the policer is placed in an orphan state. Orphan policers operate as if they are not parented and are not subject to any bandwidth constraints other than their own PIR. When a policer enters the orphan state, it is flagged as operationally degraded due to the fact that it is not operating as intended and a trap is generated. Whenever a **policer-control-policy** is added to the SAP or the existing policy is modified, the SAP's policer's parenting configurations must be reevaluated. If an orphan policer becomes parented, the degraded flag is cleared and a resulting trap is generated.

For subscribers, the policer control hierarchy is created through the **policer-control-policy** applied to the **sub-profile** used by the subscriber. A unique policer control hierarchy is created for each subscriber associated with the **sub-profile**. The QoS policy containing the policer with the parenting command comes into play through the subscriber **sla-profile** that references the QoS policy. The combining of the **sub-profile** and the **sla-profile** at the subscriber level provides the system with the proper information to create the policer control hierarchy instance for the subscriber.

Executing the **parent** command will fail if:

- The policer’s stat-mode in the QoS policy is set to no-stats
- A stat-mode no-stats override exists on an instance of the policer on a SAP or subscriber or multiservice site context

A policer with a **parent** command applied cannot be configured with **stat-mode no-stats** in either the QoS policy or as an override on an instance of the policer.

When a policer is successfully parented to an arbiter, the **parent** commands **level** and **weight** parameters are used to determine at what priority level and at which weight in the priority level that the child policer competes with other children (policers or other arbiters) for bandwidth.

The **no** form of this command is used to remove the parent association from all instances of the policer.

Parameters

{root | *arbiter-name*}

When the **parent** command is executed, either the keyword **root** or an *arbiter-name* must be specified.

Default **root**

root

The **root** keyword specifies that the policer is intended to become a child to the **root** arbiter where an instance of the policer is created. If the **root** arbiter does not exist, the policer will be placed in the orphan state.

arbiter-name

The *arbiter-name* parameter specifies that the policer is intended to become a child to one of the tiered arbiters with the given arbiter-name where an instance of the policer is created. If the specified arbiter-name does not exist, the policer will be placed in the orphan state.

weight *weight-within-level*

The **weight** *weight-within-level* keyword and parameter are optional when executing the **parent** command. When **weight** is not specified, a default level of 1 is used in the parent arbiters priority level. When **weight** is specified, the *weight-within-level* parameter must be specified as an integer value from 1 through 100.

Default 1

Platforms

7705 SAR Gen 2

parent

Syntax

parent *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]

no parent

Context

[Tree] (config>qos>qgrps>egr>qgrp>queue parent)

[Tree] (config>qos>qgrps>ing>qgrp>queue parent)

Full Context

configure qos queue-group-templates egress queue-group queue parent

configure qos queue-group-templates ingress queue-group queue parent

Description

This command defines an optional parent scheduler that further governs the available bandwidth given the queue aside from the queue's PIR setting. When multiple schedulers, policers (at egress only), and/or queues share a child status with the parent scheduler, the **weight** or **level** parameters define how this queue contends with the other children for the parent's bandwidth.

Checks are not performed to see if a *scheduler-name* exists when the parent command is defined on the queue. Scheduler names are configured in the config>qos>scheduler-policy>tier *level* context. Multiple schedulers can exist with the *scheduler-name* and the association pertains to a scheduler that should exist on the egress SAP as the policy is applied and the queue created. When the queue is created on the egress SAP, the existence of the *scheduler-name* is dependent on a scheduler policy containing the *scheduler-name* being directly or indirectly applied (through a multiservice customer site) to the egress SAP. If the *scheduler-name* does not exist, the queue is placed in the orphaned operational state. The queue will accept packets but will not be bandwidth limited by a virtual scheduler or the scheduler hierarchy applied to the SAP. The SAP that the queue belongs to must also depict an orphan queue status. The orphaned state of the queue is automatically cleared when the *scheduler-name* becomes available on the egress SAP.

The parent scheduler can be made unavailable due to the removal of a scheduler policy or scheduler. When an existing parent scheduler is removed or inoperative, the queue enters the orphaned state and automatically returns to normal operation when the parent scheduler is available again.

When a parent scheduler is defined without specifying weight or strict parameters, the default bandwidth access method is weight with a value of 1.

The **no** form of this command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. When a parent association has been removed, the former child queue attempts to operate based on its configured rate parameter. Removing the parent association on the queue within the policy takes effect immediately on all queues using the SAP egress QoS policy.

Parameters

scheduler-name

The defined *scheduler-name* conforms to the same input criteria as the schedulers defined within a scheduler policy. Scheduler names are configured in the config>qos>scheduler-policy>tier *level* context. There are no checks performed at the time of definition to ensure that the *scheduler-name* exists within an existing scheduler policy. For the queue to use the defined *scheduler-name*, the scheduler exists on each egress SAP the queue is eventually created on. For the duration where *scheduler-name* does not exist on the egress SAP, the queue operates in an orphaned state.

Values Any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

weight *weight*

weight defines the relative weight of this queue in comparison to other child schedulers and queues while vying for bandwidth on the parent *scheduler-name*. Any queues or schedulers defined as weighted receive no parental bandwidth until all policers, queues, and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

All **weight** values from all weighted active queues, policers, and schedulers with a common parent scheduler are added together. Then, each individual active weight is divided by the total, deriving the percentage of remaining bandwidth provided to the queue, policer, or scheduler. A weight is considered to be active when the pertaining queue or scheduler has not reached its maximum rate and still has packets to transmit. All child policers, queues, and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all non-zero weighted queues, policers, and schedulers at that level are operating at the maximum bandwidth or are idle.

Values 0 to 100

Default 1

level *level*

The optional **level** parameter defines the level of hierarchy when compared to other schedulers and queues when vying for bandwidth on the parent *scheduler-name*. Queues or schedulers will not receive parental bandwidth until all queues and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

Children of the parent scheduler with a lower strict priority will not receive bandwidth until all children with a higher strict priority have either reached their maximum bandwidth or are idle. Children with the same strict level are serviced relative to their weights.

Values 1 to 8

Default 1

cir-weight ***cir-weight***

Specifies the weight the queue or scheduler will use at the within-CIR port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-CIR pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 100

Default 1

cir-level ***cir-level***

Specifies the port priority the queue or scheduler will use to receive bandwidth for its within-CIR offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-CIR pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 8 (8 is the highest priority)

Default 0

Platforms

7705 SAR Gen 2

parent

Syntax

parent *scheduler-name* [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]

no parent

Context

[Tree] (config>qos>scheduler-policy>tier>scheduler parent)

Full Context

configure qos scheduler-policy tier scheduler parent

Description

This command defines an optional parent scheduler that is higher up the policy hierarchy. Only schedulers in tier levels 2 and 3 can have a parental association. When multiple schedulers, policers (at egress only), and/or queues share a child status with the scheduler on the parent, the weight or strict parameters define

how this scheduler contends with the other children for the parent's bandwidth. The parent scheduler can be removed or changed at any time and is immediately reflected on the schedulers created by association of this scheduler policy.

When a parent scheduler is defined without specifying weight or strict parameters, the default bandwidth access method is weight with a value of 1.

The **no** form of this command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. When a parent association has been removed, the former child scheduler attempts to operate based on its configured rate parameter. Removing the parent association on the scheduler within the policy will take effect immediately on all schedulers with *scheduler-name* that have been created using the *scheduler-policy-name*.

Parameters

scheduler-name

Specifies a scheduler name. The *scheduler-name* must already exist within the context of the scheduler policy in a tier that is higher (numerically lower).

Values Any valid *scheduler-name* existing on a higher tier within the scheduler policy.

weight *weight*

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict level defined by the **level** parameter. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the queue or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.

A zero (0) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.

Values 0 to 100

Default 1

level *level*

Specifies the strict priority level of this scheduler in comparison to other child schedulers and queues vying for bandwidth on the parent scheduler-name during the above-CIR distribution phase of bandwidth allocation. During the above-CIR distribution phase, any queues or schedulers defined at a lower strict level receive no parental bandwidth until all queues and schedulers defined with a higher (numerically larger) strict level on the parent have reached their maximum bandwidth or have satisfied their offered load requirements.

When the similar **cir-level** parameter default (undefined) are retained for the child scheduler, bandwidth is only allocated to the scheduler during the above-CIR distribution phase.

Children of the parent scheduler with a lower strict priority level will not receive bandwidth until all children with a higher strict priority level have either reached their maximum bandwidth or are idle. Children with the same strict level are serviced in relation to their relative weights.

Values 1 to 8

Default 1

cir-weight ***cir-weight***

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same **cir-level** defined by the **cir-level** parameter. Within the strict **cir-level**, all **cir-weight** values from active children at that level are summed and the ratio of each active child's **cir-weight** to the total is used to distribute the available bandwidth at that level. A **cir-weight** is considered to be active when the queue or scheduler that the **cir-weight** pertains to has not reached the CIR and still has packets to transmit.

A zero (0) **cir-weight** value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.

Values 0 to 100

Default 1

cir-level ***cir-level***

Specifies the strict priority CIR level of this scheduler in comparison to other child schedulers and queues vying for bandwidth on the parent *scheduler-name* during the within-CIR distribution phase of bandwidth allocation. During the within-CIR distribution phase, any queues or schedulers defined at a lower strict CIR level receive no parental bandwidth until all queues and schedulers defined with a higher (numerically larger) strict CIR level on the parent have reached their CIR bandwidth or have satisfied their offered load requirements.

If the scheduler's **cir-level** parameter retains the default (undefined) state, bandwidth is only allocated to the scheduler during the above-CIR distribution phase.

Children with the same strict cir-level are serviced according to their cir-weight.

Values 0 to 8

Default 0

Platforms

7705 SAR Gen 2

parent

Syntax

parent [**weight** *weight*] [**cir-weight** *cir-weight*]

no parent

Context

[Tree] (config>service>cust>multi-service-site>egress>sched-override>scheduler parent)

[Tree] (config>service>cust>multi-service-site>ingress>sched-override>scheduler parent)

Full Context

configure service customer multi-service-site egress scheduler-override scheduler parent

configure service customer multi-service-site ingress scheduler-override scheduler parent

Description

This command overrides the scheduler's parent weight and CIR weight information. The weights apply to the associated level or cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site.

The override weights are ignored if the scheduler does not have a **parent** command configured in the scheduler policy. This allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non-default weightings for fostered schedulers.

The **no** form of the command returns the scheduler's parent weight and CIR weight to the value configured in the applied scheduler policy.

Default

no parent

Parameters

weight

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict **level** defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the policer, queue or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit. A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.

Values 0 to 100

Default 1

cir-weight

Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same *cir-level* defined by the **cir-level** parameter in the applied scheduler policy. Within the strict **cir-level**, all **cir-weight** values from active children at that level are summed and the ratio of each active child's **cir-weight** to the total is used to distribute the available bandwidth at that level. A **cir-weight** is considered to be active when the policer, queue or scheduler that the **cir-weight** pertains to has not reached the CIR and still has packets to transmit. A 0 (zero) **cir-weight** value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.

Values 0 to 100

Default 0

Platforms

7705 SAR Gen 2

21.15 parent-location

parent-location

Syntax

parent-location {**none** | **sub** | **vport**}

no parent-location

Context

[Tree] (config>qos>scheduler-policy>tier parent-location)

Full Context

configure qos scheduler-policy tier parent-location

Description

This command determines the expected location of the parent schedulers for the tier 1 schedulers configured with a parent command within the scheduler-policy. The parent schedulers must be configured within a scheduler-policy applied at the location corresponding to the parent-location parameter.

If a parent scheduler name does not exist at the specified location, the schedulers will not be parented and will be orphaned.

The configuration of **parent-location** and **frame-based-accounting** in a scheduler policy is mutually exclusive in order to ensure consistency between the different scheduling levels.

The **no** form of this command reverts to the default.

Default

parent-location none

Parameters

none

This parameter indicates that the tier 1 schedulers do not have a parent scheduler and the configuration of the parent under a tier 1 scheduler is blocked. Conversely, this parameter is blocked when any tier 1 scheduler has a parent configured.

sub

When the scheduler-policy is applied to an sla-profile for a subscriber, the parent schedulers of the tier 1 schedulers need to be configured in the scheduler-policy applied to the subscriber's sub-profile.

If this parameter is configured within a scheduler-policy that is applied to any object except for the egress of an sla-profile, the configured parent schedulers will not be found and so the tier 1 schedulers will not be parented and will be orphaned.

vport

When the scheduler-policy is applied to an sla-profile, a sub-profile for a subscriber, or to the egress of a pseudowire SAP, the parent schedulers of the tier 1 schedulers need to be configured in the scheduler-policy applied to the Vport to which the subscriber will be assigned.

If this parameter is configured within a scheduler-policy that is applied to any object except for the egress of an sla-profile or sub-profile, or to the egress of a PW SAP, the configured parent schedulers will not be found and so the tier 1 schedulers will not be parented and will be orphaned.

Platforms

7705 SAR Gen 2

21.16 participate

participate

Syntax

[**no**] **participate**

Context

[Tree] (config>router>isis>flex-algos>flex-algo participate)

Full Context

configure router isis flexible-algorithms flex-algo participate

Description

This command enables IS-IS participation in a specific flexible algorithm.

The router advertises its capability to participate in a specific flexible algorithm within the IS-IS router-capability TLV. Router participation in a flexible algorithm assumes that segment routing and, consequently the **advertise-router-capability area** is enabled. However, a router only advertises flexible algorithm participation when it can support the corresponding winning flexible algorithm definition. The flexible algorithm participation is not enabled by default.

The **no** form of this command disables participation for a particular flexible algorithm.

Default

no participate

Platforms

7705 SAR Gen 2

participate

Syntax

[**no**] **participate**

Context

[Tree] (config>router>ospf>flex-algos>flex-algo participate)

Full Context

configure router ospf flexible-algorithms flex-algo participate

Description

This command enables OSPFv2 participation in a specific flexible algorithm.

The router advertises its capability to participate in a specific flexible algorithm within the OSPFv2 SR algorithm TLV of the router information opaque LSA. Router participation in a flexible algorithm assumes that segment routing and, consequently, the **advertise-router-capability area** is enabled. However, a router only advertises flexible algorithm participation when it can support the corresponding winning flexible algorithm definition. The flexible algorithm participation is not enabled by default.

The **no** form of this command disables participation for a specific flexible algorithm.

Default

no participate

Platforms

7705 SAR Gen 2

21.17 partner-down-delay

partner-down-delay

Syntax

partner-down-delay [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no partner-down-delay

Context

[Tree] (config>router>dhcp>server>failover partner-down-delay)

[Tree] (config>router>dhcp6>server>pool>failover partner-down-delay)

[Tree] (config>service>vprn>dhcp6>server>failover partner-down-delay)

[Tree] (config>router>dhcp>server>pool>failover partner-down-delay)

[Tree] (config>service>vprn>dhcp6>server>pool>failover partner-down-delay)

[Tree] (config>service>vprn>dhcp>server>pool>failover partner-down-delay)

[Tree] (config>service>vprn>dhcp>server>failover partner-down-delay)

[Tree] (config>router>dhcp6>server>failover partner-down-delay)

Full Context

configure router dhcp local-dhcp-server failover partner-down-delay

configure router dhcp6 local-dhcp-server pool failover partner-down-delay

configure service vprn dhcp6 local-dhcp-server failover partner-down-delay

configure router dhcp local-dhcp-server pool failover partner-down-delay

configure service vprn dhcp6 local-dhcp-server pool failover partner-down-delay

configure service vprn dhcp local-dhcp-server pool failover partner-down-delay

configure service vprn dhcp local-dhcp-server failover partner-down-delay

configure router dhcp6 local-dhcp-server failover partner-down-delay

Description

This command configures the partner down delay time. Since the DHCP lease synchronization failure can be caused by the failure of the intercommunication link (and not necessary the entire node), there is a possibility the redundant DHCP servers become isolated in the network. In other words, they can serve DHCP clients but they cannot synchronize the lease. This can lead to duplicate assignment of IP addresses, since the servers have configured overlapping IP address ranges but they are not aware of each other's leases.

The purpose of the partner down delay is to prevent the IP lease duplication during the intercommunication link failure by not allowing new IP addresses to be assigned from the remote IP address range. This timer is intended to provide the operator with enough time to remedy the failed situation and to avoid duplication of IP addresses or prefixes during the failure.

During the partner-down-delay time, the prefix designated as remote is eligible only for renewals of the existing DHCP leases that have been synchronized by the peering node. Only after the sum of the partner-down-delay and the maximum-client-lead-time will the prefix designated as remote be eligible for delegation of the new DHCP leases. When this occurs, we say that the remote IP address range has been taken over.

It is possible to expedite the takeover of a remote IP address range so that the new IP leases can start being delegated from that range shortly after the intercommunication failure is detected. This can be achieved by configuring the partner-down-delay timer to 0 seconds, along with enabling the ignore-mclt-on-takeover CLI flag. Caution must be taken before enabling this functionality. It is safe to bypass safety timers (partner-down-delay + MCLT) only in cases where the operator is certain that the intercommunication between the nodes has failed due to the entire node failure and not due to the intercommunication (MCS) link failure. Failed intercommunication due to the nodal failure would ensure that only one node is present in the network for IP address delegation (as opposed to two isolated nodes with overlapping IP address ranges where address duplication can occur). For this reason, the operator must ensure that there are redundant paths between the nodes to ensure uninterrupted synchronization of DHCP leases.

In access-driven mode of operation, partner-down-delay has no effect.

The **no** form of this command reverts to the default.

Default

partner-down-delay hrs 23 min 59 sec 59

Parameters

partner-down-delay

Specifies the partner down delay time.

Values	**hrs** *hours*	1 to 23
	min *minutes*	1 to 59
	sec *seconds*	0 to 59

Platforms

7705 SAR Gen 2

21.18 passive

passive

Syntax

[**no**] **passive**

Context

[Tree] (config>service>vprn>bgp>group>neighbor passive)

[Tree] (config>service>vprn>bgp>group passive)

Full Context

configure service vprn bgp group neighbor passive

configure service vprn bgp group passive

Description

This command enables passive mode for the BGP group or neighbor.

When in passive mode, BGP will not attempt to actively connect to the configured BGP peers but responds only when it receives a connect open request from the peer.

The **no** form of this command used at the group level disables passive mode where BGP actively attempts to connect to its peers.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no passive — BGP will actively try to connect to all the configured peers.

Platforms

7705 SAR Gen 2

passive

Syntax

[**no**] **passive**

Context

[Tree] (config>service>vprn>isis>if passive)

[Tree] (config>service>vprn>isis>if>level passive)

Full Context

configure service vprn isis interface passive

configure service vprn isis interface level passive

Description

This command adds the passive attribute which causes the interface to be advertised as an IS-IS interface without running the IS-IS protocol. Normally, only interface addresses that are configured for IS-IS are advertised as IS-IS interfaces at the level that they are configured.

When the passive mode is enabled, the interface or the interface at the level ignores ingress IS-IS protocol PDUs and does not transmit IS-IS protocol PDUs.

The **no** form of this command removes the passive attribute.

Default

no passive

Platforms

7705 SAR Gen 2

passive

Syntax

[**no**] **passive**

Context

[Tree] (config>service>vprn>ospf3>area>if passive)

[Tree] (config>service>vprn>ospf>area>if passive)

Full Context

configure service vprn ospf3 area interface passive

configure service vprn ospf area interface passive

Description

This command adds the passive property to the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol.

By default, only interface addresses that are configured for OSPF are advertised as OSPF interfaces. The **passive** parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol.

While in passive mode, the interface ignores ingress OSPF protocol packets and does not transmit any OSPF protocol packets.

The **no** form of this command removes the passive property from the OSPF interface.

Default

no passive

Platforms

7705 SAR Gen 2

passive

Syntax

[**no**] **passive**

Context

[Tree] (config>router>bgp>group passive)

[Tree] (config>router>bgp>group>neighbor passive)

Full Context

configure router bgp group passive

configure router bgp group neighbor passive

Description

Enables/disables passive mode for the BGP group or neighbor**.**

When in passive mode, BGP will not attempt to actively connect to the configured BGP peers but responds only when it receives a connect open request from the peer.

The **no** form of this command used at the group level disables passive mode where BGP actively attempts to connect to its peers.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no passive

Platforms

7705 SAR Gen 2

passive

Syntax

[**no**] **passive**

Context

[Tree] (config>router>isis>if passive)

[Tree] (config>router>isis>if>level passive)

Full Context

configure router isis interface passive

configure router isis interface level passive

Description

This command adds the passive attribute which causes the interface to be advertised as an IS-IS interface without running the IS-IS protocol. Normally, only interface addresses that are configured for IS-IS are advertised as IS-IS interfaces at the level that they are configured.

When the passive mode is enabled, the interface or the interface at the level ignores ingress IS-IS protocol PDUs and does not transmit IS-IS protocol PDUs.

The **no** form of this command removes the passive attribute.

Default

no passive

Platforms

7705 SAR Gen 2

passive

Syntax

[**no**] **passive**

Context

[Tree] (config>router>ospf>area>interface passive)

[Tree] (config>router>ospf3>area>interface passive)

Full Context

configure router ospf area interface passive

configure router ospf3 area interface passive

Description

This command adds the passive property to the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol.

By default, only interface addresses that are configured for OSPF will be advertised as OSPF interfaces. The **passive** parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol.

While in passive mode, the interface will ignore ingress OSPF protocol packets and not transmit any OSPF protocol packets.

The **no** form of this command removes the passive property from the OSPF interface.

Default

no passive

Platforms

7705 SAR Gen 2

21.19 passive-mode

passive-mode

Syntax

[**no**] **passive-mode**

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ep passive-mode)

Full Context

configure redundancy multi-chassis peer mc-endpoint passive-mode

Description

This command configures the passive mode behavior for the MC-EP protocol. When in passive mode the MC-EP pair will be dormant until two of the pseudowires in a MC-EP will be signaled as active by the remote PEs, being assumed that the remote pair is configured with regular MC-EP. As soon as more than one pseudowire is active, dormant MC-EP pair will activate. It will use the regular exchange to select the best pseudowire between the active ones and it will block the Rx and Tx directions of the other pseudowires.

The **no** form of this command will disable the passive mode behavior.

Default

no passive-mode

Platforms

7705 SAR Gen 2

21.20 password

password

Syntax

password

Context

[Tree] (password)

Full Context

password

Description

This operational command changes the local user password.

This command is automatically invoked when a user logs in after the administrator uses the **new-password-at-login** command to force a new password, or the password has expired (**aging**). At this time, the user is prompted to enter the old password, new password, and then the new password again to verify the input.

If the user fails to create a new password, CLI access is denied.

A user cannot configure a nonconforming password using the global **password** command. In this case, the CLI displays an error message and the password change fails. To configure a password value that does not conform to the minimum length or other password complexity rules, use the **config>system>security>user>password** command (for example, executed by an administrator).

Platforms

7705 SAR Gen 2

password

Syntax

password

Context

[Tree] (config>system>security password)

Full Context

configure system security password

Description

Commands in this context configure password-related parameters.

Platforms

7705 SAR Gen 2

password

Syntax

password *password* [**hash** | **hash2** | **custom**]

no password

Context

[Tree] (config>ipsec>rad-auth-plcy password)

Full Context

configure ipsec radius-authentication-policy password

Description

This command specifies the password that is used in the RADIUS access requests.

The **no** form of this command resets the password to its default of **ALU** and will be stored using hash/hash2 encryption.

Default

no password

Parameters

password

Specifies a password string up to 64characters.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

password

Syntax

password [*password*]

Context

[Tree] (config>system>security>user password)

Full Context

configure system security user password

Description

This command configures the user password for console and FTP access.

The password is stored in an encrypted format in the configuration file when specified. Passwords should be encased in double quotes (" ") at the time of the password creation. The double quote character (") is not accepted inside a password. It is interpreted as the start or stop delimiter of a string.

The password can be entered as plain text or a hashed value. SR OS can distinguish between hashed passwords and plain text passwords and take the appropriate action to store the password correctly.

```
config>system>security>user# password testuser1
```

The password is hashed by default.

For example:

```
config>system>security# user testuser1
config>system>security>user$ password xyzabcd1
config>system>security>user# exit
```

```
config>system>security# info
----------------------------------
...
            user "testuser1"
                password "$2y$10$pFoehOg/tCbBMPDJ/
kqpu.8af0AoVGY2xsR7WFqyn5fVTnwRzGmOK"
            exit
...
----------------------------------
config>system>security#
```

The **password** command allows you also to enter the password as a hashed value.

For example:

```
config>system>security# user testuser1
config>system>security>user$ password "$2y$10$pFoehOg/tCbBMPDJ/
kqpu.8af0AoVGY2xsR7WFqyn5fVTnwRzGmOK"
config>system>security>user# exit
config>system>security# info
----------------------------------
...
user "testuser1"
password "$2y$10$pFoehOg/tCbBMPDJ/kqpu.8af0AoVGY2xsR7WFqyn5fVTnwRzGmOK"
exit
...
----------------------------------
config>system>security#
```

Parameters

password

This is the password for the user that must be entered by this user during the login procedure. The minimum length of the password is determined by the **minimum-length** command. The maximum length can be up to 20 chars if unhashed, 32 characters if hashed. The complexity requirements for the password is determined by the **complexity-rules** command and must be followed; otherwise, the password will not be accepted.

All password special characters (#, ?, space) must be enclosed within double quotes.

For example: config>system>security>user# password "south#bay?”

The question mark character (?) cannot be directly inserted as input during a telnet connection because the character is bound to the **help** command during a normal Telnet/ console connection.

To insert a # or ? characters, they must be entered inside a notepad or clipboard program and then cut and pasted into the Telnet session in the password field that is encased in the double quotes as delimiters for the password.

If a **password** is entered without any parameters, a password length of zero is implied: (carriage return).

Platforms

7705 SAR Gen 2

password

Syntax

password *password* [**hash** | **hash2** | **custom**]

no password

Context

[Tree] (bof password)

Full Context

bof password

Description

This command configures the password to access the BOF interactive menu at startup.

If a password is configured, the BOF interactive menu is accessible only when the correct password is entered. If the correct password is not entered in 30 s, the node reboots.

The **no** form of this command removes the configured password.

Default

no password

Parameters

password

Specifies the password.

If the **hash**, **hash2**, or **custom** parameter is not configured, the password is entered in plaintext and the password length must be between 8 and 32 characters. A plaintext password cannot contain embedded nulls or end with " hash", " hash2", or " custom".

If the **hash**, **hash2**, or **custom** parameter is configured, the password is hashed and the password length must be between 1 and 64 characters.

hash

Keyword to specify that the password is entered in an encrypted form.

hash2

Keyword to specify that the password is entered in a more complex encrypted form. The **hash2** encryption scheme is node-specific and the password cannot be transferred between nodes.

custom

Keyword to specify that the password uses custom encryption.

Platforms

7705 SAR Gen 2

21.21 password-history

password-history

Syntax

password-history {**user** *user-name* | **all**}

Context

[Tree] (admin>clear password-history)

Full Context

admin clear password-history

Description

This command is used to clear old passwords used by a specific user, or for all users.

Parameters

user-name

Clears the password history information about the specified user, up to 32 characters.

all

Clears the password history information for all users.

Platforms

7705 SAR Gen 2

21.22 path

path

Syntax

path *name*

no path

Context

[Tree] (config>service>epipe>spoke-sdp-fec path)

Full Context

configure service epipe spoke-sdp-fec path

Description

This command specifies the explicit path, containing a list of S-PE hops, that should be used for this spoke SDP. The path-name should correspond to the name of an explicit path configured in the **config>service>pw-routing** context.

If no path is configured, then each next-hop of the MS-PW used by the spoke SDP will be chosen locally at each T-PE and S-PE.

Default

no path

Parameters

name

The name of the explicit path to be used, as configured under the **config>service>pw-routing** context.

Platforms

7705 SAR Gen 2

path

Syntax

[**no**] **path** *path-name*

Context

[Tree] (config>router>mpls path)

Full Context

configure router mpls path

Description

This command creates the path to be used for an LSP. A path can be used by multiple LSPs. A path can specify some or all hops from ingress to egress and they can be either **strict** or **loose**. A path can also be empty (no *path-name* specified) in which case the LSP is set up based on IGP (best effort) calculated shortest path to the egress router. Paths are created in a **shutdown** state. A path must be shutdown before making any changes (adding or deleting hops) to the path. When a path is shutdown, any LSP using the path becomes operationally down.

To create a strict path from the ingress to the egress router, the ingress and the egress routers must be included in the path statement.

The **no** form of this command deletes the path and all its associated configuration information. All the LSPs that are currently using this path will be affected. Additionally all the services that are actively using these LSPs will be affected. A path must be **shutdown** and unbound from all LSPs using the path before it can be deleted. The **no path** *path-name* command will not result in any action except a warning message on the console indicating that the path may be in use.

Parameters

path-name

Specifies a unique case-sensitive alphanumeric name label for the LSP path up to 32 characters in length.

Platforms

7705 SAR Gen 2

path

Syntax

path [**detail**]

no path

Context

[Tree] (debug>router>rsvp>event path)

Full Context

debug router rsvp event path

Description

This command debugs path-related events.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about path-related events.

Platforms

7705 SAR Gen 2

path

Syntax

path [**detail**]

no path

Context

[Tree] (debug>router>rsvp>packet path)

Full Context

debug router rsvp packet path

Description

This command enables debugging for RSVP path packets.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about path-related events.

Platforms

7705 SAR Gen 2

path

Syntax

path *name* [**create**]

no path *name*

Context

[Tree] (config>service>pw-routing path)

Full Context

configure service pw-routing path

Description

This command configures an explicit path between this T-PE and a remote T-PE. For each path, one or more intermediate S-PE hops must be configured. A path can be used by multiple multi-segment pseudowires. Paths are used by a 7705 SAR Gen 2 T-PE to populate the list of Explicit Route TLVs included in the signaling of a dynamic MS-PW.

A path may specify all or only some of the hops along the route to reach a T-PE.

The **no** form of the command removes a specified explicit path from the configuration.

Parameters

path-name

Specifies a locally-unique case-sensitive alphanumeric name label for the MS-PW path of up to 32 characters in length.

Platforms

7705 SAR Gen 2

path

Syntax

path *path-name* [**create**]

no path *name*

Context

[Tree] (config>system>telemetry>sensor-groups>sensor-group path)

Full Context

configure system telemetry sensor-groups sensor-group path

Description

This command configures a sensor path for the specified sensor-group. Multiple sensor paths can be defined for a single sensor-group. The path is defined in the form of an XML Path (XPath) syntax that refers to single or multiple objects within the YANG model.

The **no** form of the command removes the specified explicit path from the configuration.

Parameters

path-name

Specifies a sensor path, up to 512 characters.

create

Keyword used to create the sensor path.

Platforms

7705 SAR Gen 2

21.23 path-computation-method

path-computation-method

Syntax

path-computation-method *path-computation-method*

no path-computation-method

Context

[Tree] (config>router>mpls>lsp path-computation-method)

[Tree] (config>router>mpls>lsp-template path-computation-method)

Full Context

configure router mpls lsp path-computation-method

configure router mpls lsp-template path-computation-method

Description

This command configures the path computation method of a RSVP-TE or SR-TE LSP.

The user can select among the **hop-to-label** translation, the local CSPF or the PCE for a configured SR-TE LSP. For SR-TE LSP templates, the PCE option is supported with the SR-TE LSP template type **on-demand-p2p-srte** and not other template types.

The user can select among the IGP-based path, the local CSPF, or the PCE for a configured RSVP-TE LSP. The PCE option is not supported with the RSVP-TE LSP template.

By default, the IGP-based path is used for an RSVP-TE LSP and the **hop-to-label** path computation method is used for an SR-TE LSP.

The **no** form of this command returns to the default path computation method for the type of LSP.

Default

no path-computation-method

Parameters

path-computation-method

Specifies the path computation method for the LSP.

Values local-cspf — Selects the local router CSPF for path computation.

pce — Selects the PCE for path computation.

Platforms

7705 SAR Gen 2

21.24 path-cost

path-cost

Syntax

path-cost *sap-path-cost*

no path-cost [*sap-path-cost*]

Context

[Tree] (config>service>vpls>sap>stp path-cost)

[Tree] (config>service>template>vpls-sap-template>stp path-cost)

[Tree] (config>service>vpls>spoke-sdp>stp path-cost)

Full Context

configure service vpls sap stp path-cost

configure service template vpls-sap-template stp path-cost

configure service vpls spoke-sdp stp path-cost

Description

This command configures the Spanning Tree Protocol (STP) path cost for the SAP or spoke-SDP.

The path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that SAP or spoke-SDP. When BPDUs are sent out of other egress SAPs or spoke-SDPs, the newly calculated root path cost is used. These are the values used for CIST when running MSTP.

STP suggests that the path cost is defined as a function of the link bandwidth. Since SAPs and spoke-SDPs are controlled by complex queuing dynamics, in the 7705 SAR Gen 2 the STP path cost is a purely static configuration.

The **no** form of this command returns the path cost to the default value.

Parameters

path-cost

The path cost for the SAP or spoke-SDP

Values 1 to 200000000 (1 is the lowest cost)

Default 10

Platforms

7705 SAR Gen 2

path-cost

Syntax

path-cost *sap-path-cost*

no path-cost

Context

[Tree] (config>service>pw-template>stp path-cost)

Full Context

configure service pw-template stp path-cost

Description

This command configures the Spanning Tree Protocol (STP) path cost for the SAP or spoke SDP.

The path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that SAP or spoke SDP. When BPDUs are sent out other egress SAPs or spoke SDPs, the newly calculated root path cost is used. These are the values used for CIST when running MSTP.

STP suggests that the path cost is defined as a function of the link bandwidth. Since SAPs and spoke SDPs are controlled by complex queuing dynamics, the STP path cost is a purely static configuration.

The **no** form of this command returns the path cost to the default value.

Default

path-cost 10

Parameters

path-cost

Specifies the path cost for the SAP or spoke SDP.

Values 1 to 200000000 (1 is the lowest cost)

Default 10

Platforms

7705 SAR Gen 2

21.25 path-destination

path-destination

Syntax

path-destination *ip-address* **interface** *if-name*

path-destination *ip-address* [**next-hop** *ip-address*]

no path-destination

Context

[Tree] (config>saa>test>type-multi-line>lsp-ping>sr-policy path-destination)

Full Context

configure saa test type-multi-line lsp-ping sr-policy path-destination

Description

This command configures the IP address of the path destination from the range 127/8. When the LDP FEC prefix is IPv6, the user must enter a 127/8 IPv4 mapped IPv6 address, that is, in the range ::ffff:127/104.

The **no** form of this command removes the configuration.

Parameters

ip-address

Specifies the IP address.

Values ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

if-name

Specifies the name of an IP interface, up 32 characters, to send the MPLS echo request to. The name must already exist in the **config>router>interface** context.

Platforms

7705 SAR Gen 2

21.26 path-mtu

path-mtu

Syntax

path-mtu [*bytes*]

no path-mtu *bytes*

Context

[Tree] (config>service>pw-template path-mtu)

[Tree] (config>service>sdp path-mtu)

Full Context

configure service pw-template path-mtu

configure service sdp path-mtu

Description

This command configures the Maximum Transmission Unit (MTU) in bytes that the SDP can transmit to the far-end device router without packet dropping or IP fragmentation overriding the SDP-type default MTU.

The SDP-type default path-mtu can be overridden on a per SDP basis. Dynamic maintenance protocols on the SDP, like RSVP, may override this setting.

If the physical MTU on an egress interface indicates the next hop on an SDP path cannot support the current path-mtu, the system modifies the operational path-mtu on that SDP to a value that can be transmitted without fragmentation.

The **no** form of this command removes any path-mtu configured on the SDP, and the SDP uses the system default for the SDP type.

Default

the default **path-mtu** defined on the system for the type of SDP

Parameters

bytes

Specifies the bytes.

Values 576 to 9800

Platforms

7705 SAR Gen 2

21.27 path-mtu-discovery

path-mtu-discovery

Syntax

[**no**] **path-mtu-discovery**

Context

[Tree] (config>router>ldp>tcp-session-params>peer-transport path-mtu-discovery)

Full Context

configure router ldp tcp-session-parameters peer-transport path-mtu-discovery

Description

This command enables Path MTU discovery for the associated TCP connections. When enabled, the MTU for the associated TCP session is initially set to the egress interface MTU. The DF bit is also set so that if a router along the path of the TCP connection cannot handle a packet of a particular size without fragmenting, it sends back an ICMP message to set the path MTU for the given session to a lower value that can be forwarded without fragmenting.

If one or more transport addresses used in the Hello adjacencies to the same peer LSR are different from the LSR-ID value, the user must add each of the transport addresses to the path MTU discovery configuration as a separate peer. This means when the TCP connection is bootstrapped by a given Hello adjacency, the path MTU discovery can operate over that specific TCP connection by using its specific transport address.

Default

no path-mtu-discovery

Platforms

7705 SAR Gen 2

path-mtu-discovery

Syntax

[**no**] **path-mtu-discovery**

Context

[Tree] (config>router>bgp>group path-mtu-discovery)

[Tree] (config>router>bgp>group>neighbor path-mtu-discovery)

[Tree] (config>router>bgp path-mtu-discovery)

Full Context

configure router bgp group path-mtu-discovery

configure router bgp group neighbor path-mtu-discovery

configure router bgp path-mtu-discovery

Description

This command enables Path MTU Discovery (PMTUD) for the associated TCP connections.

When enabled, PMTUD is activated toward an IPv4 BGP neighbor. The Don't Fragment (DF) bit is set in the IP header of all IPv4 packets sent to the peer. If any device along the path toward the peer cannot forward the packet because the IP MTU of the interface is smaller than the IP packet size, the device drops the packet and sends an ICMP or ICMPv6 error message encoding the interface MTU. When the router receives the ICMP or ICMPv6 message, it lowers the TCP maximum segment size limit from the previous value to accomodate the IP MTU constraint.

When PMTUD is disabled and there is no **tcp-mss** configuration to associate with a BGP neighbor (in either the BGP configuration or the first-hop IP interface configuration), the router advertises a TCP MSS option of only 1024 bytes, limiting received TCP segments to that size.

The **no** form of this command disables PMTUD.

Default

no path-mtu-discovery

Platforms

7705 SAR Gen 2

21.28 path-preference

path-preference

Syntax

path-preference *value*

no path-preference

Context

[Tree] (config>router>mpls>lsp>secondary path-preference)

Full Context

configure router mpls lsp secondary path-preference

Description

This command enables the use of path preference among configured standby secondary paths per LSP. If all standby secondary paths have a default path-preference value then a non-standby secondary path will remain the active path while a standby secondary is available. A standby secondary path configured with the highest priority (for example, the lowest path-preference value) is made the active path when the primary is not in use. If multiple standby secondary paths have the same, lowest, path-preference value then the system will select the path with the highest up-time. Path preference can only be configured on the standby secondary paths.

The **no** form of this command resets the path-preference to the default value.

Default

path-preference 255

Parameters

value

Specifies an alternate path for the LSP if the primary path is not available.

Values 1 to 255

Platforms

7705 SAR Gen 2

21.29 path-profile

path-profile

Syntax

path-profile *profile-id* [**path-group** *group-id*]

no path-profile *profile-id*

Context

[Tree] (config>router>mpls>lsp-template path-profile)

[Tree] (config>router>mpls>lsp path-profile)

Full Context

configure router mpls lsp-template path-profile

configure router mpls lsp path-profile

Description

This command configures the PCE path profile and path group ID.

The PCE supports the computation of disjoint paths for two different LSPs originating and/or terminating on the same or different PE routers. To indicate this constraint to the PCE, the user must configure the PCE path profile ID and path group ID to which the PCE computed or PCE controlled LSP belongs to. These parameters are passed transparently by the PCC to the PCE and are thus opaque data to the router.

The association of the optional path-group ID is to allow the PCE to determine the profile ID that must be used with this path-group ID. One path-group ID is allowed per profile ID. The user can, however, enter the same path-group ID with multiple profile IDs by executing this command multiple times. A maximum of five **path-profile** [**path-group**] entries can be associated with the same LSP.

The **no** form of this command removes the path profile association with the LSP.

Parameters

profile-id

Specifies the profile ID.

Values 1 to 4294967295

path-group ***group-id***

Specifies the path group ID.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

21.30 path-type

path-type

Syntax

path-type {**ibgp** | **ebgp**}

no path-type

Context

[Tree] (config>router>policy-options>policy-statement>entry>from path-type)

Full Context

configure router policy-options policy-statement entry from path-type

Description

This command matches BGP routes based on their path type (EBGP or IBGP). A route learned from an EBGP peer has path-type **ebgp**. A route learned from an IBGP or confed-EBGP peer has path-type **ibgp**.

A non-BGP route does not match a policy entry if it contains the **path-type** command.

Default

no path-type

Parameters

ibgp

Matches routes from internal BGP peers.

ebgp

Matches routes from external BGP peers.

Platforms

7705 SAR Gen 2

21.31 patherr

patherr

Syntax

patherr [**detail**]

no patherr

Context

[Tree] (debug>router>rsvp>packet patherr)

Full Context

debug router rsvp packet patherr

Description

This command debugs path error packets.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about path error packets.

Platforms

7705 SAR Gen 2

21.32 pathtear

pathtear

Syntax

pathtear [**detail**]

no pathtear

Context

[Tree] (debug>router>rsvp>packet pathtear)

Full Context

debug router rsvp packet pathtear

Description

This command debugs path tear packets.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about path tear packets.

Platforms

7705 SAR Gen 2

21.33 pattern

pattern

Syntax

pattern *pad-value*

no pattern

Context

[Tree] (config>oam-pm>session>ip pattern)

Full Context

configure oam-pm session ip pattern

Description

This command configures the pattern value to be repeated in the padding portion of the TWAMP Light packet.

The **no** form of this command uses an incrementing byte pattern beginning with 00 and ending with FF, wrapping back to 00.

Default

pattern 0

Parameters

pad-value

Specifies the specific pattern to use.

Values 0 to 65535

Default 0

Platforms

7705 SAR Gen 2

21.34 pbr-down-action-override

pbr-down-action-override

Syntax

pbr-down-action-override *filter-action*

no pbr-down-action-override

Context

[Tree] (config>filter>ipv6-filter>entry pbr-down-action-override)

[Tree] (config>filter>ip-filter>entry pbr-down-action-override)

Full Context

configure filter ipv6-filter entry pbr-down-action-override

configure filter ip-filter entry pbr-down-action-override

Description

This command allows overriding the default action that is applied for entries with PBR/PBF action defined, when the PBR/PBF target is down.

The **no** form of the command preserves default behavior when PBR/PBF target is down.

Default

no pbr-down-action-override

Parameters

filter-action

Specifies the packets matching the entry.

drop — Specifies that packets matching the entry will be dropped if PBR/PBF target is down.

forward — Specifies that packets matching the entry will be forwarded if PBR/PBF target is down.

filter-default-action — Specifies that packets matching the entry will be processed as per **default-action** configuration for this filter if PBR/PBF target is down.

Platforms

7705 SAR Gen 2

21.35 pcap

pcap

Syntax

pcap *session-name* [**create**]

no pcap *session-name*

Context

[Tree] (config>mirror>mirror-dest pcap)

Full Context

configure mirror mirror-dest pcap

Description

This command specifies a PCAP instance used for packet capture.

The **no** form of this command removes the PCAP instance and stops the packet capture and file transfer session.

Parameters

session-name

Specifies the session name, up to 32 characters.

Platforms

7705 SAR Gen 2

pcap

Syntax

pcap *session-name*

Context

[Tree] (debug pcap)

Full Context

debug pcap

Description

This command specifies the session for the packet capture process.

Parameters

session-name

Specifies the session name, up to 32 characters.

Platforms

7705 SAR Gen 2

21.36 pcc

pcc

Syntax

[**no**] **pcc**

Context

[Tree] (debug>router>mpls>event pcc)

Full Context

debug router mpls event pcc

Description

This command debugs pcc events.

The **no** form of the command disables the debugging.

Platforms

7705 SAR Gen 2

pcc

Syntax

pcc

Context

[Tree] (config>router>pcep pcc)

Full Context

configure router pcep pcc

Description

Commands in this context configure PCC parameters.

Platforms
7705 SAR Gen 2

pcc

Syntax
[**no**] **pcc**

Context
[Tree] (debug>router>pcep pcc)

Full Context
debug router pcep pcc

Description
This command enables debugging for the PCEP Path Computation Client (PCC).

The **no** form of this command disables PCEP PCC debugging.

Platforms
7705 SAR Gen 2

21.37 pce-associations

pce-associations

Syntax
pce-associations

Context
[Tree] (config>router>pcep>pcc pce-associations)

Full Context
configure router pcep pcc pce-associations

Description
Commands in this context configure PCE association groups.

Platforms
7705 SAR Gen 2

pce-associations

Syntax

pce-associations

Context

[Tree] (config>router>mpls>lsp-template pce-associations)

[Tree] (config>router>mpls>lsp pce-associations)

Full Context

configure router mpls lsp-template pce-associations

configure router mpls lsp pce-associations

Description

Commands in this context configure LSP binding with one or more PCEP association groups.

Platforms

7705 SAR Gen 2

21.38 pce-control

pce-control

Syntax

[**no**] **pce-control**

Context

[Tree] (config>router>mpls>lsp pce-control)

[Tree] (config>router>mpls>lsp-template pce-control)

Full Context

configure router mpls lsp pce-control

configure router mpls lsp-template pce-control

Description

This command enables a PCE controlled LSP mode of operation. The **pce-control** option means the router delegates full control of the LSP to the PCE (PCE controlled). Enabling it means the PCE is acting in stateful-active mode for this LSP and the PCE will be able to reroute the path following a failure or re-optimize the path and update the router without a request from the router.

The user can delegate CSPF and non-CSPF LSPs, or LSPs that have the **path-computation-method pce** option enabled or disabled. The LSP maintains its latest active path computed by PCE or the router at the time it is delegated. The PCE only makes an update to the path at the next network event or reoptimization.

When configured to no, the PCE controlled mode of operation for the LSP has not effect.

Default

no pce-control

Platforms

7705 SAR Gen 2

21.39 pce-initiated-lsp

pce-initiated-lsp

Syntax

[**no**] **pce-initiated-lsp**

Context

[Tree] (config>router>mpls pce-initiated-lsp)

Full Context

configure router mpls pce-initiated-lsp

Description

This command creates a context to configure support for PCE-initiated LSPs.

The **no** form of this command removes PCE-initiated LSP support. All PCE-initiated LSPs are deleted.

Platforms

7705 SAR Gen 2

21.40 pce-report

pce-report

Syntax

pce-report rsvp-te {**enable** | **disable**}

pce-report sr-te {**enable** | **disable**}

Context

[Tree] (config>router>mpls pce-report)

Full Context

configure router mpls pce-report

Description

This command separately configures the reporting modes to a PCE for RSVP-TE or SR-TE LSPs. The PCC LSP database is synchronized with the PCE LSP database using the PCEP PCRpt (PCE Report) message for PCC-controlled, PCE-computed, and PCE-controlled LSPs.

The global MPLS level **pce-report** command can be used to enable or disable PCE reporting for all SR-TE LSPs or RSVP-TE LSPs during PCE LSP database synchronization. This configuration is inherited by all LSPs of the specified type. The PCC reports both CSPF and non-CSPF LSPs. The default value is disabled for both types of LSP. This default value is meant to control the introduction of the PCE into an existing network and to let the operator decide if all LSPs of a particular type need to be reported.

The LSP-level **pce-report** command overrides the global configuration for the reporting of LSPs to the PCE. The default value is to inherit the global MPLS level value. The **enable** or **disable** value allows for the override of the inherited value. The **inherit** value explicitly resets the LSP to inherit the global configuration for that LSP type.

If PCE reporting is disabled for the LSP, either due to inheritance or due to LSP-level configuration, then enabling the **pce-control** option for the LSP has no effect.

Default

pce-report rsvp-te disable

pce-report sr-te disable

Parameters

rsvp-te

Specifies the PCE reporting mode for all TE LSPs of RSVP-TE type.

Values **enable** — enables PCE reporting for all TE LSPs of RSVP-TE type

disable — disables PCE reporting for all TE LSPs of RSVP-TE type

sr-te

Specifies the PCE reporting mode for all TE LSPs of SR-TE type.

Values **enable** — enables PCE reporting for all TE LSPs of SR-TE type

disable — disables PCE reporting for all TE LSPs of SR-TE type

Platforms

7705 SAR Gen 2

pce-report

Syntax

pce-report {**enable** | **disable** | **inherit**}

Context

[Tree] (config>router>mpls>lsp-template pce-report)

[Tree] (config>router>mpls>lsp pce-report)

Full Context

configure router mpls lsp-template pce-report

configure router mpls lsp pce-report

Description

This command separately configures the reporting modes to a PCE for RSVP-TE or SR-TE LSPs.

The PCC LSP database is synchronized with the PCE LSP database using the PCEP PCRpt (PCE Report) message for PCC-controlled, PCE-computed and PCE-controlled LSPs.

The global MPLS-level **pce-report** command can be used to enable or disable PCE reporting for all SR-TE LSPs or RSVP-TE LSPs during PCE LSP database synchronization. This configuration is inherited by all LSPs of the specified type. The PCC reports both CSPF and non-CSPF LSPs. The default value is disabled for both types of LSP. This default value is meant to control the introduction of the PCE into an existing network and to let the operator decide if all LSPs of a particular type need to be reported.

The LSP-level **pce-report** command overrides the global configuration for the reporting of LSP to the PCE. The default value is to inherit the global MPLS level value. The **enable** or **disable** value allows for the override of the inherited value. The **inherit** value explicitly resets the LSP to inherit the global configuration for that LSP type.

If PCE reporting is disabled for the LSP, either due to inheritance or due to LSP-level configuration, then enabling the **pce-control** option for the LSP has no effect.

Default

pce-report inherit

Parameters

enable

Enables PCE reporting.

disable

Disables PCE reporting.

inherit

Inherits the global configuration for PCE reporting.

Platforms

7705 SAR Gen 2

21.41 pcep

pcep

Syntax

[**no**] **pcep**

Context

[Tree] (config>router pcep)

Full Context

configure router pcep

Description

This command enables Path Computation Element communications Protocol (PCEP), and enters the context to configure PCEP parameters.

The **no** form of the command disables PCEP.

Platforms

7705 SAR Gen 2

pcep

Syntax

[**no**] **pcep**

Context

[Tree] (debug>router pcep)

Full Context

debug router pcep

Description

This command enables debugging for the Path Computation Element Protocol (PCEP).

The **no** form of this command disables PCEP debugging.

Platforms

7705 SAR Gen 2

21.42 pe-id-mac-flush-interop

pe-id-mac-flush-interop

Syntax

[**no**] **pe-id-mac-flush-interop**

Context

[Tree] (config>router>ldp>session-params>peer pe-id-mac-flush-interop)

Full Context

configure router ldp session-parameters peer pe-id-mac-flush-interop

Description

This command enables the addition of the PE-ID TLV in the LDP MAC withdrawal (mac-flush) message, under certain conditions, and modifies the mac-flush behavior for interoperability with other vendors that do not support the flush-all-from-me vendor-specific TLV. This flag can be enabled on a per LDP peer basis and allows the flush-all-from-me interoperability with other vendors. When the pe-id-mac-flush-interop flag is enabled for a given peer, the current mac-flush behavior is modified in terms of mac-flush generation, mac-flush propagation and behavior upon receiving a mac-flush.

The mac-flush generation will be changed depending on the type of event and according to the following rules:

- Any all-from-me mac-flush event will trigger a mac-flush all-but-mine message (RFC 4762 compliant format) with the addition of a PE-ID TLV. The PE-ID TLV contains the IP address of the sending PE.
- Any all-but-mine mac-flush event will trigger a mac-flush all-but-mine message WITHOUT the addition of the PE-ID TLV, as long as the source spoke SDP is not part of an end-point.
- Any all-but-mine mac-flush event will trigger a mac-flush all-but-mine message WITH the addition of the PE-ID TLV, if the source spoke SDP is part of an end-point and the spoke-sdp goes from down/standby state to active state. In this case, the PE-ID TLV will contain the IP address of the PE to which the previous active spoke-sdp was connected to.

Any other case will follow the existing mac-flush procedures.

When the pe-id-mac-flush-interop flag is enabled for a given LDP peer, the mac-flush ingress processing is modified according to the following rules:

- Any received all-from-me mac-flush will follow the existing mac-flush all-from-me rules regardless of the existence of the PE-ID.
- Any received all-but-mine mac-flush will take into account the received PE-ID, that is all the mac addresses associated to the PE-ID will be flushed. If the PE-ID is not included, the mac addresses associated to the sending PE will be flushed.
- Any other case will follow the existing mac-flush procedures.

When a mac-flush message has to be propagated (for an ingress sdp-binding to an egress sdp-binding) and the pe-id-mac-flush-interop flag is enabled for the ingress and egress TLDP peers, the following behavior is observed:

- If the ingress and egress bindings are spoke SDP, the PE will propagate the mac-flush message with its own PE-ID.
- If the ingress binding is an spoke SDP and the egress binding a mesh SDP, the PE will propagate the mac-flush message without modifying the PE-ID included in the PE-ID TLV.
- If the ingress binding is a mesh SDP and the egress binding an spoke SDP, the PE will propagate the mac-flush message with its own PE-ID.
- When ingress and egress bindings are mesh-sdp, the mac-flush message is never propagated. This is the behavior regardless of the pe-id-mac-flush-interop flag configuration.

The PE-ID TLV is never added when generating a mac-flush message on a B-VPLS if the **send-bvpls-flush** command is enabled in the I-VPLS. In the same way, no PE-ID is added when propagating mac-flush from a B-VPLS to a I-VPLS when the **propagate-mac-flush-from-bvpls** command is enabled. Mac-flush messages for peers within the same I-VPLS or within the same B-VPLS domain follow the procedures described above.

Default

no pe-id-mac-flush-interop

Platforms

7705 SAR Gen 2

21.43 peer

peer

Syntax

peer *ip-address* **tag** *sync-tag*

no peer

Context

[Tree] (config>router>dhcp6>server>failover peer)

[Tree] (config>router>dhcp6>server>pool>failover peer)

[Tree] (config>router>dhcp>server>pool>failover peer)

[Tree] (config>router>dhcp>server>failover peer)

Full Context

configure router dhcp6 local-dhcp-server failover peer

configure router dhcp6 local-dhcp-server pool failover peer

configure router dhcp local-dhcp-server pool failover peer

configure router dhcp local-dhcp-server failover peer

Description

This command creates a sync tag. DHCP leases can be synchronized per DHCP server or DHCP pool. The pair of synchronizing servers or pools is identified by a tag. The synchronization information is carried over the Multi-Chassis Synchronization (MCS) link between the two peers. MCS link is a logical link (IP, or MPLS).

MCS runs over TCP, port 45067 and it is using either data traffic or keepalives to detect failure on the communication link between the two nodes. In the absence of any MCS data traffic for more than 0.5sec, MCS will send its own keepalive to the peer. If a reply is **not** received within three sec, MCS will declare its operation state as DOWN and the DB Sync state as out-of-sync. MCS will consequently notify its clients (DHCP Server being one of them) of this. It can take up to three seconds before the DHCP client realizes that the inter-chassis communication link has failed.

The inter-chassis communication link failure does not necessarily assume the same failed fate for the access links. The two redundant nodes can become isolated from each other in the network. This occurs when only the intercommunication (MCS) link fails. It is important that this MCS link be highly redundant.

The **no** form of this command reverts to the default.

Parameters

ip-address

Specifies the IPv4 or IPv6 address of the peer.

Values

ipv4-address:	a.b.c.d	
:	x:ipv6-addressx:x:x:x:x:x:x (eight 16-bit pieces)	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 to FFFF]H
	d:	[0 to 255]D

tag *sync-tag*

Specifies a tag, up to 32 characters, that identifies the synchronizing DHCP servers or pools.

Platforms

7705 SAR Gen 2

peer

Syntax

peer *ip-address* [**create**]

no peer *ip-address*

Context

[Tree] (config>redundancy>multi-chassis peer)

Full Context

configure redundancy multi-chassis peer

Description

This command configures the IP address of the peer in a redundant multi-chassis setup, and enters the context for further, application-specific configuration options.

Parameters

ip-address

Specifies a peer IP address. Multicast addresses are not allowed.

Values ipv4-address: a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF] H
- d: [0 to 255] D

Platforms

7705 SAR Gen 2

peer

Syntax

[**no**] **peer** *ip-address*

Context

[Tree] (config>router>ldp>session-parameters peer)

Full Context

configure router ldp session-parameters peer

Description

This command configures parameters for an LDP peer.

Parameters

ip-address

Specifies the IPv4 or IPv6 address of the LDP peer in dotted decimal notation.

Platforms

7705 SAR Gen 2

peer

Syntax

[**no**] **peer** *ip-address*

Context

[Tree] (config>router>ldp>targeted-session peer)

Full Context

configure router ldp targeted-session peer

Description

This command configures parameters for an LDP peer.

The **no** form of this command removes the LDP peer parameters.

Parameters

ip-address

Specifies a peer IP address.

Values ipv4-address: a.b.c.d; 0 to 255, decimal

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]; hexadecimal
- d: [0 to 255]; decimal

Platforms

7705 SAR Gen 2

peer

Syntax

[**no**] **peer** *ip-address*

Context

[Tree] (debug>router>ldp peer)

Full Context

debug router ldp peer

Description

Use this command for debugging an LDP peer.

Parameters

ip-address

The IP address of the LDP peer.

Platforms

7705 SAR Gen 2

peer

Syntax

peer *ip-address* [**preference** *preference*]

no peer *ip-address*

Context

[Tree] (config>router>pcep>pcc peer)

Full Context

configure router pcep pcc peer

Description

This command configures the IP address of a peer PCEP speaker. The address is used as the destination address in the PCEP session messages to a PCEP peer.

The **preference** parameter allows the PCC to select the preferred PCE when both have their PCEP sessions successfully established. A maximum of two PCEP peers is supported.

The PCE peer that is not in overload is always selected by the PCC as the active PCE. However, if neither of the PCEs are signaling the overload state, the PCE with the higher numerical preference value is selected, and in case of a tie, the PCE with the lower IP address is selected.

Note:
The system does not support two or more simultaneously active PCEs.

The **no** form of the command removes the specified peer PCEP speaker.

Parameters

ip-address

The IP address of the PCEP peer to be used as the destination address in the PCEP session.

preference

The preference value of the peer.

Values 0 to 100

Default 1

Platforms

7705 SAR Gen 2

peer

Syntax

peer [**router** *router-instance* | **service-name** *service-name*] {*ip-address* | *ipv6-address*} [**key-id** *key-id* | **authentication-keychain** *keychain-name*] [**version** *version*] [**prefer**]

no peer [**router** *router-instance* | **service-name** *service-name*] {*ip-address* | *ipv6-address*}

Context

[Tree] (config>system>time>ntp peer)

Full Context

configure system time ntp peer

Description

This command configures symmetric active mode for an NTP peer. It is recommended to configure authentication and to only configure known time servers as peers. Peers may exist within a VPRN service.

Note:
For symmetric peering to operate correctly with a peer accessible through a VPRN, local NTP server functionality must be enabled within the VPRN using the **configure service vprn ntp** command.

The **no** form of the command removes the configured peer.

Parameters

router-instance

Specifies the routing context that contains the interface.

Values *router-name* — Base | Management

service-id — 1 to 2147483647

Default Base

service name

Specifies the service name for the VPRN, up to 64 characters. CPM routing instances are not supported.

ip-address

Configures the IPv4 address of the peer that requires a peering relationship to be set up.

Values a.b.c.d

ipv6-address

Configures the IPv6 address of the peer that requires a peering relationship to be set up.

Values
- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF] H
- d: [0 to 255] D

key-id

Specifies the key ID. Successful authentication requires that both peers must have the same authentication key-id, type, and key value.

Specify the *key-id* that identifies the configured authentication key and authentication type used by this node to transmit NTP packets to an NTP peer. If an NTP packet is received by these nodes, the authentication key-id, type, and key value must be valid, otherwise the packet is rejected and an event or trap is generated.

Values 1 to 255

keychain-name

Identifies the keychain name, up to 32 characters.

version

Specifies the NTP version number that is generated by this node. This parameter does not need to be configured when in client mode, in which case all versions are accepted.

Values 2 to 4

Default 4

prefer

When configuring more than one peer, one remote system can be configured as the preferred peer. When a second peer is configured as preferred, the new entry overrides the old entry.

Platforms

7705 SAR Gen 2

peer

Syntax

peer *ip-address* [**create**]

no peer *ip-address*

Context

[Tree] (config>router>ipsec>mc-shunt-profile peer)

[Tree] (config>service>vprn>ipsec>mc-shunt-profile peer)

Full Context

configure router ipsec multi-chassis-shunting-profile peer

configure service vprn ipsec multi-chassis-shunting-profile peer

Description

Commands in this context configure a multi-chassis IPsec peer IP address for the **multi-chassis-shunting-profile**.

The **no** form of this command removes the peer IP address from the configuration.

Default

no command

Parameters

ip-address

Specifies a peer IP address.

Values ipv4-address: a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF] H
- d: [0 to 255] D

create

Keyword used to create the command instance.

Platforms

7705 SAR Gen 2

21.44 peer-as

peer-as

Syntax

peer-as *as-number*

Context

[Tree] (config>service>vprn>bgp>group peer-as)

[Tree] (config>service>vprn>bgp>group>neighbor peer-as)

Full Context

configure service vprn bgp group peer-as

configure service vprn bgp group neighbor peer-as

Description

This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.

For EBGP peers, the peer AS number configured must be different from the autonomous system number configured for this router under the global level since the peer will be in a different autonomous system than this router

For IBGP peers, the peer AS number must be the same as the autonomous system number of this router configured under the global level.

This is a required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.

Default

No AS numbers are defined.

Parameters

as-number

The autonomous system number, expressed as a decimal integer.

Values 1 to 65535

Platforms

7705 SAR Gen 2

peer-as

Syntax

peer-as *as-number*

Context

[Tree] (config>router>bgp>group peer-as)

[Tree] (config>router>bgp>group>neighbor peer-as)

Full Context

configure router bgp group peer-as

configure router bgp group neighbor peer-as

Description

This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.

For EBGP peers, the peer AS number configured must be different from the autonomous system number configured for this router under the global level since the peer will be in a different autonomous system than this router.

For IBGP peers, the peer AS number must be the same as the autonomous system number of this router configured under the global level.

This is required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.

Parameters

as-number

Specifies the autonomous system number expressed as a decimal integer.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

21.45 peer-group

peer-group

Syntax

peer-group *tunnel-group-id*

no peer-group

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ipsec>tunnel-group peer-group)

Full Context

configure redundancy multi-chassis peer mc-ipsec tunnel-group peer-group

Description

This command specifies the corresponding tunnel-group ID on peer node. The peer tunnel-group ID does not necessary equals to local tunnel-group ID.

The **no** form of this command removes the tunnel-group ID from the configuration.

Parameters

tunnel-group-id

Specifies the tunnel-group identifier.

Values 1 to 16

Platforms

7705 SAR Gen 2

21.46 peer-ip-prefix

peer-ip-prefix

Syntax

peer-ip-prefix *ip-prefix/ip-prefix-length*

peer-ip-prefix ipv4-any

peer-ip-prefix ipv6-any

no peer-ip-prefix

Context

[Tree] (config>ipsec>client-db>client>client-id peer-ip-prefix)

Full Context

configure ipsec client-db client client-identification peer-ip-prefix

Description

This command specifies match criteria that uses the peer's tunnel IP address as the input. Only one peer-ip-prefix criteria can be configured for a given client entry.

The **no** form of this command reverts to the default.

Default

no peer-ip-prefix

Parameters

ip-prefix/ip-prefix-length

Specifies an IPv4 or IPv6 prefix. It is considered a match if the peer's tunnel IP address is within the specified prefix.

ipv4-any

Matches any IPv4 address.

ipv6-any

Matches any IPv6 address.

Platforms

7705 SAR Gen 2

peer-ip-prefix

Syntax

[**no**] **peer-ip-prefix**

Context

[Tree] (config>ipsec>client-db>match-list peer-ip-prefix)

Full Context

configure ipsec client-db match-list peer-ip-prefix

Description

This command enables the use of the peer's tunnel IP address as the match input.

The **no** form of this command disables the peer IP prefix matching process.

Default

no peer-ip-prefix

Platforms

7705 SAR Gen 2

21.47 peer-name

peer-name

Syntax

peer-name *name*

no peer-name

Context

[Tree] (config>redundancy>multi-chassis>peer peer-name)

Full Context

configure redundancy multi-chassis peer peer-name

Description

This command specifies a peer name.

Default

no peer-name

Parameters

name

Specifies the string up to 32 characters. Any printable, seven-bit ASCII characters can be used within the string. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

21.48 peer-template

peer-template

Syntax

[**no**] **peer-template** *template-name*

Context

[Tree] (config>router>ldp>targeted-session peer-template)

Full Context

configure router ldp targeted-session peer-template

Description

This command creates a targeted session peer parameter template that can be referenced in the automatic creation of targeted Hello adjacency and LDP session to a discovered peer.

The **no** form of this command deletes the peer template. A peer template cannot be deleted if it is bound to a peer prefix list.

Parameters

template-name

Specifies the template name to identify targeted peer template. It must be 32 characters maximum.

Platforms

7705 SAR Gen 2

21.49 peer-template-map

peer-template-map

Syntax

peer-template-map *template-name* **policy** *peer-prefix-policy1* [*peer-prefix-policy2..up to 5*]

no peer-template-map peer-template *template-name*

Context

[Tree] (config>router>ldp>targeted-session peer-template-map)

Full Context

configure router ldp targeted-session peer-template-map

Description

This command enables the automatic creation of a targeted Hello adjacency and LDP session to a discovered peer. The user configures a targeted session peer parameter template and binds it to a peer prefix policy.

Each application of a targeted session template to a given prefix in the prefix list will result in the establishment of a targeted Hello adjacency to an LDP peer using the template parameters as long as the prefix corresponds to a router-id for a node in the TE database. As a result of this, the user must enable the traffic-engineering option in ISIS or OSPF. The targeted Hello adjacency will either trigger a new LDP session or will be associated with an existing LDP session to that peer.

Up to 5 peer prefix policies can be associated with a single peer template at all times. Also, the user can associate multiple templates with the same or different peer prefix policies. Thus multiple templates can match with a given peer prefix. In all cases, the targeted session parameters applied to a given peer prefix are taken from the first created template by the user. This provides a more deterministic behavior regardless of the order in which the templates are associated with the prefix policies.

Each time the user executes the above command, with the same or different prefix policy associations, or the user changes a prefix policy associated with a targeted peer template, the system re-evaluates the prefix policy. The outcome of the re-evaluation will tell LDP if an existing targeted Hello adjacency needs to be torn down or if an existing targeted Hello adjacency needs to have its parameters updated on the fly.

If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with a targeted peer template, the same prefix policy re-evaluation described above is performed.

The template comes up in the **no shutdown** state and as such it takes effect immediately. Once a template is in use, the user can change any of the parameters on the fly without shutting down the template. In this case, all targeted Hello adjacencies are updated.

The SR OS supports multiple ways of establishing a targeted Hello adjacency to a peer LSR:

- User configuration of the peer with the targeted session parameters inherited from the **config>router>ldp>targeted-session** in the top level context or explicitly configured for this peer in the **config>router>ldp>targ-session>peer** context and which overrides the top level parameters shared by all targeted peers. Let us refer to the top level configuration context as the global context. Some

parameters only exist in the global context; their value will always be inherited by all targeted peers regardless of which event triggered it.

- User configuration of an SDP of any type to a peer with the signaling tldp option enabled (default configuration). In this case the targeted session parameter values are taken from the global context.
- User configuration of a (FEC 129) PW template binding in a BGP-VPLS service. In this case the targeted session parameter values are taken from the global context.
- User configuration of a (FEC 129 type II) PW template binding in a VLL service (dynamic multi-segment PW). In this case the target session parameter values are taken from the global context
- User configuration of a mapping of a targeted session peer parameter template to a prefix policy when the peer address exists in the TE database (this feature). In this case, the targeted session parameter values are taken from the template.

Since the above triggering events can occur simultaneously or in any arbitrary order, the LDP code implements a priority handling mechanism in order to decide which event overrides the active targeted session parameters. The overriding trigger will become the owner of the targeted adjacency to a given peer. The following is the priority order:

- Priority 1: manual configuration of session parameters
- Priority 2: mapping of targeted session template to prefix policy.
- Priority 3: auto-tx parameters
- Priority 4: auto-rx parameters
- Priority 5: manual configuration of SDP, PW template binding in BGP-AD VPLS and in FEC 129 VLL.

Any parameter value change to an active targeted Hello adjacency caused by any of the above triggering events is performed on the fly by having LDP immediately send a Hello message with the new parameters to the peer without waiting for the next scheduled time for the Hello message. This allows the peer to adjust its local state machine immediately and maintains both the Hello adjacency and the LDP session in UP state. The only exceptions are the following:

- The triggering event caused a change to the local-lsr-id parameter value. In this case, the Hello adjacency is brought down which will also cause the LDP session to be brought down if this is the last Hello adjacency associated with the session. A new Hello adjacency and LDP session will then get established to the peer using the new value of the local LSR ID.
- The triggering event caused the targeted peer shutdown option to be enabled. In this case, the Hello adjacency is brought down which will also cause the LDP session to be brought down if this is the last Hello adjacency associated with the session.

Finally, the value of any LDP parameter which is specific to the LDP/TCP session to a peer is inherited from the **config>router>ldp>session-params>peer** context. This includes MD5 authentication, LDP prefix per-peer policies, label distribution mode (DU or DOD), and so on.

The **no** form of this command deletes the binding of the template to the peer prefix list and brings down all Hello adjacencies to the discovered LDP peers.

Platforms

7705 SAR Gen 2

21.50 peer-tracking-policy

peer-tracking-policy

Syntax

peer-tracking-policy *policy-name*

no peer-tracking-policy

Context

[Tree] (config>service>vprn>bgp peer-tracking-policy)

Full Context

configure service vprn bgp peer-tracking-policy

Description

This command specifies the name of a policy statement to use with the BGP peer-tracking function on the BGP sessions where this is enabled. The policy controls which IP routes in RTM are eligible to indicate reachability of IPv4 and IPv6 BGP neighbor addresses. If the longest matching route in RTM for a BGP neighbor address is an IP route that is rejected by the policy, or it is a BGP route accepted by the policy, or if there is no matching route, the neighbor is considered unreachable and BGP tears down the peering session and holds it in the idle state until a valid route is once again available and accepted by the policy.

The default peer-tracking policy (when the **no peer-tracking-policy** command is configured) is to use the longest matching active route in RTM that is not an LDP shortcut route or an aggregate route.

Note:
When **peer-tracking** is configured, the peer-tracking policy should only permit one of **direct-interface** or **direct** routes to be advertised to a BGP peer. Advertising both routes will cause the best route to oscillate.

Default

no peer-tracking-policy

Parameters

policy-name

Specifies the route policy name. Allowed values are any string up to 64 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes. Route policies are configured in the **config>router>policy-options** context.

Platforms

7705 SAR Gen 2

peer-tracking-policy

Syntax

peer-tracking-policy *policy-name*

no peer-tracking-policy

Context

[Tree] (config>router>bgp peer-tracking-policy)

Full Context

configure router bgp peer-tracking-policy

Description

This command specifies the name of a policy statement to use with the BGP peer-tracking function on the BGP sessions where this is enabled. The policy controls which IP routes in RTM are eligible to indicate reachability of IPv4 and IPv6 BGP neighbor addresses. If the longest matching route in RTM for a BGP neighbor address is an IP route that is rejected by the policy, or it is a BGP route accepted by the policy, or if there is no matching route, the neighbor is considered unreachable and BGP tears down the peering session and holds it in the idle state until a valid route is once again available and accepted by the policy.

The default peer-tracking policy (when the **no peer-tracking-policy** command is configured) is to use the longest matching active route in RTM that is not an LDP shortcut route or an aggregate route.

Note:
When **peer-tracking** is configured, the peer-tracking policy should only permit one of **direct-interface** or **direct** routes to be advertised to a BGP peer. Advertising both routes will cause the best route to oscillate.

Default

no peer-tracking-policy

Parameters

policy-name

Specifies the route policy name. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

Platforms

7705 SAR Gen 2

21.51 peer-transport

peer-transport

Syntax

peer-transport *ip-address*

no peer transport

Context

[Tree] (config>router>ldp>tcp-session-parameters peer-transport)

Full Context

configure router ldp tcp-session-parameters peer-transport

Description

This command configures the peer transport address, that is, the destination address of the TCP connection, and not the address corresponding to the LDP LSR-ID of the peer.

Parameters

ip-address

Specifies the IPv4 or IPv6 address of the TCP connection to the LDP peer in dotted decimal notation.

Platforms

7705 SAR Gen 2

21.52 pending-requests-limit

pending-requests-limit

Syntax

pending-request-limit *limit*

no pending-request-limit

Context

[Tree] (config>router>radius-server>server pending-requests-limit)

[Tree] (config>service>vprn>radius-server>server pending-requests-limit)

Full Context

configure router radius-server server pending-requests-limit

configure service vprn radius-server server pending-requests-limit

Description

This command specifies the per-server maximum number of outstanding requests sent to the RADIUS server. If the maximum number is exceeded, the next RADIUS server in the pool is selected.

The **no** form of this command removes the limit value from the configuration.

Default

pending-requests-limit 4096

Parameters

limit

Specifies the maximum number of outstanding requests sent to the RADIUS server.

Values 1 to 4096

Platforms

7705 SAR Gen 2

21.53 per-host-authentication

per-host-authentication

Syntax

[**no**] **per-host-authentication**

Context

[Tree] (config>port>ethernet>dot1x per-host-authentication)

Full Context

configure port ethernet dot1x per-host-authentication

Description

This command enables dot1x authenticating per host source mac or VLAN. The port does not allow traffic from any hosts or any MAC. When a host is authenticated via RADIUS policy, its source mac is then allowed through the port, while the port is closed for any other mac. Any traffic from the allowed host is forwarded on the port, including untagged and tagged traffic.

Default

no per-host-authentication

Platforms

7705 SAR Gen 2

21.54 per-peer-queuing

per-peer-queuing

Syntax

[**no**] **per-peer-queuing**

Context

[Tree] (config>system>security per-peer-queuing)

Full Context

configure system security per-peer-queuing

Description

This command enables CPM hardware queuing per peer. This means that when a peering session is established, the router will automatically allocate a separate CPM hardware queue for that peer.

The **no** form of this command disables CPM hardware queuing per peer.

Default

per-peer-queuing

Platforms

7705 SAR Gen 2

21.55 per-user

per-user

Syntax

per-user user-directory *dir-url* **file-name** *file-name*

no per-user

Context

[Tree] (config>system>login-control>login-scripts per-user)

Full Context

configure system login-control login-scripts per-user

Description

This command allows users to define their own login scripts that can be executed each time they first login to a CLI session. The command executes the script " *file-url* / *username* / *file-name*" when the user *username* logs into a CLI session (authenticated by any means including local user database, TACACS+, or RADIUS).

For example:

per-user user-directory "cf1:/local/users" file-name "login-script.txt"

would search for the following script when user "admin" logs in and authenticates via RADIUS:

cf1:/local/users/admin/login-script.txt

The per user login script is executed after any global script executes and before any login-exec script configured against a local user is executed. This allows users, for example, who are authenticated via TACACS+ or RADIUS to define their own login scripts.

This CLI script executes in the context of the user who opens the CLI session. Any commands in the script that the user is not authorized to execute will fail.

The **no** form of this command disables the execution of any per user login-scripts.

Default

no per-user

Parameters

dir-url

Specifies the path or directory name.

file-name

Specifies the name of the file (located in the *dir-url* directory) including the extension.

Platforms

7705 SAR Gen 2

21.56 percent-rate

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*]

Context

[Tree] (config>port>eth>access>egr>qgrp>qover>q percent-rate)

[Tree] (config>port>ethernet>network>egr>qgrp>qover>q percent-rate)

Full Context

configure port ethernet access egress queue-group queue-overrides queue percent-rate

configure port ethernet network egress queue-group queue-overrides queue percent-rate

Description

This command specifies percent rates (CIR and PIR).

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the **percent-rate** is performed under the **hs-wrr-group** within the egress queue group template.

Parameters

pir-percent

Specifies the PIR as a percentage.

Values 0.01 to 100.00

cir-percent

Specifies the CIR as a percentage.

Values 0.00 to 100.00

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*]

no percent-rate

Context

[Tree] (config>service>epipe>sap>egress>policer-over>plcr percent-rate)

Full Context

configure service epipe sap egress policer-override policer percent-rate

Description

This command configures the percent rates (CIR and PIR) override.

Parameters

pir-percent

Specifies the policer's PIR as a percentage of the policers's parent arbiter rate.

Values 0.01 to 100.00

Default 100.00

cir-percent

Specifies the policer's CIR as a percentage of the policers's parent arbiter rate.

Values 0.00 to 100.00

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*]

no percent-rate

Context

[Tree] (config>service>epipe>sap>egress>queue-override>queue percent-rate)

Full Context

configure service epipe sap egress queue-override queue percent-rate

Description

The percent-rate command within the SAP ingress and egress QoS policy enables supports for a queue's PIR and CIR rate to be configured as a percentage of the egress port's line rate or of its parent scheduler's rate.

When the rates are expressed as a port-limit, the actual rates used per instance of the queue will vary based on the port speed. For example, when the same QoS policy is used on a 1-Gb and a 10-Gb Ethernet port, the queue's rates will be 10 times greater on the 10 Gb port due to the difference in port speeds. This enables the same QOS policy to be used on SAPs on different ports without needing to use SAP based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's PIR and CIR rates will be recalculated based on the defined percentage value.

When the rates are expressed as a local-limit, the actual rates used per instance of the queue are relative to the queue's parent scheduler rate. This enables the same QOS policy to be used on SAPs with different parent scheduler rates without needing to use SAP based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the parent scheduler rate changes after the queue is created, the queue's PIR and CIR rates will be recalculated based on the defined percentage value.

Queue rate overrides can only be specified in the form as configured in the QoS policy (a SAP override can only be specified as a percent-rate if the associated QoS policy was also defined as percent-rate).

Likewise, a SAP override can only be specified as a rate (kb/s) if the associated QoS policy was also defined as a rate. Queue-overrides are relative to the limit type specified in the QOS policy.

When no percent-rate is defined within a SAP ingress or egress queue-override, the queue reverts to the defined shaping and CIR rates within the SAP ingress and egress QOS policy associated with the queue.

Parameters

percent-of-line-rate

The percent-of-line-rate parameter is used to express the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

pir-percent

Specifies the queue's PIR as a percentage dependent on the use of the port-limit or local-limit.

Values 0.01 to 100.00

Default 100.00

cir-percent

Specifies the queue's CIR as a percentage dependent on the use of the port-limit or local-limit.

Values 0.00 to 100.00

Default 100.00

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*]

no percent-rate

Context

[Tree] (config>service>vpls>sap>egress>policer-override>plcr percent-rate)

[Tree] (config>service>vpls>sap>ingress>policer-override>plcr percent-rate)

Full Context

configure service vpls sap egress policer-override policer percent-rate

configure service vpls sap ingress policer-override policer percent-rate

Description

This command configures the percent rates (CIR and PIR) override and can only be used when the rate for the associated policer in the applied SAP ingress QoS policy is also configured with the **percent-rate** command.

The **no** form of this command removes the **percent-rate** override so that the **percent-rate** configured for the policer in the applied SAP egress QoS policy is used.

Parameters

pir-percent

Specifies the policer's PIR as a percentage of the policers' parent arbiter rate.

Values 0.01 to 100.00

Default 100.00

cir-percent

Specifies the policer's CIR as a percentage of the policers' parent arbiter rate.

Values 0.00 to 100.00

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*]

Context

[Tree] (config>service>vpls>sap>ingress>queue-override>queue percent-rate)

[Tree] (config>service>vpls>sap>egress>queue-override>queue percent-rate)

Full Context

configure service vpls sap ingress queue-override queue percent-rate

configure service vpls sap egress queue-override queue percent-rate

Description

The **percent-rate** command supports a queue’s shaping rate and CIR rate as a percentage of the egress port’s line rate. When the rates are expressed as a percentage within the template, the actual rate used per instance of the queue group queue-id will vary based on the port speed. For example, when the same template is used to create a queue group on a 1-Gb and a 10-Gb Ethernet port, the queue’s rates will be 10 times greater on the 10-Gb port due to the difference in port speeds. This enables the same template to be used on multiple ports without needing to use port based queue overrides to modify a queue’s rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's shaping and CIR rates will be recalculated based on the defined percentage value.

The **rate** and **percent-rate** commands override one another. If the current rate for a queue is defined using the percent-rate command and the **rate** command is executed, the percent-rate values are deleted. In a similar fashion, the **percent-rate** command causes any **rate** command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at any time.

An egress port queue group queue rate override may be expressed as either a percentage or an explicit rate independent on how the queue's template rate is expressed.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case the configuration of the **percent-rate** is performed under the **hs-wrr-group** within the SAP egress QoS policy.

The **no** form of this command returns the queue to its default shaping **rate** and **cir** rate. When **no percent-rate** is defined within a port egress queue group queue override, the queue reverts to the defined shaping and CIR rates within the egress queue group template associated with the queue.

Parameters

pir-percent

Specifies the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

Values 0.01 to 100.00

Default 100.00

cir-percent

Specifies the queue's committed scheduling rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

Values 0.00 to 100.00

Default 100.00

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*]

no percent-rate

Context

[Tree] (config>service>ies>if>sap>egress>policer-override>plcr percent-rate)

[Tree] (config>service>ies>if>sap>ingress>policer-override>plcr percent-rate)

Full Context

configure service ies interface sap egress policer-override policer percent-rate

configure service ies interface sap ingress policer-override policer percent-rate

Description

This command configures the percent rates (CIR and PIR) override and can only be used when the rate for the associated policer in the applied SAP ingress QoS policy is also configured with the **percent-rate** command.

The **no** form of this command removes the **percent-rate** override so that the **percent-rate** configured for the policer in the applied SAP egress QoS policy is used.

Parameters

pir-percent

Specifies the policer's PIR as a percentage of the policers's parent arbiter rate.

Values 0.01 to 100.00

Default 100.00

cir-percent

Specifies the policer's CIR as a percentage of the policers's parent arbiter rate.

Values 0.00 to 100.00

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*]

no percent-rate

Context

[Tree] (config>service>ies>if>sap>ingress>queue-override>queue percent-rate)

[Tree] (config>service>ies>if>sap>egress>queue-override>queue percent-rate)

Full Context

configure service ies interface sap ingress queue-override queue percent-rate

configure service ies interface sap egress queue-override queue percent-rate

Description

The **percent-rate** command supports a queue's shaping rate and CIR rate as a percentage of the egress port's line rate. When the rates are expressed as a percentage within the template, the actual rate used per instance of the queue group queue-id will vary based on the port speed. For example, when the same template is used to create a queue group on a 1-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit port due to the difference in port speeds. This enables the same template to be used on multiple ports without needing to use port based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's shaping and CIR rates will be recalculated based on the defined percentage value.

The rate and percent-rate commands override one another. If the current rate for a queue is defined using the percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime.

An egress port queue group queue rate override may be expressed as either a percentage or an explicit rate independent on how the queue's template rate is expressed.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the **percent-rate** is performed under the **hs-wrr-group** within the SAP egress QoS policy.

The **no** form of this command returns the queue to its default shaping rate and cir rate. When **no percent-rate** is defined within a port egress queue group queue override, the queue reverts to the defined shaping and CIR rates within the egress queue group template associated with the queue.

Parameters

pir-percent

Specifies the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

Values 0.01 to 100.00

Default 100.00

cir-percent

Specifies the queue's committed scheduling rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

Values 0.00 to 100.00

Default 100.00

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*]

no percent-rate

Context

[Tree] (config>service>vprn>if>sap>egress>policer-override>plcr percent-rate)

[Tree] (config>service>vprn>if>sap>ingress>policer-override>plcr percent-rate)

Full Context

configure service vprn interface sap egress policer-override policer percent-rate

configure service vprn interface sap ingress policer-override policer percent-rate

Description

This command configures the percent rates (CIR and PIR) override and can only be used when the rate for the associated policer in the applied SAP ingress QoS policy is also configured with the **percent-rate** command.

The **no** form of this command removes the **percent-rate** override so that the **percent-rate** configured for the policer in the applied SAP egress QoS policy is used.

Parameters

pir-percent

Specifies the policer's PIR as a percentage of the policers's parent arbiter rate.

Values 0.01 to 100.00

Default 100.00

cir-percent

Specifies the policer's CIR as a percentage of the policers's parent arbiter rate.

Values 0.00 to 100.00

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*]

no percent-rate

Context

[Tree] (config>service>vprn>if>sap>egress>queue-override>queue percent-rate)

Full Context

configure service vprn interface sap egress queue-override queue percent-rate

Description

The **percent-rate** command supports a queue's shaping rate and CIR rate as a percentage of the egress port's line rate. When the rates are expressed as a percentage within the template, the actual rate used per instance of the queue group queue-id will vary based on the port speed. For example, when the same template is used to create a queue group on a 1-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit port due to the difference in port speeds. This enables the same template to be used on multiple ports without needing to use port based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's shaping and CIR rates will be recalculated based on the defined percentage value.

The rate and percent-rate commands override one another. If the current rate for a queue is defined using the percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime.

An egress port queue group queue rate override may be expressed as either a percentage or an explicit rate independent on how the queue's template rate is expressed.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the **percent-rate** is performed under the **hs-wrr-group** within the SAP egress QoS policy.

The **no** form of this command returns the queue to its default shaping rate and cir rate. When **no percent-rate** is defined within a port egress queue group queue override, the queue reverts to the defined shaping and CIR rates within the egress queue group template associated with the queue.

Parameters

pir-percent

Specifies the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

Values 0.01 to 100.00

Default 100.00

cir-percent

Specifies the queue's committed scheduling rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

Values 0.00 to 100.00

Default 100.00

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *percentage* [**local-limit** | **reference-port-limit**]

no percent-rate

Context

[Tree] (config>qos>plcr-ctrl-plcy>tier>arbiter percent-rate)

Full Context

configure qos policer-control-policy tier arbiter percent-rate

Description

This command configures the percent rate of this contexts policer policy.

The **no** form of this command removes the configuration.

Parameters

percentage

Specifies the percentage.

Values 0.01 to 100.00

local-limit

Keyword used to specify the local limit.

reference-port-limit

Keyword used to specify the reference port limit.

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*] [**local-limit** | **reference-port-limit**]

no percent-rate

Context

[Tree] (config>qos>sap-egress>policer percent-rate)

[Tree] (config>qos>sap-ingress>policer percent-rate)

Full Context

configure qos sap-egress policer percent-rate

configure qos sap-ingress policer percent-rate

Description

The percent-rate command within the SAP ingress and egress QoS policies enables supports for a policer's PIR and CIR rate to be configured as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

This enables the same QoS policy to be used on SAPs on different FPs without needing to use SAP-based policer overrides to modify a policer's rate to get the same relative performance from the policer.

If the parent arbiter rate changes after the policer is created, the policer's PIR and CIR rates will be recalculated based on the defined percentage value.

The rate and percent-rate commands override one another. If the current rate for a policer is defined using the percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A policer's rate may dynamically be changed back and forth from a percentage to an explicit rate at any time.

The **no** form of this command returns the queue to its default shaping rate and CIR rate.

Parameters

pir-percent

Specifies the policer's PIR as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

Values Percentage ranging from 0.01 to 100.00

Default 100.00

cir-percent

The **cir** keyword is optional and, when defined, the required cir-percent CIR parameter expresses the policer's CIR as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

Values Percentage ranging from 0.00 to 100.00

Default 100.00

local-limit

Keyword used to specify the local limit.

reference-port-limit

Keyword used to specify the reference port limit.

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*] [**fir** *fir-percent*] [{ **port-limit** | **local-limit** | **reference-port-limit**}]

percent-rate *pir-percent* **police** [{**port-limit** | **local-limit** | **reference-port-limit**}]

no percent-rate

Context

[Tree] (config>qos>sap-ingress>queue percent-rate)

Full Context

configure qos sap-ingress queue percent-rate

Description

This command configures a queue's PIR and CIR as a percentage of the ingress port line rate or as a percentage of its parent scheduler rate. When the rates are expressed as a **port-limit**, the actual rates used per instance of the queue will vary based on the port speed. For example, when the same QoS policy is used on a 1 Gb and a 10 Gb Ethernet port, the queue's rates will be 10 times greater on the 10 Gb port due to the difference in port speeds. This enables the same QoS policy to be used on SAPs on different ports without needing to use SAP-based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port’s speed changes after the queue is created, the queue’s PIR, CIR, and FIR rates will be recalculated based on the defined percentage value.

When the rates are expressed as a **local-limit**, the actual rates used per instance of the queue are relative to the queue's parent scheduler rate. This enables the same QoS policy to be used on SAPs with different parent scheduler rates without needing to use SAP-based queue overrides to modify a queue's rate to get the same relative performance from the queue. If the parent scheduler rate changes after the queue is created, the queue's PIR, CIR, and FIR will be recalculated based on the defined percentage value.

The **rate** and **percent-rate** commands override one another. If the current rate for a queue is defined using the **percent-rate** command and the **rate** command is executed, the **percent-rate** values are deleted. Similarly, the **percent-rate** command causes any **rate** command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at any time.

Queue rate overrides can only be specified in the form configured in the QoS policy (for example, a SAP override can only be specified as a **percent-rate** if the associated QoS policy was also defined as percent-rate). Likewise, a SAP override can only be specified as a **rate** (kb/s) if the associated QoS policy was also defined as a rate. Queue-overrides are relative to the limit type specified in the QoS policy.

When no **percent-rate** is defined within a SAP ingress queue-override, the queue uses the defined shaping rate, CIR, and FIR within the SAP ingress QoS policy associated with the queue.

The **no** form of this command returns the queue to its default shaping rate, CIR, and FIR.

Parameters

pir-percent

Specifies the queue's PIR as a percentage dependent on the use of the **port-limit** or **local-limit**.

Values 0.01 to 100.00

Default 100.00

cir-percent

Specifies the queue's CIR as a percentage dependent on the use of the **port-limit** or **local-limit**.

Values 0.00 to 100.00

Default 100.00

fir-percent

Specifies the queue's FIR as a percentage dependent on the use of the **port-limit** or **local-limit**. FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.

Values 0.00 to 100.00

Default 100.00

police

Keyword used to specify that traffic feeding into the physical queue instance above the specified PIR rate is dropped. When the **police** keyword is defined, only the PIR rate may be overridden. The **police** keyword is only applicable to SAP ingress.

port-limit

Keyword used to specify that the configured PIR, CIR, and FIR percentages are relative to the rate of the port (including the ingress-rate setting) to which the queue is attached. The **port-limit** is the default.

local-limit

Keyword used to specify that the configured PIR, CIR, and FIR percentages are relative to the queue's parent scheduler rate. If there is no parent scheduler rate, or its rate is **max**, the **port-limit** is used.

reference-port-limit

Keyword used to specify that the configured PIR, CIR, and FIR percentages are relative to the rate of the reference port (including the ingress-rate setting) to which the queue is attached.

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*] [{**port-limit** | **local-limit** | **reference-port-limit**}]

no percent-rate

Context

[Tree] (config>qos>sap-egress>queue percent-rate)

Full Context

configure qos sap-egress queue percent-rate

Description

This command configures a queue's PIR and CIR as a percentage of the egress port line rate or as a percentage of its parent scheduler rate or **agg-rate** rate. When the rates are expressed as a **port-limit**, the actual rates used per instance of the queue will vary based on the port speed. For example, when the same QoS policy is used on a 1 Gb and a 10 Gb Ethernet port, the queue's rates will be 10 times greater on the 10 Gb port due to the difference in port speeds. This enables the same QoS policy to be used on SAPs on different ports without needing to use SAP-based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's PIR and CIR will be recalculated based on the defined percentage value.

When the rates are expressed as a **local-limit**, the actual rates used per instance of the queue are relative to the queue's parent scheduler rate or **agg-rate** rate. This enables the same QoS policy to be used on SAPs with different parent scheduler rates without needing to use SAP-based queue overrides to modify a queue's rate to get the same relative performance from the queue. If the parent scheduler rate changes after the queue is created, the queue's PIR and CIR will be recalculated based on the defined percentage value.

The **rate** and **percent-rate** commands override one another. If the current rate for a queue is defined using the **percent-rate** command and the **rate** command is executed, the **percent-rate** values are deleted. Similarly, the **percent-rate** command causes any **rate** command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at any time.

Queue rate overrides can only be specified in the form as configured in the QoS policy (a SAP override can only be specified as a **percent-rate** if the associated QoS policy was also defined as percent-rate). Likewise, a SAP override can only be specified as a rate (kb/s) if the associated QoS policy was also defined as a rate. Queue-overrides are relative to the limit type specified in the QoS policy.

When configured on an egress HSQ queue group queue, the **cir** keyword is ignored.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the **percent-rate** is performed under the **hs-wrr-group** within the SAP egress QoS policy.

When no **percent-rate** is defined within a SAP egress queue-override, the queue uses the defined shaping rate and CIR within the SAP egress QoS policy associated with the queue.

The **no** form of this command returns the queue to its default shaping rate and CIR.

Parameters

pir-percent

Specifies the queue's PIR as a percentage dependent on the use of the **port-limit** or **local-limit**.

Values 0.01 to 100.00

Default 100.00

cir-percent

Specifies the queue's CIR as a percentage dependent on the use of the **port-limit** or **local-limit.**

Values 0.00 to 100.00

Default 100.00

port-limit

Keyword used to specify that the configure PIR and CIR percentages are relative to the rate of the port (including the **egress-rate** setting) to which this queue connects. The **port-limit** is the default.

local-limit

Keyword used to specify that the configure PIR and CIR percentages are relative to the rate of the queue's parent scheduler **rate** or **agg-rate** rate at egress. If there is no parent scheduler rate or **agg-rate** rate, or those rates are **max**, the **port-limit** is used.

reference-port-limit

Keyword used to specify that the configure PIR and CIR percentages are relative to the rate of the reference port (including the **egress-rate** setting) to which this queue connects.

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percentage*] [**local-limit** | **reference-port-limit**]

no percent-rate

Context

[Tree] (config>qos>qgrps>egr>qgrp>policer percent-rate)

Full Context

configure qos queue-group-templates egress queue-group policer percent-rate

Description

This command configures the percent rate for this contexts policer.

The **no** form of this command removes the configuration.

Parameters

pir-percent

Specifies the policer's PIR as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

Values 0.01 to 100.00

cir-percent

The **cir** keyword is optional and, when defined, the required cir-percent CIR parameter expresses the policer's CIR as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

Values 0.00 to 100.00, sum

local-limit

Keyword used to specify the local limit.

reference-port-limit

Keyword used to specify the reference port limit.

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percentage*]

no percent-rate

Context

[Tree] (config>qos>qgrps>ing>qgrp>policer percent-rate)

Full Context

configure qos queue-group-templates ingress queue-group policer percent-rate

Description

This command configures the percent rate for this contexts policer.

The **no** form of this command removes the configuration.

Parameters

pir-percent

Specifies the policer's PIR as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

Values 0.01 to 100.00

cir-percent

The **cir** keyword is optional and, when defined, the required cir-percent CIR parameter expresses the policer's CIR as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

Values 0.00 to 100.00, sum

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percent*] [{**port-limit** | **local-limit** | **reference-port-limit**}]

no percent-rate

Context

[Tree] (config>qos>qgrps>egr>queue-group>queue percent-rate)

Full Context

configure qos queue-group-templates egress queue-group queue percent-rate

Description

The **percent-rate** command within the egress queue group template enables support for a queue's PIR and CIR rate to be configured as a percentage of the egress port's line rate. When the rates are expressed as a percentage within the template, the actual rate used per instance of the queue group *queue-id* will vary based on the port speed. For example, when the same template is used to create a queue group on a 1-Gb and a 10-Gb Ethernet port, the queue's rates will be 10 times greater on the 10 Gb port due to the difference in port speeds. This enables the same template to be used on multiple ports without needing to use port-based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's shaping and CIR rates will be recalculated based on the defined percentage value.

When configured on an egress HSQ queue group queue, the **cir** keyword is ignored.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case the configuration of the rate is performed under the **hs-wrr-group** within the egress queue group template.

The **rate** and **percent-rate** commands override one another. If the current rate for a queue is defined using the **percent-rate** command and the **rate** command is executed, the **percent-rate** values are deleted. Similarly, the **percent-rate** command causes any **rate** command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at any time.

The **no** form of this command returns the queue to its default shaping rate and CIR rate.

Parameters

pir-percent

Expresses the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation and the egress rate setting.

Values 0.01 to 100.00 percent

Default 100.0

cir-percent

The **cir** keyword is optional and when defined, the required *pir-percent* parameter expresses the queue's committed scheduling rate as a percentage of line rate. The line rate associated with the queue's port may change dynamically due to configuration or auto-negotiation and the egress rate setting.

Values 0.01 to 100.00 percent

Default 100.0

port-limit

Keyword used to specify that the configure PIR and CIR percentages are relative to the rate of the port (including the **egress-rate** setting) to which this queue connects. The **port-limit** is the default.

local-limit

Keyword used to specify that the configure PIR and CIR percentages are relative to the rate of the queue's parent scheduler **rate** or **agg-rate** rate at egress. If there is no parent scheduler rate or **agg-rate** rate, or those rates are **max**, the **port-limit** is used.

reference-port-limit

Keyword used to specify that the configure PIR and CIR percentages are relative to the rate of the reference port (including the **egress-rate** setting) to which this queue connects.

Platforms

7705 SAR Gen 2

percent-rate

Syntax

percent-rate *pir-percent* [**cir** *cir-percentage*] [**local-limit** | **reference-port-limit**]

no percent-rate

Context

[Tree] (config>qos>scheduler-policy>tier>scheduler percent-rate)

Full Context

configure qos scheduler-policy tier scheduler percent-rate

Description

This command configures the percentage rate for the scheduler policy.

The **no** form of this command removes the configuration.

Parameters

pir-percent

Specifies the policer's PIR as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

Values 0.01 to 100.00

cir-percent

The **cir** keyword is optional and, when defined, the required *cir-percent* CIR parameter expresses the policer's CIR as a percentage of the immediate parent root policer/arbiter rate or the FP capacity.

Values 0.00 to 100.00, sum

local-limit

Keyword used to specify the local limit.

reference-port-limit

Keyword used to specify the reference port limit.

Platforms

7705 SAR Gen 2

21.57 percent-reduction-from-mbs

percent-reduction-from-mbs

Syntax

percent-reduction-from-mbs *percent*

no percent-reduction-from-mbs

Context

[Tree] (config>service>ies>if>sap>egress>queue-override>queue>drop-tail>low percent-reduction-from-mbs)

[Tree] (config>service>ies>if>sap>ingress>queue-override>queue>drop-tail>low percent-reduction-from-mbs)

[Tree] (config>service>vpls>sap>ingress>queue-override>queue>drop-tail>low percent-reduction-from-mbs)

[Tree] (config>service>vpls>sap>egress>queue-override>queue>drop-tail>low percent-reduction-from-mbs)

Full Context

configure service ies interface sap egress queue-override queue drop-tail low percent-reduction-from-mbs

configure service ies interface sap ingress queue-override queue drop-tail low percent-reduction-from-mbs

configure service vpls sap ingress queue-override queue drop-tail low percent-reduction-from-mbs

configure service vpls sap egress queue-override queue drop-tail low percent-reduction-from-mbs

Description

This command overrides the low queue drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and this percentage is configured to be 30% for the low drop tail, then the low drop tail is at 420 kbytes and out-of-profile packets are not accepted into the queue if its depth is greater than this value, and discarded.

Parameters

percent

Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, default

Platforms

7705 SAR Gen 2

percent-reduction-from-mbs

Syntax

percent-reduction-from-mbs *percent*

no percent-reduction-from-mbs

Context

[Tree] (config>port>ethernet>network>egr>qgrp>qover>q>drop-tail>low percent-reduction-from-mbs)

[Tree] (config>port>ethernet>access>egr>qgrp>qover>q>drop-tail>low percent-reduction-from-mbs)

[Tree] (config>port>ethernet>access>ing>qgrp>qover>q>drop-tail>low percent-reduction-from-mbs)

Full Context

configure port ethernet network egress queue-group queue-overrides queue drop-tail low percent-reduction-from-mbs

configure port ethernet access egress queue-group queue-overrides queue drop-tail low percent-reduction-from-mbs

configure port ethernet access ingress queue-group queue-overrides queue drop-tail low percent-reduction-from-mbs

Description

This command overrides the low queue drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and this percentage is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets will not be accepted into the queue if its depth is greater than this value, and so will be discarded.

Parameters

percent

Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, default

Platforms

7705 SAR Gen 2

percent-reduction-from-mbs

Syntax

percent-reduction-from-mbs *percent*

no percent-reduction-from-mbs

Context

[Tree] (config>service>epipe>sap>egress>queue-override>queue>drop-tail>low percent-reduction-from-mbs)

[Tree] (config>service>epipe>sap>ingress>queue-override>queue>drop-tail>low percent-reduction-from-mbs)

Full Context

configure service epipe sap egress queue-override queue drop-tail low percent-reduction-from-mbs

configure service epipe sap ingress queue-override queue drop-tail low percent-reduction-from-mbs

Description

This command overrides the low queue drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and the percentage reduction is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes. Any out-of-profile packets will not be accepted into the queue if its depth is greater than this value, and so will be discarded.

Parameters

percent

Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, default

Platforms

7705 SAR Gen 2

percent-reduction-from-mbs

Syntax

percent-reduction-from-mbs *percent*

no percent-reduction-from-mbs

Context

[Tree] (config>service>vprn>if>sap>egress>queue-override>queue>drop-tail>low percent-reduction-from-mbs)

[Tree] (config>service>vprn>if>sap>ingress>queue-override>queue>drop-tail>low percent-reduction-from-mbs)

Full Context

configure service vprn interface sap egress queue-override queue drop-tail low percent-reduction-from-mbs

configure service vprn interface sap ingress queue-override queue drop-tail low percent-reduction-from-mbs

Description

This command overrides the low queue drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and the percentage reduction is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets will not be accepted into the queue if its depth is greater than this value, and so will be discarded.

Parameters

percent

Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, default

Platforms

7705 SAR Gen 2

percent-reduction-from-mbs

Syntax

percent-reduction-from-mbs *percent*

no percent-reduction-from-mbs

Context

[Tree] (config>qos>sap-ingress>queue>drop-tail>low percent-reduction-from-mbs)

Full Context

configure qos sap-ingress queue drop-tail low percent-reduction-from-mbs

Description

This command configures the ingress SAP low drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and this percentage is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets will not be accepted into the queue if its depth is greater than this value and will be discarded.

Default

percent-reduction-from-mbs 10

Parameters

percent

Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, default

Platforms

7705 SAR Gen 2

percent-reduction-from-mbs

Syntax

percent-reduction-from-mbs *percent*

no percent-reduction-from-mbs

Context

[Tree] (config>qos>sap-egress>queue>drop-tail>exceed percent-reduction-from-mbs)

[Tree] (config>qos>sap-egress>queue>drop-tail>highplus percent-reduction-from-mbs)

[Tree] (config>qos>sap-egress>queue>drop-tail>low percent-reduction-from-mbs)

[Tree] (config>qos>sap-egress>queue>drop-tail>high percent-reduction-from-mbs)

Full Context

configure qos sap-egress queue drop-tail exceed percent-reduction-from-mbs

configure qos sap-egress queue drop-tail highplus percent-reduction-from-mbs

configure qos sap-egress queue drop-tail low percent-reduction-from-mbs

configure qos sap-egress queue drop-tail high percent-reduction-from-mbs

Description

This command configures the egress SAP queue drop tails as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and this percentage is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets will not be accepted into the queue if its depth is greater than this value and will be discarded.

The drop tails apply to packets with the following profile state:

- Exceed drop tail: exceed-profile
- High drop tail: in-profile
- Highplus drop tail: inplus-profile
- Low drop tail: out-of-profile

Default

Exceed drop tail: 20%

Low drop tail: 10%

High drop tail: 0%

Highplus drop tail: 0%

Parameters

percent

Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, default

Platforms

7705 SAR Gen 2

percent-reduction-from-mbs

Syntax

percent-reduction-from-mbs *percent*

no percent-reduction-from-mbs

Context

[Tree] (config>qos>network-queue>queue>drop-tail>low percent-reduction-from-mbs)

Full Context

configure qos network-queue queue drop-tail low percent-reduction-from-mbs

Description

This command configures the ingress and egress network queue low drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and **percent-reduction-from-mbs** is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets will not be accepted into the queue if its depth is greater than this value and will be discarded.

The exceed drop tail is not configurable for network queues, however, it is set to a value of 10% in addition to low drop tail and capped by the MBS.

Default

percent-reduction-from-mbs 10

Parameters

percent

Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, default

Platforms

7705 SAR Gen 2

percent-reduction-from-mbs

Syntax

percent-reduction-from-mbs *percent*

no percent-reduction-from-mbs

Context

[Tree] (config>qos>qgrps>ing>qgrp>queue>drop-tail>low percent-reduction-from-mbs)

Full Context

configure qos queue-group-templates ingress queue-group queue drop-tail low percent-reduction-from-mbs

Description

This command configures the ingress queue group queue low drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and the percentage reduction is configured to be 30% for the low drop tail, the low drop tail will be at 420 kbytes. Out-of-profile packets will not be accepted into the queue and will be discarded if the queue depth is greater than this value.

Default

10%

Parameters

percent

Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, **default**

Platforms

7705 SAR Gen 2

percent-reduction-from-mbs

Syntax

percent-reduction-from-mbs *percent*

no percent-reduction-from-mbs

Context

[Tree] (config>qos>qgrps>egr>qgrp>queue>drop-tail>low percent-reduction-from-mbs)

[Tree] (config>qos>qgrps>egr>qgrp>queue>drop-tail>highplus percent-reduction-from-mbs)

[Tree] (config>qos>qgrps>egr>qgrp>queue>drop-tail>high percent-reduction-from-mbs)

[Tree] (config>qos>qgrps>egr>qgrp>queue>drop-tail>exceed percent-reduction-from-mbs)

Full Context

configure qos queue-group-templates egress queue-group queue drop-tail low percent-reduction-from-mbs

configure qos queue-group-templates egress queue-group queue drop-tail highplus percent-reduction-from-mbs

configure qos queue-group-templates egress queue-group queue drop-tail high percent-reduction-from-mbs

configure qos queue-group-templates egress queue-group queue drop-tail exceed percent-reduction-from-mbs

Description

This command configures the egress queue group queue drop tails as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and the percentage reduction is configured to be 30% for the low drop tail, the low drop tail will be at 420 kbytes. Out-of-profile packets will not be accepted into the queue and will be discarded if the queue depth is greater than this value.

The drop tails apply to packets with the following profile states:

- exceed drop tail: exceed-profile
- high drop tail: in-profile

- highplus drop tail: inplus-profile
- low drop tail: out-of-profile

Default

exceed drop tail: 20%

low drop tail: 10%

high drop tail: 0%

highplus drop tail: 0%

Parameters

percent

Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, default

Platforms

7705 SAR Gen 2

21.58 period

period

Syntax

period *milli-seconds*

no period

Context

[Tree] (config>router>rsvp>msg-pacing period)

Full Context

configure router rsvp msg-pacing period

Description

This command specifies the time interval (in ms), when the router can send the specified number of RSVP messages which is specified in the **max-burst** command.

Default

period 100

Parameters

milli-seconds

Specifies the time interval in increments of 10 ms.

Values 100 to 1000

Platforms

7705 SAR Gen 2

21.59 periodic-update-interval

periodic-update-interval

Syntax

periodic-update-interval [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

Context

[Tree] (config>system>security>pki>ca-prof>auto-crl-update periodic-update-interval)

Full Context

configure system security pki ca-profile auto-crl-update periodic-update-interval

Description

This command specifies the interval for periodic updates. The minimal interval is 1 hour. The maximum interval is 366 days.

Default

periodic-update-interval days 1

Parameters

days *days*

Specifies the number of days for periodic updates.

Values 0 to 366

hours

Specifies the number of hours for periodic updates.

Values 0 to 23

minutes

Specifies the number of minutes for periodic updates.

Values 0 to 59

seconds

Specifies the number of seconds for periodic updates.

Values 0 to 59

Platforms

7705 SAR Gen 2

21.60 permit-empty-passwords

permit-empty-passwords

Syntax

[**no**] **configure system security ssh permit-empty-passwords**

Context

[Tree] (config>system>security>ssh permit-empty-passwords)

Full Context

configure system security ssh permit-empty-passwords

Description

This command configures the permission of users with empty password strings to log in.

The **no** form of this command prevents users with empty password strings from logging in.

Default

permit-empty-passwords

Platforms

7705 SAR Gen 2

21.61 persist

persist

Syntax

persist {**on** | **off**}

Context

[Tree] (bof persist)

Full Context

bof persist

Description

This command specifies whether the system will preserve system indexes when a **save** command is executed in classic or mixed configuration mode. During a subsequent boot, the index file is read along with the configuration file. As a result, a number of system indexes are preserved between reboots, including the interface index, LSP IDs, path IDs, and so on. This reduces resynchronizations of the Network Management System (NMS) with the affected network element.

This command is ignored in model-driven configuration mode. In model-driven mode, system indices are always saved and they are embedded in the configuration file.

In the event that persist is **on** and the reboot with the appropriate index file fails in classic or mixed configuration mode, SNMP is operationally shut down to prevent the management system from accessing and possibly synchronizing with a partially booted or incomplete network element. To enable SNMP access, enter the **config>system>snmp>no shutdown** command.

If **persist** is enabled and the **admin save** *url* command is executed with an FTP path used as the *url* parameter, two FTP sessions simultaneously open to the FTP server. The FTP server must be configured to allow multiple sessions from the same login, otherwise, the configuration and index files will not be saved correctly.

Note:

- In classic or mixed configuration mode, persistency files (.ndx) are saved on the same disk as the configuration files and the image files.
- When an operator sets the location for the persistency file in classic or mixed configuration mode, the system will check to ensure that the disk has enough free space. If there is not enough free space, the persistency will not become active and a trap will be generated. Then, it is up to the operator to free adequate disk space. In the meantime, the system will perform a space availability check every 30 seconds. As soon as the space is available the persistency will become active on the next (30 second) check.

Default

persist off

Parameters

on

Enables the system index saves between reboots.

off

Disables the system index saves between reboots.

Platforms

7705 SAR Gen 2

21.62 persistence

persistence

Syntax

persistence

Context

[Tree] (config>system persistence)

Full Context

configure system persistence

Description

Commands in this context configure persistence parameters on the system.

The persistence feature enables state information learned through applications such as subscriber management, DHCP server, or application assurance to be retained across reboots.

Platforms

7705 SAR Gen 2

persistence

Syntax

persistence [*persistence-client*]

no persistence

Context

[Tree] (debug>system persistence)

Full Context

debug system persistence

Description

This command displays persistence debug information.

Parameters

persistence-client

Displays persistence debug information.

Values		
	ancp	ANCP
	application-assurance	application-assurance
	dhcp-server	local DHCP server
	nat-fwds	NAT port forwarding
	python-policy-cache	Python Cache
	submgt	subscriber management

Platforms
7705 SAR Gen 2

21.63 persistent-subscriptions

persistent-subscriptions

Syntax
persistent-subscriptions

Context
[Tree] (config>system>telemetry persistent-subscriptions)

Full Context
configure system telemetry persistent-subscriptions

Description
Commands in this context configure persistent subscriptions.

Platforms
7705 SAR Gen 2

21.64 pfs

pfs

Syntax
pfs [**dh-group** {**1** | **2** | **5** | **14** | **15** | **19** | **20** | **21**}]

no pfs

Context

[Tree] (config>ipsec>ike-policy pfs)

Full Context

configure ipsec ike-policy pfs

Description

This command enables perfect forward secrecy on the IPsec tunnel using this policy. PFS provides for a new Diffie-Hellman key exchange each time the SA key is renegotiated. After that SA expires, the key is forgotten and another key is generated (if the SA remains up). This means that an attacker who cracks part of the exchange can only read the part that used the key before the key changed. There is no advantage in cracking the other parts if they attacker has already cracked one.

The **no** form of this command disables PFS. If this it turned off during an active SA, when the SA expires and it is time to re-key the session, the original Diffie-Hellman primes will be used to generate the new keys.

Default

no pfs

Parameters

dh-group {1 | 2 | 5 | 14 | 15 | 19 | 20 | 21}

Specifies which Diffie-Hellman group to use for calculating session keys. More bits provide a higher level of security, but require more processing. Three groups are supported with IKE-v1:

Group 1: 768 bits

Group 2: 1024 bits

Group 5: 1536 bits

Group 14: 2048 bits

Group 15: 3072 bits

Group 19: P-256 ECC Curve, 256 bits

Group 20: P-384 ECC Curve, 384 bits

Group 21: P-512 ECC Curve, 512 bits

Platforms

7705 SAR Gen 2

21.65 pfs-dh-group

pfs-dh-group

Syntax

pfs-dh-group {**1** | **2** | **5** | **14** | **15** | **19** | **20** | **21**}

pfs-dh-group inherit

no pfs-dh-group

Context

[Tree] (config>ipsec>ipsec-transform pfs-dh-group)

Full Context

configure ipsec ipsec-transform pfs-dh-group

Description

This command specifies the Diffie-Hellman group to be used for Perfect Forward Secrecy (PFS) computation during CHILD_SA rekeying.

The **no** form of this command reverts to the default.

Default

pfs-dh-group inherit

Parameters

{1 | 2 | 5 | 14 | 15 | 19 | 20 | 21}

Specifies the Diffie-Hellman group to achieve PFS.

inherit

Specifies that the value of the DH group used by the system is inherited from the IPsec gateway or IPsec tunnel.

Platforms

7705 SAR Gen 2

21.66 phone

phone

Syntax

[**no**] **phone** *phone-number*

Context

[Tree] (config>service>cust phone)

Full Context

configure service customer phone

Description

This command adds telephone number information for a customer ID. The **no** form of this command removes the phone number value from the customer ID.

Parameters

string

Specifies the customer phone number entered as an ASCII string up to 80 characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Any printable, seven bit ASCII characters may be used within the string.

Platforms

7705 SAR Gen 2

21.67 pim

pim

Syntax

[**no**] **pim**

Context

[Tree] (config>service>vprn pim)

Full Context

configure service vprn pim

Description

This command configures a Protocol Independent Multicast (PIM) instance in the VPRN service. When an PIM instance is created, the protocol is enabled. PIM is used for multicast routing within the network. Devices in the network can receive the multicast feed requested and non-participating routers can be pruned. The router supports PIM sparse mode (PIM-SM).

The **no** form of this command deletes the PIM protocol instance removing all associated configuration parameters.

Platforms

7705 SAR Gen 2

pim

Syntax

[**no**] **pim**

Context

[Tree] (config>router pim)

Full Context

configure router pim

Description

This command enables a Protocol Independent Multicast (PIM) instance.

PIM is used for multicast routing within the network. Devices in the network can receive the multicast feed requested and non-participating routers can be pruned. The router OS supports PIM sparse mode (PIM-SM).

The **no** form of this command disables the PIM instance.

Default

no pim

Platforms

7705 SAR Gen 2

22 p Commands – Part II

22.1 pim-ssm-scaling

pim-ssm-scaling

Syntax

[**no**] **pim-ssm-scaling**

Context

[Tree] (config>router>pim pim-ssm-scaling)

Full Context

configure router pim pim-ssm-scaling

Description

This command enables an increase of PIM SSM (S,G) scaling to a maximum of 256kper system. The per-complex (FP) multicast scaling limit is still in place, but multiple complexes can be used to achieve the 256k per-system (S,G) scaling.

The **no** form of this command disables the increase in PIM SSM scaling.

Default

no pim-ssm-scaling

Platforms

7705 SAR Gen 2

22.2 ping

ping

Syntax

ping *ip-address* | *dns-name* [**bypass-routing** | {**interface** *interface-name*} | {**next-hop** *ip-address*}] [{**router** *router-or-service*} | {**router-instance** *router-instance*} | {**service-name** *service-name*}] [**source** *ip-address*] [**count** *requests*] [**detail** | **rapid**] [**do-not-fragment**] [**fc** *fc-name*] [**interval centisecs** | **secs**] [**pattern** *pattern*] [**size** *bytes*] [**timeout** *timeout*] [**tos** *type-of-service*] [**ttl** *time-to-live*]

ping *ipv4-address* **subscriber-id** *sub-ident-string* [{**router** *router-or-service*} | {**router-instance** *router-instance*} | {**service-name** *service-name*}] [**source** *ip-address*] [**count** *requests*] [**detail** | **rapid**] [**do-not-fragment**] [**fc** *fc-name*] [**interval centisecs** | **secs**] [**pattern** *pattern*] [**size** *bytes*] [**timeout** *timeout*] [**tos** *type-of-service*] [**ttl** *time-to-live*]

ping srv6-policy color *color* **endpoint** *ipv6-address* [**segment-list** *segment-list*] [**candidate-path protocol-owner static** | **bgp** [**preference** *preference*] [**distinguisher** *distinguisher*]] [**count** *requests*] [**detail** | **rapid**] [**do-not-fragment**] [**fc** *fc-name*] [**interval centisecs** | **secs**] [**pattern** *pattern*] [**size** *bytes*] [**timeout** *timeout*] [**tos** *type-of-service*] [**ttl** *time-to-live*]

Context

[Tree] (ping)

Full Context

ping

Description

This command sends a ping to a destination to verify IP reachability.

Use the **ping** {*ip-address* | *dns-name*} [{**bypass-routing** | {**interface** *interface-name*} | {**next-hop** *ip-address*}}] command syntax to send a generic ping. This is the basic TCP/IP utility to verify IP reachability.

Use the **ping** *ipv4-address* **subscriber-id** *sub-ident-string* command syntax to send a ping to verify L2-Aware remote host reachability.

The L2-Aware form of this command can be initiated from the gateway IPv4 address in the inside routing context or from any IPv4 address in the outside routing context. If the gateway IPv4 address is used as the source address, it must be explicitly specified in conjunction with the L2-Aware syntax.

Any source address can be used for the ping to test the relevant NAT policy. If the source address refers to a policy that is not configured on the router, the message "MINOR: OAM #2160 router ID is not an outside router for this subscriber" is displayed. The source address does not need to belong to the system.

If the outside routing context is not specified, by default, the Base router is selected. If the specified or default Base router instance is not the outside routing context for the subscriber, the L2-Aware form of this command fails to execute, and the message "MINOR: OAM #2160 router ID is not an outside router for this subscriber" is displayed.

The NAT application shares query IDs between L2-Aware pings and ICMP or GRE traffic that has undergone NAT and is destined for a DMZ host. If there is query ID space exhaustion, ICMP or GRE flows destined for DMZ hosts are deleted, so their query IDs can be reused for the requested L2-Aware pings.

Use the **ping srv6-policy color** *color* **endpoint** *ipv6-address* command syntax to launch a ping for an SRv6 policy matching a specific color and endpoint. The ping probe may optionally be targeted at a specific segment list of the SRv6 policy. When the segment list is not specified, the ping probe is sent on the lowest available segment list.

Parameters

bypass-routing

Specifies whether to send the ping request to a host on a directly attached network, bypassing the routing table.

bytes

Specifies the request packet size in bytes, expressed as a decimal integer.

Values 0 to 16384

Default 56

candidate-path

Specifies a candidate path of the SRv6 policy to ping. The candidate path does not need to be the currently active candidate path.

centiseconds

Sets the interval in centiseconds.

Values 1 to 10000 centiseconds if **rapid** is selected.

Default 1 centisecond if **rapid** is selected.

detail

Displays detailed information.

distinguisher

Specifies the distinguisher of the SRv6 policy candidate path to send the ping probe on. This parameter must be configured if **protocol-owner** is configured to **bgp**.

Values 1 to 4294967295

do-not-fragment

Sets the DF (Do Not Fragment) bit in the ICMP ping packet (does not apply to ICMPv6).

dns-name

Specifies the DNS name of the far-end device on which to send the **svc-ping** request message, expressed as a character string.

fc-name

Specifies the forwarding class of the MPLS echo request packets.

Values be, l2, af, l1, h2, ef, h1, nc

Default nc

interface-name

Specifies the name of an IP interface. The name must already exist in the **configure router interface** context.

ip-address

Specifies the far-end IP address, in dotted decimal notation, on which to send the **svc-ping** request message.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x[-*interface*]
	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x: [0 to FFFF]H

d: [0 to 255]D

interface: up to 32 characters, mandatory for link local addresses

next-hop *ip-address*

Displays only static routes with the specified next hop IP address.

Values

ipv4-address: a.b.c.d (host bits must be 0)

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]

pattern

Specifies that the date portion in a ping packet that is filled with the pattern value specified. If not specified, position information is filled instead.

Values 0 to 65535

Default system-generated sequential pattern

preference

Specifies the preference of the SRv6 policy candidate path to send the ping probe on.

Values 0 to 4294967295

Default 100

protocol-owner

Specifies the protocol owner of the SRv6 policy candidate path to ping.

Values bgp — Specifies a BGP SRv6 policy.

static — Specifies a locally configured static SRv6 policy.

rapid

Specifies that packets be generated as fast as possible instead of the default 1 per second. Changes the units for the **interval** command from seconds to centiseconds.

requests

Specifies the number of times to perform an OAM ping probe operation. Each OAM echo message request must either time out or receive a reply before the next message request is sent.

Values 1 to 100000

Default 5

router-instance

Specifies the preferred method for entering a service name. Stored as the service name. This is the only service-linking function allowed for both mixed-mode and model-driven configuration modes.

Values router-name: Base, management, *cpm-vr-name*, vpls-management

vprn-svc-name: The service name, up to 64 characters

cpm-vr-name: The CPM VR name, up to 32 characters

router-or-service

Specifies the routing instance or service, by number. The *router-instance* parameter is preferred for specifying the router or service.

Values router-name: Base, management, vpls-management

vprn-svc-id: 1 to 2147483647

Default Base

seconds

Overrides the default *timeout* value and is the amount of time that the router waits for a message reply after sending the message request. Upon the expiration of the message time out, the requesting router assumes that the message response is not received. A **request timeout** message is displayed by the CLI for each message request sent that expires. Any response received after the request times out is silently discarded.

Default 5

Values 1 to 10

secs

Sets the interval in seconds.

Values 1 to 1000 seconds if **secs** is selected.

Default 1 second if **secs** is selected.

service-name

Specifies the alias function that allows the service-name to be used, converted, and stored as a service ID.

sub-ident-string

Specifies the L2-Aware NAT subscriber to which ICMP-ping is sent, up to 32 characters. The **subscriber-id** keyword serves as a differentiator between the subscribers with the same IP address in the same routing context (which is allowed in L2-Aware NAT). The **subscriber-id** keyword is mandatory for L2-Aware IPv4 ping, but optional in generic ping framework.

source *ip-address*

Specifies the IP address to be used.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

timeout

Specifies the time out, in seconds.

Values 1 to 10

Default 5

type-of-service

Specifies the service type.

Values 0 to 255

Default 0

time-to-live

Specifies the TTL value for the MPLS label, expressed as a decimal integer.

Values 1 to 128

Default 64

srv6-policy

Keyword to specify that the ping probe is applied to an SRv6 policy.

color-id

Specifies the SRv6 policy color ID.

Values 0 to 4294967295

endpoint *ipv6-address*

Specifies an endpoint as the target of the ping.

Values

ipv6-address:	x:x:x:x:x:x:x:x
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

segment-list

Specifies the segment list to trace.

Values 1 to 32

Platforms

7705 SAR Gen 2

22.3 ping-reply

ping-reply

Syntax

[**no**] **ping-reply**

Context

[Tree] (config>service>ies>if>ipv6>vrrp ping-reply)

Full Context

configure service ies interface ipv6 vrrp ping-reply

Description

This command enables the non-owner master to reply to ICMP echo requests directed at the virtual router instances IP addresses. The ping request can be received on any routed interface.

Ping must not have been disabled at the management security level (either on the parental Ip interface or based on the ping source host address). when ping-reply is not enabled, icmp Echo Requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP echo requests regardless of the setting of ping-reply configuration.

The ping-reply command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the ping-reply command is not executed, ICMP echo requests to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all ICMP echo request messages destined to the non-owner virtual router instance IP addresses.

Default

no ping-reply

Platforms

7705 SAR Gen 2

ping-reply

Syntax

[**no**] **ping-reply**

Context

[Tree] (config>service>ies>if>vrrp ping-reply)

Full Context

configure service ies interface vrrp ping-reply

Description

This command enables the non-owner master to reply to ICMP Echo Requests directed at the virtual router instances IP addresses. The ping request can be received on any routed interface.

Ping must not have been disabled at the management security level (either on the parental IP interface or based on the Ping source host address). When ping-reply is not enabled, ICMP Echo Requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP Echo Requests regardless of the setting of ping-reply configuration.

The ping-reply command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the ping-reply command is not executed, ICMP Echo Requests to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all ICMP Echo Request messages destined to the non-owner virtual router instance IP addresses.

Default

no ping-reply

Platforms

7705 SAR Gen 2

ping-reply

Syntax

[**no**] **ping-reply**

Context

[Tree] (config>service>vprn>if>ipv6>vrrp ping-reply)

[Tree] (config>service>vprn>if>vrrp ping-reply)

Full Context

configure service vprn interface ipv6 vrrp ping-reply

configure service vprn interface vrrp ping-reply

Description

This command enables the non-owner master to reply to ICMP Echo Requests directed at the virtual router instances IP addresses. The ping request can be received on any routed interface.

Ping must not have been disabled at the management security level (either on the parental IP interface or based on the Ping source host address). When ping-reply is not enabled, ICMP Echo Requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP Echo Requests regardless of the setting of ping-reply configuration.

The ping-reply command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the ping-reply command is not executed, ICMP Echo Requests to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all ICMP Echo Request messages destined to the non-owner virtual router instance IP addresses.

Default

no ping-reply

Platforms

7705 SAR Gen 2

ping-reply

Syntax

[**no**] **ping-reply**

Context

[Tree] (config>router>if>vrrp ping-reply)

[Tree] (config>router>if>ipv6>vrrp ping-reply)

Full Context

configure router interface vrrp ping-reply

configure router interface ipv6 vrrp ping-reply

Description

This command enables the non-owner master to reply to ICMP echo requests directed at the virtual router instances IP addresses.

Non-owner virtual router instances are limited by the VRRP specifications to responding to ARP requests destined to the virtual router IP addresses and routing IP packets not addressed to the virtual router

IP addresses. Many network administrators find this limitation frustrating when troubleshooting VRRP connectivity issues.

SR OS allows this access limitation to be selectively lifted for certain applications. Ping, Telnet and SSH can be individually enabled or disabled on a per-virtual-router-instance basis.

The **ping-reply** command enables the non-owner master to reply to ICMP echo requests directed at the virtual router instances IP addresses. The Ping request can be received on any routed interface. Ping must not have been disabled at the management security level (either on the parental IP interface or based on the Ping source host address).

When **ping-reply** is not enabled, ICMP echo requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP echo requests regardless of the **ping-reply** setting.

The **ping-reply** command is only available in non-owner **vrrp** nodal context.

By default, ICMP echo requests to the virtual router instance IP addresses are silently discarded.

The **no** form of the command configures discarding all ICMP echo request messages destined to the non-owner virtual router instance IP addresses.

Default

no ping-reply — ICMP echo requests to the virtual router instance IP addresses are discarded.

Platforms

7705 SAR Gen 2

22.4 ping-test

ping-test

Syntax

[**no**] **ping-test**

Context

[Tree] (config>filter>redirect-policy>dest ping-test)

Full Context

configure filter redirect-policy destination ping-test

Description

This command configures parameters to perform connectivity ping tests to validate the ability for the destination to receive redirected traffic.

Default
no ping-test

Platforms
7705 SAR Gen 2

22.5 pki

pki

Syntax
pki

Context
[Tree] (config>system>security pki)

Full Context
configure system security pki

Description
Commands in this context configure PKI related parameters.

Platforms
7705 SAR Gen 2

22.6 pkt-too-big

pkt-too-big

Syntax
[**no**] **pkt-too-big**

Context
[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>icmp6-gen pkt-too-big)
[Tree] (config>ipsec>tnl-temp>icmp6-gen pkt-too-big)
[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>icmp6-gen pkt-too-big)
[Tree] (config>service>vprn>if>sap>ipsec-tun>icmp6-gen pkt-too-big)

Full Context

configure service vprn interface ipsec ipsec-tunnel icmp6-generation pkt-too-big

configure ipsec tunnel-template icmp6-generation pkt-too-big

configure service ies interface ipsec ipsec-tunnel icmp6-generation pkt-too-big

configure service vprn interface sap ipsec-tunnel icmp6-generation pkt-too-big

Description

This command enables the system to send ICMPv6 PTB (Packet Too Big) messages on the private side and optionally specifies the rate.

With this command configured, the system sends PTB back if it received an IPv6 packet on the private side that is bigger than 1280 bytes and also exceeds the private MTU of the tunnel.

The **ip-mtu** command (under **ipsec-tunnel** or **tunnel-template**) specifies the private MTU for the ipsec-tunnel or dynamic tunnel.

The **no** form of this command reverts **interval** and **message-count** values to their default values.

Platforms

7705 SAR Gen 2

22.7 platform-type

platform-type

Syntax

platform-type *type*

no platform type

Context

[Tree] (config>system>ned>profile platform-type)

Full Context

configure system network-element-discovery profile platform-type

Description

This command configures the platform name and chassis type to be advertised.

The **no** form of this command removes any explicitly defined type and the default type of "chassis-name, chassis-type" is used.

Default

no platform-type

Parameters

type

Specifies the platform type to be associates with the profile, up to 255 characters.

Platforms

7705 SAR Gen 2

22.8 pmtu-discovery-aging

pmtu-discovery-aging

Syntax

pmtu-discovery-aging *seconds*

no pmtu-discovery-aging

Context

[Tree] (config>service>vprn>if>sap>ip-tunnel pmtu-discovery-aging)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel pmtu-discovery-aging)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel pmtu-discovery-aging)

[Tree] (config>router>if>ipsec>ipsec-tunnel pmtu-discovery-aging)

[Tree] (config>ipsec>tnl-temp pmtu-discovery-aging)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel pmtu-discovery-aging)

[Tree] (config>service>ies>if>sap>ip-tunnel pmtu-discovery-aging)

Full Context

configure service vprn interface sap ip-tunnel pmtu-discovery-aging

configure service ies interface ipsec ipsec-tunnel pmtu-discovery-aging

configure service vprn interface sap ipsec-tunnel pmtu-discovery-aging

configure router interface ipsec ipsec-tunnel pmtu-discovery-aging

configure ipsec tunnel-template pmtu-discovery-aging

configure service vprn interface ipsec ipsec-tunnel pmtu-discovery-aging

configure service ies interface sap ip-tunnel pmtu-discovery-aging

Description

This command configures the time used to age out the learned temporary MTU which is from the public network. The temporary MTU is used for MTU propagation.

The **no** form of of this command reverts to the default value.

Default

pmtu-discovery-aging 900

Parameters

seconds

specifies the time, in seconds, used to age out the learned MTU

Values 900 to 3600

Platforms

7705 SAR Gen 2

22.9 poi-tlv-enable

poi-tlv-enable

Syntax

[no] poi-tlv-enable

Context

[Tree] (config>service>vprn>isis poi-tlv-enable)

Full Context

configure service vprn isis poi-tlv-enable

Description

Enable use of Purge Originator Identification (POI) TLV for this IS-IS instance. The POI is added to purges and contains the system ID of the router that generated the purge, which simplifies troubleshooting and determining what caused the purge.

The **no** form of this command removes the POI functionality from the configuration.

Default

no poi-tlv-enable

Platforms

7705 SAR Gen 2

poi-tlv-enable

Syntax

[no] poi-tlv-enable

Context

[Tree] (config>router>isis poi-tlv-enable)

Full Context

configure router isis poi-tlv-enable

Description

Enable use of Purge Originator Identification (POI) TLV for this IS-IS instance. The POI is added to purges and contains the system ID of the router that generated the purge, which simplifies troubleshooting and determining what caused the purge.

The **no** form of this command removes the POI functionality from the configuration.

Default

no poi-tlv-enable

Platforms

7705 SAR Gen 2

22.10 policer

policer

Syntax

policer *policer-id* [**create**]

no policer *policer-id*

Context

[Tree] (config>card>fp>ingress>network>qgrp>policer-over policer)

[Tree] (config>card>fp>ingress>access>qgrp>policer-over policer)

Full Context

configure card fp ingress network queue-group policer-override policer

configure card fp ingress access queue-group policer-override policer

Description

This command creates, modifies or deletes a policer. Policers are created and used in a similar manner to queues. The policer ID space is separate from the queue ID space, allowing both a queue and a policer to share the same ID. The sap-ingress policy may have up to 32 policers (numbered 1 through 32) may be defined while the sap-egress QoS policy supports a maximum of 8 (numbered 1 through 8). While a policer may be defined within a QoS policy, it is not actually created on SAPs or subscribers associated with the policy until a forwarding class is mapped to the policer's ID.

All policers must be created within the QoS policies. A default policer is not created when a sap-ingress or sap-egress QoS policy is created.

Once a policer is created, the policer's metering rate and profiling rates may be defined as well as the policer's maximum and committed burst sizes (MBS and CBS respectively). Unlike queues which have dedicated counters, policers allow various stat-mode settings that define the counters that will be associated with the policer. Another supported feature—packet-byte-offset—provides a policer with the ability to modify the size of each packet based on a defined number of bytes.

Once a policer is created, it cannot be deleted from the QoS policy unless any forwarding classes that are mapped to the policer are first moved to other policers or queues.

The system will allow a policer to be created on a SAP QoS policy regardless of the ability to support policers on objects where the policy is currently applied. The system only scans the current objects for policer support and sufficient resources to create the policer when a forwarding class is first mapped to the policer ID. If the policer cannot be created due to one or more instances of the policy not supporting policing or having insufficient resources to create the policer, the forwarding class mapping fails.

The **no** form of this command deletes a policer from a sap-ingress or sap-egress QoS policy. The specified policer cannot currently have any forwarding class mappings for the removal of the policer to succeed. It is not necessary to actually delete the policer ID for the policer instances to be removed from SAPs or subscribers associated with the QoS policy once all forwarding classes have been moved away from the policer. It is automatically deleted from each policing instance although it still appears in the QoS policy.

Parameters

policer-id

Specifies that the *policer-id* must be specified when executing the policer command. If the specified ID already exists, the system enters that policer's context to allow the policer's parameters to be modified. If the ID does not exist and is within the allowed range for the QoS policy type, a context for the policer ID will be created (depending on the system's current create keyword requirements which may require the create keyword to actually add the new policer ID to the QoS policy) and the system will enter that new policer's context for possible parameter modification.

Values 1 to 32

Platforms

7705 SAR Gen 2

policer

Syntax

policer *policer-id* [**create**]

no policer *policer-id*

Context

[Tree] (config>service>epipe>sap>egress>policer-over policer)

Full Context

configure service epipe sap egress policer-override policer

Description

This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to a specific policer created on the SAP through a sap-ingress or sap-egress QoS policy.

The **no** form of this command is used to remove any existing overrides for the specified policer-id.

Parameters

policer-id

The *policer-id* parameter is required when executing the policer command within the policer-overrides context. The specified *policer-id* must exist within the sap-ingress or sap-egress QoS policy applied to the SAP. If the policer is not currently used by any forwarding class or forwarding type mappings, the policer will not actually exist on the SAP. This does not preclude creating an override context for the policer-id.

create

The create keyword is required when a **policer** policer-id override node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

Platforms

7705 SAR Gen 2

policer

Syntax

policer *policer-id* [**create**]

no policer *policer-id*

Context

[Tree] (config>service>vpls>sap>egress>policer-override policer)

[Tree] (config>service>vpls>sap>ingress>policer-override policer)

Full Context

configure service vpls sap egress policer-override policer

configure service vpls sap ingress policer-override policer

Description

This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to a specific policer created on the SAP through a sap-ingress or sap-egress QoS policy.

The **no** form of this command is used to remove any existing overrides for the specified policer-id.

Parameters

policer-id

The policer-id parameter is required when executing the policer command within the policer-overrides context. The specified policer-id must exist within the sap-ingress or sap-egress QoS policy applied to the SAP. If the policer is not currently used by any forwarding class or forwarding type mappings, the policer will not actually exist on the SAP. This does not preclude creating an override context for the policer-id.

create

The create keyword is required when a policer policer-id override node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

Platforms

7705 SAR Gen 2

policer

Syntax

policer *policer-id* [**create**]

no policer *policer-id*

Context

[Tree] (config>service>ies>if>sap>ingress>policer-override policer)

[Tree] (config>service>ies>if>sap>egress>policer-override policer)

Full Context

configure service ies interface sap ingress policer-override policer

configure service ies interface sap egress policer-override policer

Description

This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to a specific policer created on the SAP through a sap-ingress or sap-egress QoS policy.

The **no** form of this command is used to remove any existing overrides for the specified policer-id.

Parameters

policer-id

This parameter is required when executing the policer command within the policer-override context. The specified *policer-id* must exist within the sap-ingress or sap-egress QoS policy applied to the SAP. If the policer is not currently used by any forwarding class or forwarding type mappings, the policer will not actually exist on the SAP. This does not preclude creating an override context for the *policer-id*.

create

The create keyword is required when a policer override node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit configuration, the **create** keyword is not required.

Platforms

7705 SAR Gen 2

policer

Syntax

policer *policer-id* [**create**]

no policer *policer-id*

Context

[Tree] (config>service>vprn>if>sap>ingress>policer-override policer)

[Tree] (config>service>vprn>if>sap>egress>policer-override policer)

Full Context

configure service vprn interface sap ingress policer-override policer

configure service vprn interface sap egress policer-override policer

Description

This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to a specific policer created on the SAP through a sap-ingress or sap-egress QoS policy.

The **no** form of this command is used to remove any existing overrides for the specified policer-id.

Parameters

policer-id

This parameter is required when executing the policer command within the policer-override context. The specified *policer-id* must exist within the sap-ingress or sap-egress QoS policy applied to the SAP. If the policer is not currently used by any forwarding class or forwarding type mappings, the policer will not actually exist on the SAP. This does not preclude creating an override context for the *policer-id*.

create

The create keyword is required when a policer override node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit configuration, the **create** keyword is not required.

Platforms

7705 SAR Gen 2

policer

Syntax

policer *policer-id* [**fp-redirect-group**]

no policer

Context

[Tree] (config>qos>sap-ingress>fc policer)

Full Context

configure qos sap-ingress fc policer

Description

Within a sap-ingress QoS policy forwarding class context, the **policer** command is used to map packets that match the forwarding class and are considered unicast in nature to the specified policer-id. The specified policer-id must already exist within the **sap-ingress** QoS policy. While the system is determining the forwarding class of a packet, it is also looking up its forwarding destination. If ingress forwarding logic has resolved a unicast destination (the packet does not need to be sent to multiple destinations), it is considered to be a unicast packet and will be mapped to either an ingress queue (using the **queue** *queue-id* or **queue** *queue-id* **group** *ingress-queue-group* commands) or an ingress policer (**policer** *policer-id*). The **queue** and **policer** commands within the forwarding class context are mutually exclusive. By default, the unicast forwarding type is mapped to the SAP ingress default queue (queue 1). If the **policer** *policer-id* command is executed, any previous policer mapping or queue mapping for the unicast forwarding type within the forwarding class is overridden if the policer mapping is successful.

A policer defined within the **sap-ingress** policy is not actually created on an ingress SAP or a subscriber using an **sla-profile** where the policy is applied until at least one forwarding type (unicast, broadcast, unknown, or multicast) from one of the forwarding classes is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or multiservice site or ingress policing is not supported on the port associated with the SAP or subscriber or multiservice site, the initial forwarding class forwarding type mapping will fail.

When the unicast forwarding type within a forwarding class is mapped to a policer, the unicast packets classified to the subclasses within the forwarding class are also mapped to the policer.

The **no** form of this command is used to restore the mapping of the unicast forwarding type within the forwarding class to the default queue. If all forwarding class forwarding types had been removed from the default queue, the queue will not exist on the SAPs or subscriber or multiservice sites associated with the QoS policy and the **no policer** command will cause the system to attempt to create the default queue on each object. If the system cannot create the default queue in each instance, the **no policer** command will fail and the unicast forwarding type within the forwarding class will continue its mapping to the existing *policer-id*. If the **no policer** command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscriber will be lost.

Parameters

policer-id

When the forwarding class **policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the **sap-ingress** QoS policy.

Values 1 to 63

fp-redirect-group

Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

Platforms

7705 SAR Gen 2

policer

Syntax

policer *policer-id* [**create**]

no policer *policer-id*

Context

[Tree] (config>qos>sap-ingress policer)

Full Context

configure qos sap-ingress policer

Description

This command is used in the sap-ingress and sap-egress QoS policies to create, modify, or delete a policer. Policers are created and used in a similar manner to queues. The policer ID space is separate from the queue ID space, allowing both a queue and a policer to share the same ID. The sap-ingress policy may be defined to have up to 63 policers (numbered 1 through 63) while the sap-egress QoS policy supports a maximum of 8 (numbered 1 through 8). While a policer may be defined within a QoS policy, it is not actually created on SAPs or subscribers or multiservice sites associated with the policy until a forwarding class is mapped to the policer's ID.

All policers must be created within the QoS policies. A default policer is not created when a sap-ingress or sap-egress QoS policy is created.

When a policer is created, the policer's metering rate and profiling rates may be defined as well as the policer's maximum and committed burst sizes (MBS and CBS, respectively). Unlike queues that have dedicated counters, policers allow various stat-mode settings that define the counters that will be associated with the policer. Another supported feature—packet-byte-offset—provides a policer with the ability to modify the size of each packet, based on a defined number of bytes.

When a policer is created, it cannot be deleted from the QoS policy unless any forwarding classes that are mapped to the policer are first moved to other policers or queues.

The system will allow a policer to be created on a SAP QoS policy regardless of the ability to support policers on objects where the policy is currently applied. The system only scans the current objects for policer support and sufficient resources to create the policer when a forwarding class is first mapped to the policer ID. If the policer cannot be created due to one or more instances of the policy not supporting policing or having insufficient resources to create the policer, the forwarding class mapping will fail.

The **no** form of this command is used to delete a policer from a sap-ingress or sap-egress QoS policy. The specified policer cannot currently have any forwarding class mappings for the removal of the policer to succeed. It is not necessary to actually delete the policer ID for the policer instances to be removed from SAPs or subscriber or multiservice sites associated with the QoS policy when all forwarding classes have been moved away from the policer. It is automatically deleted from each policing instance although it still appears in the QoS policy.

Parameters

policer-id

The *policer-id* must be specified when executing the policer command. If the specified ID already exists, the system enters that policer's context to allow the policer's parameters to be modified. If the ID does not exist and is within the allowed range for the QoS policy type, a context for the policer ID will be created (depending on the system's current create keyword requirements, which may require the create keyword to actually add the new policer ID to the QoS policy) and the system will enter that new policer's context for possible parameter modification.

Values 1 to 63

Platforms

7705 SAR Gen 2

policer

Syntax

policer *policer-id* [{[**port-redirect-group-queue**] [**queue** *queue-id*] | **group** *group-name* [**instance** *instance-id*] [**queue** *queue-id*]}]

no policer

Context

[Tree] (config>qos>sap-egress>fc policer)

Full Context

configure qos sap-egress fc policer

Description

Within a sap-egress QoS policy forwarding class context, the policer command is used to map packets that match the forwarding class to the specified policer-id. The specified policer-id must already exist within the sap-egress QoS policy. The forwarding class of the packet is first discovered at ingress, based on the ingress classification rules. When the packet arrives at egress, the sap-egress QoS policy may match a forwarding class reclassification rule that overrides the ingress derived forwarding class. The forwarding

class context within the sap-egress QoS policy is then used to map the packet to an egress queue (using the queue queue-id, or port-redirect-group queue queue-id, or group queue-group-name instance instance-id queue queue-id commands) or an egress policer (policer policer-id). The queue and policer commands within the forwarding class context are mutually exclusive. By default, the forwarding class is mapped to the SAP egress default queue (queue 1). If the **policer** *policer- id* command is executed, any previous policer mapping or queue mapping for the forwarding class is overridden if the policer mapping is successful.

A policer defined within the sap-egress policy is not actually created on an egress SAP, or a subscriber using an SLA profile where the policy is applied, until at least one forwarding class is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber, or egress policing is not supported on the port associated with the SAP or subscriber, the initial forwarding class mapping will fail.

Packets that are mapped to an egress policer that are not discarded by the policer must be placed into a default queue on the packet's destination port. The system uses egress port queue groups for this purpose. An egress queue group named policer-output-queues is automatically created on each port that supports egress policers. By default, the system uses the forwarding class mappings within this queue group to decide which queue within the group will receive each packet output from the policer. This default policer output queuing behavior may be overridden for non-subscriber packets by redirection to a queue group. The name and instance of the queue group to redirect to is either specified in the QoS policy, or the fact that a forwarding class must be redirected is identified in the QoS policy and the specific queue group instance is only identified at the time the QoS policy is applied:

- If the **policer** *policer-id* command is successfully executed, the default egress queuing is performed for the forwarding class using the policer-output-queues queue group and the *queue-id* within the group based on the forwarding class map from the group template.
- If the **policer** *policer-id* **queue** *queue-id* command is successfully executed, the specified SAP *queue-id* within the egress QoS policy is used instead of the default policer output queues.
- If the **policer** *policer-id* **port-redirect-group-queue** keyword is successfully executed, the system will map the forwarding class to the queue within the egress queue group instance specified at the time the QoS policy is applied to the SAP, using the forwarding class map from the queue group template.
- If the **policer** *policer-id* **port-redirect-group queue** *queue-id* command is successfully executed, the system will map the forwarding class to the configured *queue-id* within the egress queue group instance that is specified at the time the QoS policy is applied to the SAP (ignoring using the forwarding class map from the queue group template).
- If the **policer** *policer-id* **group** *queue-group-name* **instance** *instance-id* command is successfully executed, the system will map the forwarding class to the queue within the specified egress queue group instance using the forwarding class map from the group template.
- If the **policer** *policer-id* **group** *queue-group-name* **instance** *instance-id* **queue** *queue-id* command is successfully executed, the system will map the forwarding class to the specified *queue-id* within the specified egress queue group instance (ignoring the forwarding class map in the group template).

If the specified **group** *group-name* is not defined as an egress queue-group-template, the **policer** command will fail. Also, if the specified group does not exist on the port for the SAPs or subscribers associated with the **sap-egress** QoS policy, the policer command will fail. While a group queue-group-name is specified in a **sap-egress** QoS policy, the groups corresponding egress template cannot be deleted. While a port egress queue group is associated with a policer instance, the port queue group cannot be deleted.

If the specified **queue** *queue-id* is not defined in the egress queue-group-template queue-group- name, the policer command will fail. While a *queue-id* within an egress queue group template is referenced by a **sap-**

egress QoS policy forwarding class policer command, the queue cannot be deleted from the queue group template.

If an egress policed packet is discarded by the egress port queue group queue, the source policer discard stats are incremented. This means that the discard counters for the policer represent both the policer discard events and the destination queue drop tail events associated with the policer.

The **no** form of this command is used to restore the mapping of the forwarding class to the default queue. If all forwarding classes have been removed from the default queue, the queue will not exist on the SAPs or subscribers associated with the QoS policy and the **no policer** command will cause the system to attempt to create the default queue on each object. If the system cannot create the default queue in each instance, the **no policer** command will fail and the forwarding class will continue its mapping to the existing *policer-id*. If the **no policer** command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscribers will be lost.

Default

no policer

Parameters

policer-id

When the forwarding class **policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the **sap-egress** QoS policy.

Values 1 to 63

port-redirect-group-queue

Used to override the forwarding class default egress queue destination to an egress port queue group. The specific egress queue group instance to use is specified at the time the QoS policy is applied to the SAP. Therefore, this parameter is only valid if SAP-based redirection is required.

queue *queue-id*

This parameter overrides the forwarding class default egress queue destination to a specified *queue-id*. If port-redirect-group is not configured, this will be a local SAP queue of that *queue-id*. A queue of ID *queue-id* must exist within the egress QoS policy. If **port-redirect-group-queue** is configured, the **queue** *queue-id* in the egress port queue group instance is used.

Values 1 to 8

Default Derived from forwarding class assignment in queue-group definition.

group *group-name*

The **group** *queue-group-name* is optional and is used to override the forwarding class's default egress queue destination. If the queue group-queue-id parameter is not specified, the forwarding class map within the specified group's template is used to derive which queue within the group will receive the forwarding class's packets. An egress queue group template must exist for the specified queue-group-name or the policer command will fail. The specified queue-group-name must also exist as an egress queue group on the ports

where SAPs and subscribers associated with the sap-egress policy are applied or the policer command will fail.

Values Any qualifying egress queue group name

Default policer-output-queues

queue *queue-id*

The **queue** *group-queue-id* is optional when the group queue-group-name parameter is specified and is used to override the forwarding class mapping within the group's egress queue group template. The specified group-queue-id must exist within the group's egress queue group template or the policer command will fail.

Values 1 to 8

Default Derived from forwarding class assignment in queue-group definition

instance *instance-id*

This parameter is used to specify the specific instance of a queue group with template queue-group-name to which this queue should be redirected. This parameter is only valid for queue groups on egress ports where policy-based redirection is required.

Values 1 to 40960

Default 1

Platforms

7705 SAR Gen 2

policer

Syntax

policer *policer-id* [**create**]

no policer

Context

[Tree] (config>qos>sap-egress policer)

Full Context

configure qos sap-egress policer

Description

A policer defined within the sap-egress policy is not actually created on an egress SAP, or a subscriber using an SLA profile where the policy is applied, until at least one forwarding class is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber, or egress policing is not supported on the port associated with the SAP or subscriber, the initial forwarding class mapping will fail.

Packets that are mapped to an egress policer that are not discarded by the policer must be placed into a default queue on the packet's destination port. The system uses egress port queue groups for this purpose. An egress queue group named policer-output-queues is automatically created on each port that supports egress policers. By default, the system uses the forwarding class mappings within this queue group to decide which queue within the group will receive each packet output from the policer. This default policer output queuing behavior may be overridden for non-subscriber packets by redirection to a queue group. The name and instance of the queue group to redirect to is either specified in the QoS policy, or the fact that a forwarding class must be redirected is identified in the QoS policy and the specific queue group instance is only identified at the time the QoS policy is applied:

- If the **policer** *policer-id* command is successfully executed, the default egress queuing is performed for the forwarding class using the policer-output-queues queue group and the *queue-id* within the group based on the forwarding class map from the group template.
- If the **policer** *policer-id* **queue** *queue-id* command is successfully executed, the specified SAP *queue-id* within the egress QoS policy is used instead of the default policer output queues.
- If the **policer** *policer-id* **port-redirect-group-queue** keyword is successfully executed, the system will map the forwarding class to the queue within the egress queue group instance specified at the time the QoS policy is applied to the SAP, using the forwarding class map from the queue group template.
- If the **policer** *policer-id* **port-redirect-group queue** *queue-id* command is successfully executed, the system will map the forwarding class to the configured *queue-id* within the egress queue group instance that is specified at the time the QoS policy is applied to the SAP (ignoring using the forwarding class map from the queue group template).
- If the **policer** *policer-id* **group** *queue-group-name* **instance** *instance-id* command is successfully executed, the system will map the forwarding class to the queue within the specified egress queue group instance using the forwarding class map from the group template.
- If the **policer** *policer-id* **group** *queue-group-name* **instance** *instance-id* **queue** *queue-id* command is successfully executed, the system will map the forwarding class to the specified *queue-id* within the specified egress queue group instance (ignoring the forwarding class map in the group template).

If the specified **group** *group-name* is not defined as an egress queue-group-template, the **policer** command will fail. Also, if the specified group does not exist on the port for the SAPs or subscribers associated with the **sap-egress** QoS policy, the policer command will fail. While a group queue-group-name is specified in a **sap-egress** QoS policy, the groups corresponding egress template cannot be deleted. While a port egress queue group is associated with a policer instance, the port queue group cannot be deleted.

If the specified **queue** *queue-id* is not defined in the egress queue-group-template queue-group- name, the policer command will fail. While a *queue-id* within an egress queue group template is referenced by a **sap-egress** QoS policy forwarding class policer command, the queue cannot be deleted from the queue group template.

If an egress policed packet is discarded by the egress port queue group queue, the source policer discard stats are incremented. This means that the discard counters for the policer represent both the policer discard events and the destination queue drop tail events associated with the policer.

The **no** form of this command is used to restore the mapping of the forwarding class to the default queue. If all forwarding classes have been removed from the default queue, the queue will not exist on the SAPs or subscribers associated with the QoS policy and the **no policer** command will cause the system to attempt to create the default queue on each object. If the system cannot create the default queue in each instance, the **no policer** command will fail and the forwarding class will continue its mapping to the existing *policer-id*. If the **no policer** command results in a policer without any current mappings, the policer will be removed

from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscribers will be lost.

Default

no policer

Parameters

policer-id

When the forwarding class **policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the **sap-egress** QoS policy.

Values 1 to 63

Platforms

7705 SAR Gen 2

policer

Syntax

policer *policer-id* [**create**]

no policer *policer-id*

Context

[Tree] (config>qos>qgrps>ing>queue-group policer)

[Tree] (cfg>qos>qgrps>egr>queue-group policer)

Full Context

configure qos queue-group-templates ingress queue-group policer

configure qos queue-group-templates egress queue-group policer

Description

This command is used in ingress and egress queue-group templates to create, modify, or delete a policer.

Policers are created and used in a similar manner to queues. The policer ID space is separate from the queue ID space, allowing both a queue and a policer to share the same ID. While a policer may be defined in a queue-group template, it is not actually created until the queue-group template is instantiated on the ingress context of a forwarding plane or on the egress context of a port.

When a policer is created, the policer's metering rate and profiling rates may be defined, as well as the policer's maximum and committed burst sizes (MBS and CBS, respectively). Unlike queues that have dedicated counters, policers allow various stat-mode settings that define the counters that will be associated with the policer. Another supported feature—packet-byte-offset—provides a policer with the ability to modify the size of each packet based on a defined number of bytes.

When a policer is created, it cannot be deleted from the queue-group template unless any forwarding classes that are redirected to the policer are first removed.

The **no** version of this command deletes the policer.

Parameters

policer-id

The *policer-id* must be specified when executing the policer command. If the specified ID already exists, the system enters that policer's context to allow the policer's parameters to be modified. If the ID does not exist and is within the allowed range for the QoS policy type, a context for the policer ID will be created (depending on the system's current create keyword requirements, which may require the **create** keyword to actually add the new policer ID to the QoS policy) and the system enters that new policer's context for possible parameter modification.

Values			
	ingress	all platforms	1 to 32
	egress	all other platforms	1 to 16

Platforms

7705 SAR Gen 2

policer

Syntax

[**no**] **policer** *policer-id*

Context

[Tree] (config>log>acct-policy>cr policer)

Full Context

configure log accounting-policy custom-record policer

Description

This command creates a policer context for which counters should be included in the custom-record.

The **no** form of this command deletes the policer and its counters from the custom-record.

Parameters

policer-id

Specifies the policer for which counters should be included in or deleted from the custom-record.

Values 1 to 63

Platforms

7705 SAR Gen 2

22.11 policer-control-override

policer-control-override

Syntax

policer-control-override [**create**]

no policer-control-override

Context

[Tree] (config>card>fp>ingress>network>queue-group policer-control-override)

[Tree] (config>card>fp>ingress>access>queue-group policer-control-override)

Full Context

configure card fp ingress network queue-group policer-control-override

configure card fp ingress access queue-group policer-control-override

Description

This command configures policer control overrides.

Parameters

create

Keyword required to create a new policer control override instance.

Platforms

7705 SAR Gen 2

policer-control-override

Syntax

policer-control-override [**create**]

no policer-control-override

Context

[Tree] (config>service>epipe>sap>ingress policer-control-override)

[Tree] (config>service>epipe>sap>egress policer-control-override)

Full Context

configure service epipe sap ingress policer-control-override

configure service epipe sap egress policer-control-override

Description

This command, within the SAP ingress or egress contexts, creates a CLI node for specific overrides to the applied policer-control-policy. A policy must be applied for a policer-control-overrides node to be created. If the policer-control-policy is removed or changed, the policer-control-overrides node is automatically deleted from the SAP.

The **no** form of this command removes any existing policer-control-policy overrides and the policer-control-overrides node from the SAP.

Default

no policer-control-override

Parameters

create

The create keyword is required when the policer-control-overrides node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

Platforms

7705 SAR Gen 2

policer-control-override

Syntax

policer-control-override [**create**]

no policer-control-override

Context

[Tree] (config>service>vpls>sap>egress policer-control-override)

[Tree] (config>service>vpls>sap>ingress policer-control-override)

Full Context

configure service vpls sap egress policer-control-override

configure service vpls sap ingress policer-control-override

Description

This command, within the SAP ingress or egress contexts, creates a CLI node for specific overrides to the applied policer-control-policy. A policy must be applied for a policer-control-overrides node to be created.

If the policer-control-policy is removed or changed, the policer-control-overrides node is automatically deleted from the SAP.

The **no** form of this command removes any existing policer-control-policy overrides and the policer-control-overrides node from the SAP.

Default

no policer-control-override

Parameters

create

The create keyword is required when the policer-control-overrides node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

Platforms

7705 SAR Gen 2

policer-control-override

Syntax

policer-control-override [**create**]

no policer-control-override

Context

[Tree] (config>service>ies>if>sap>ingress policer-control-override)

[Tree] (config>service>ies>if>sap>egress policer-control-override)

Full Context

configure service ies interface sap ingress policer-control-override

configure service ies interface sap egress policer-control-override

Description

This command, within the SAP ingress or egress contexts, creates a CLI node for specific overrides to the applied policer-control-policy. A policy must be applied for a policer-control-overrides node to be created. If the policer-control-policy is removed or changed, the policer-control-overrides node is automatically deleted from the SAP.

The **no** form of this command removes any existing policer-control-policy overrides and the policer-control-overrides node from the SAP.

Default

no policer-control-override

Parameters

create

The create keyword is required when the policer-control-overrides node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

Platforms

7705 SAR Gen 2

policer-control-override

Syntax

policer-control-override [**create**]

no policer-control-override

Context

[Tree] (config>service>vprn>if>sap>egress policer-control-override)

[Tree] (config>service>vprn>if>sap>ingress policer-control-override)

Full Context

configure service vprn interface sap egress policer-control-override

configure service vprn interface sap ingress policer-control-override

Description

This command, within the SAP ingress or egress contexts, creates a CLI node for specific overrides to the applied policer-control-policy. A policy must be applied for a policer-control-overrides node to be created. If the policer-control-policy is removed or changed, the policer-control-overrides node is automatically deleted from the SAP.

The **no** form of this command removes any existing policer-control-policy overrides and the policer-control-overrides node from the SAP.

Default

no policer-control-override

Parameters

create

The create keyword is required when the policer-control-overrides node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

Platforms

7705 SAR Gen 2

22.12 policer-control-policy

policer-control-policy

Syntax

policer-control-policy *policer-control-policy-name*

no policer-control-policy

Context

[Tree] (config>card>fp>ingress>access>queue-group policer-control-policy)

[Tree] (config>card>fp>ingress>network>queue-group policer-control-policy)

Full Context

configure card fp ingress access queue-group policer-control-policy

configure card fp ingress network queue-group policer-control-policy

Description

This command configures an policer-control policy that can apply to a queue-group on the forwarding plane.

The **no** form of this command removes the policer-control policy association from the queue-group.

Default

no policer-control-policy

Parameters

policer-control-policy-name

Specifies the name of the policer-control policy to use for the queue-group. The name can be up to 32 characters long.

Platforms

7705 SAR Gen 2

policer-control-policy

Syntax

policer-control-policy *policy-name*

no policer-control-policy

Context

[Tree] (config>port>ethernet>network>egress>queue-group policer-control-policy)

Full Context

configure port ethernet network egress queue-group policer-control-policy

Description

This command configures the policer control policy for the QoS egress queue-group.

Parameters

policy-name

Specifies the name of the policer control policy, up to 32 characters.

Platforms

7705 SAR Gen 2

policer-control-policy

Syntax

policer-control-policy *policy-name*

no policer-control-policy

Context

[Tree] (config>service>epipe>sap>egress policer-control-policy)

[Tree] (config>service>epipe>sap>ingress policer-control-policy)

Full Context

configure service epipe sap egress policer-control-policy

configure service epipe sap ingress policer-control-policy

Description

This command, within the QoS CLI node, is used to create, delete or modify policer control policies. A policer control policy is very similar to the scheduler-policy which is used to manage a set of queues by defining a hierarchy of virtual schedulers and specifying how the virtual schedulers interact to provide an aggregate SLA. In a similar fashion, the policer-control-policy controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs.

Policer Control Policy Instances

On the SAP side, an instance of a policy is created each time a policy is applied.

When applied to a sub-profile, an instance of the policy is created each time a subscriber successfully maps one or more hosts to the profile per ingress SAP.

Each instance of the policer-control-policy manages the policers associated with the object that owns the policy instance (SAP or subscriber). If a policer on the object is parented to an appropriate arbiter name that exists within the policy, the policer will be managed by the instance. If a policer is not parented or is parented to a non-existent arbiter, the policer will be orphaned and will not be subject to bandwidth control by the policy instance.

Maximum Rate and Root Arbiter

The policer-control-policy supports an overall maximum rate (max-rate) that defines the total amount of bandwidth that may be distributed to all associated child policers. By default, that rate is set to max which provides an unlimited amount of bandwidth to the policers. Once the policy is created, an actual rate should be configured in order for the policy instances to be effective. At the SAP level, the maximum rate may be overridden on a per instance basis.

For subscribers, the maximum rate may only be overridden on the subscriber profile which will then be applied to all instances associated with the profile.

The maximum rate is defined within the context of the root arbiter which is always present in a policer-control-policy. The system creates a parent policer which polices the output of all child policers attached to the policy instance to the configured rate. Child policers may be parented directly to the root arbiter (parent root) or parented to one of the tiered arbiters (parent arbiter-name). Since each tiered arbiter must be parented to either another tiered arbiter or the root arbiter (default), every parented child policer is associated with the root arbiter and therefore the root arbiter's parent policer.

Parent Policer PIR Leaky Bucket Operation

The parent policer is a single leaky bucket that monitors the aggregate throughput rate of the associated child policers. Forwarded packets increment the bucket by the size of each packet. The rate of the parent policer is implemented as a bucket decrement function which attempts to drain the bucket. If the rate of the packets flowing through the bucket is less than the decrement rate, the bucket does not accumulate depth. Each packet that flows through the bucket is accompanied by a derived discard threshold. If the current depth of the bucket is less than the discard threshold, the packet is allowed to pass through, retaining the colors derived from the packet's child policer. If the current depth is equal to or greater than the threshold value, the packet is colored red and the bucket depth is not incremented by the packet size. Also, any increased bucket depths in the child policer are canceled making any discard event an atomic function between the child and the parent.

Due to the fact that multiple thresholds are supported by the parent policer, the policer control policy is able to protect the throughput of higher priority child policers from the throughput of the lower priority child policers within the aggregate rate.

Tier 1 and Tier 2 Arbiters

As previously stated, each child is attached either to the always available root arbiter or to an explicitly created tier 1 or tier 2 arbiter. Unlike the hardware parent policer based root arbiter, the arbiters at tier 1 and tier 2 are only represented in software and are meant to provide an arbitrary hierarchical bandwidth distribution capability. An arbiter created on tier 2 must parent to either to an arbiter on tier 1 or to the root arbiter. Arbiters created on tier 1 always parent to the root arbiter. In this manner, every arbiter ultimately is parented or grand-parented by the root arbiter.

Each tiered arbiter supports an optional rate parameter that defines a rate limit for all child arbiters or child policers associated with the arbiter. Child arbiters and policers attached to the arbiter have a level attribute that defines the strict level at which the child is given bandwidth by the arbiter. Level 8 is the highest and 1 is the lowest. Also a weight attribute defines each child's weight at that strict level in order to determine how bandwidth is distributed to multiple children at that level when insufficient bandwidth is available to meet each child's required bandwidth.

Fair and Unfair Bandwidth Control

Each child policer supports three leaky buckets. The PIR bucket manages the policer's peak rate and maximum burst size, the CIR leaky bucket manages the policer's committed rate and committed burst size. The third leaky bucket is used by the policer control policy instance to manage the child policer's fair rate (FIR). When multiple child policers are attached to the root arbiter at the same priority level, the policy instance uses each child's FIR bucket rate to control how much of the traffic forwarded by the policer is fair and how much is unfair.

In the simplest case where all the child policers in the same priority level are directly attached to the root arbiter, each child's FIR rate is set according to the child's weight divided by the sum of the active children's weights multiplied by the available bandwidth at the priority level. The result is that the FIR bucket will mark the appropriate amount of traffic for each child as fair-based on the weighted fair output of the policy instance.

The fair/unfair forwarding control in the root parent policer is accomplished by implementing two different discard thresholds for the priority. The first threshold is discard-unfair and the second is discard-all for packet associated with the priority level. As the parent policer PIR bucket fills (due the aggregate forwarded rate being greater than the parent policers PIR decrement rate) and the bucket depth reaches the first threshold, all unfair packets within the priority are discarded. This leaves room in the bucket for the fair packets to be forwarded.

In the more complex case where one or more tiered arbiters are attached at the priority level, the policer control policy instance must consider more than just the child policer weights associated with the attached arbiter. If the arbiter is configured with an aggregate rate limit that its children cannot exceed, the policer control policy instance will switch to calculating the rate each child serviced by the arbiter should receive and enforces that rate using each child policers PIR leaky bucket.

When the child policer PIR leaky bucket is used to limit the bandwidth for the child policer and the child's PIR bucket discard threshold is reached, packets associated with the child policer are discarded. The child policer's discarded packets do not consume depth in the child policer's CIR or FIR buckets. The child policers discarded packets are also prevented from impacting the parent policer and will not consume the aggregate bandwidth managed by the parent policer.

Parent Policer Priority Level Thresholds

As stated in the Tier 1 and Tier 2 Arbiter subsection, each child policer is attached either to the root arbiter or explicitly to one of the tier 1 or tier 2 arbiters. When attached directly to the root arbiter, its priority relative to all other child policers is indicated by the parenting level parameter. When attached through one of the tiered arbiters, the parenting hierarchy of the arbiters must be traced through to the ultimate attachment to the root arbiter. The parenting level parameter of the arbiter parented to the root arbiter defines the child policer's priority level within the parent policer.

The priority level is important since it defines the parent policer discard thresholds that will be applied at the parent policer. The parent policer has 8 levels of strict priority and each priority level has its own discard-unfair and discard-all thresholds. Each priority's thresholds are larger than the thresholds of the lower priority levels. This ensures that when the parent policer is discarding, it will be priority sensitive.

To visualize the behavior of the parent policer, picture that when the aggregate forwarding rate of all child policers is currently above the decrement rate of the parent PIR leaky bucket, the bucket depth will increase over time. As the bucket depth increases, it will eventually cross the lowest priority's discard-unfair threshold. If this amount of discard sufficiently lowers the remaining aggregate child policer rate, the parent PIR bucket will hover around this bucket depth. If however, the remaining aggregate child rate is still greater than the decrement rate, the bucket will continue to rise and eventually reach the lowest priority's discard-all threshold which will cause all packets associated with the priority level to be discarded (fair and unfair). Again, if the remaining aggregate child rate is less than or equal to the bucket decrement rate, the parent PIR bucket will hover around this higher bucket depth. If the remaining aggregate child rate is still

higher than the decrement rate, the bucket will continue to rise through the remaining priority level discards until equilibrium is achieved.

Each child's rate feeding into the parent policer is governed by the child policer's PIR bucket decrement rate. The amount of bandwidth the child policer offers to the parent policer will not exceed the child policer's configured maximum rate.

Root Arbiter's Parent Policer's Priority Aggregate Thresholds

Each policer-control-policy root arbiter supports configurable aggregate priority thresholds which are used to control burst tolerance within each priority level. Two values are maintained per priority level; the shared-portion and the fair-portion. The shared-portion represents the amount of parent PIR bucket depth that is allowed to be consumed by both fair and unfair child packets at the priority level. The fair-portion represents the amount of parent PIR bucket depth that only the fair child policer packets may consume within the priority level. It should be noted that the fair and unfair child packets associated with a higher parent policer priority level may also consume the bucket depth set aside for this priority.

While the policy maintains a parent policer default or explicit configurable values for shared-portion and fair-portion within each priority level, it is possible that some priority levels will not be used within the parent policer. Most parent policer use cases require fewer than eight strict priority levels.

To derive the actual priority level discard-unfair and discard-all thresholds while only accounting for the actual in-use priority levels, the system maintains a child policer to parent policer association counter per priority level for each policer control policy instance. As a child policer is parented to either the root or a tiered arbiter, the system determines the parent policer priority level for the child policer and increments the association counter for that priority level on the parent policer instance.

The shared-portion for each priority level is affected by the parent policer global min-thresh-separation parameter that defines the minimum separation between any in-use discard thresholds. When more than one child policer is associated with a parent policer priority level, the shared-portion for that priority level will be the current value of min-thresh-separation. When only a single child policer is associated, the priority level's shared-portion is zero since all packets from the child will be marked fair and the discard-unfair threshold is meaningless. When the association counter is zero, both the shared-portion and the fair-portion for that priority level are zero since neither discard thresholds will be used. Whenever the association counter is greater than 0, the fair-portion for that priority level will be derived from the current value of the priority's mbs-contribution parameter and the global min-thresh-separation parameter.

Each priority level's discard-unfair and discard-all thresholds are calculated based on an accumulation of lower priorities shared-portions and fair-portions and the priority level's own shared-portion and fair-portion. The base threshold value for each priority level is equal to the sum of all lower priority level's shared-portions and fair-portions. The discard-unfair threshold is the priority level's base threshold plus the priority level's shared-portion. The discard-all threshold for the priority level is the priority level's base threshold plus both the shared-portion and fair-portion values of the priority. As can be seen, an in-use priority level's thresholds are always greater than the thresholds of lower priority levels.

Policer Control Policy Application

A policer-control-policy may be applied on any Ethernet ingress or egress SAP that is associated with a port (or ports in the case of LAG).

The **no** form of this command removes a non-associated policer control policy from the system. The command will not execute when policer-name is currently associated with any SAP context.

Parameters

policy-name

Each policer-control-policy must be created with a unique policy name. The name given as policy-name must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system will enter that policy's context for editing purposes. If policy-name does not exist, the system will attempt to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode.

create

The keyword is required when a new policy is being created and the system is configured for explicit object creation mode.

Platforms

7705 SAR Gen 2

policer-control-policy

Syntax

policer-control-policy *policy-name*

no policer-control-policy

Context

[Tree] (config>service>template>vpls-sap-template>egress policer-control-policy)

[Tree] (config>service>template>vpls-sap-template>ingress policer-control-policy)

[Tree] (config>service>vpls>sap>ingress policer-control-policy)

[Tree] (config>service>vpls>sap>egress policer-control-policy)

Full Context

configure service template vpls-sap-template egress policer-control-policy

configure service template vpls-sap-template ingress policer-control-policy

configure service vpls sap ingress policer-control-policy

configure service vpls sap egress policer-control-policy

Description

This command, within the qos CLI node, is used to create, delete or modify policer control policies. A policer control policy is very similar to the scheduler-policy which is used to manage a set of queues by defining a hierarchy of virtual schedulers and specifying how the virtual schedulers interact to provide an aggregate SLA. In a similar fashion, the policer-control-policy controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs. The policy may also be applied to the ingress or egress context of a sub-profile.

Policer Control Policy Instances

On the SAP side, an instance of a policy is created each time a policy is applied. When applied to a sub-profile, an instance of the policy is created each time a subscriber successfully maps one or more hosts to the profile per ingress SAP.

Each instance of the policer-control-policy manages the policers associated with the object that owns the policy instance (SAP or subscriber). If a policer on the object is parented to an appropriate arbiter name that exists within the policy, the policer will be managed by the instance. If a policer is not parented or is parented to a non-existent arbiter, the policer will be orphaned and will not be subject to bandwidth control by the policy instance.

Maximum Rate and Root Arbiter

The policer-control-policy supports an overall maximum rate (max-rate) that defines the total amount of bandwidth that may be distributed to all associated child policers. By default, that rate is set to max which provides an unlimited amount of bandwidth to the policers. Once the policy is created, an actual rate should be configured in order for the policy instances to be effective. At the SAP level, the maximum rate may be overridden on a per instance basis. For subscribers, the maximum rate may only be overridden on the subscriber profile which will then be applied to all instances associated with the profile.

The maximum rate is defined within the context of the root arbiter which is always present in a policer-control-policy. The system creates a parent policer which polices the output of all child policers attached to the policy instance to the configured rate. Child policers may be parented directly to the root arbiter (parent root) or parented to one of the tiered arbiters (parent arbiter-name). Since each tiered arbiter must be parented to either another tiered arbiter or the root arbiter (default), every parented child policer is associated with the root arbiter and therefore the root arbiter's parent policer.

Parent Policer PIR Leaky Bucket Operation

The parent policer is a single leaky bucket that monitors the aggregate throughput rate of the associated child policers. Forwarded packets increment the bucket by the size of each packet. The rate of the parent policer is implemented as a bucket decrement function which attempts to drain the bucket. If the rate of the packets flowing through the bucket is less than the decrement rate, the bucket does not accumulate depth. Each packet that flows through the bucket is accompanied by a derived discard threshold. If the current depth of the bucket is less than the discard threshold, the packet is allowed to pass through, retaining the colors derived from the packet's child policer. If the current depth is equal to or greater than the threshold value, the packet is colored red and the bucket depth is not incremented by the packet size. Also, any increased bucket depths in the child policer are canceled making any discard event an atomic function between the child and the parent.

Due to the fact that multiple thresholds are supported by the parent policer, the policer control policy is able to protect the throughput of higher priority child policers from the throughput of the lower priority child policers within the aggregate rate.

Tier 1 and Tier 2 Arbiters

As stated above, each child is attached either to the always available root arbiter or to an explicitly created tier 1 or tier 2 arbiter. Unlike the hardware parent policer based root arbiter, the arbiters at tier 1 and tier 2 are only represented in software and are meant to provide an arbitrary hierarchical bandwidth distribution capability. An arbiter created on tier 2 must parent to either to an arbiter on tier 1 or to the root arbiter. Arbiters created on tier 1 always parent to the root arbiter. In this manner, every arbiter ultimately is parented or grand-parented by the root arbiter.

Each tiered arbiter supports an optional rate parameter that defines a rate limit for all child arbiters or child policers associated with the arbiter. Child arbiters and policers attached to the arbiter have a level attribute that defines the strict level at which the child is given bandwidth by the arbiter. Level 8 is the highest and 1 is the lowest. Also a weight attribute defines each child's weight at that strict level in order to determine

how bandwidth is distributed to multiple children at that level when insufficient bandwidth is available to meet each child's required bandwidth.

Fair and Unfair Bandwidth Control

Each child policer supports three leaky buckets. The PIR bucket manages the policer's peak rate and maximum burst size, the CIR leaky bucket manages the policer's committed rate and committed burst size. The third leaky bucket is used by the policer control policy instance to manage the child policer's fair rate (FIR). When multiple child policers are attached to the root arbiter at the same priority level, the policy instance uses each child's FIR bucket rate to control how much of the traffic forwarded by the policer is fair and how much is unfair.

In the simplest case where all the child policers in the same priority level are directly attached to the root arbiter, each child's FIR rate is set according to the child's weight divided by the sum of the active children's weights multiplied by the available bandwidth at the priority level. The result is that the FIR bucket will mark the appropriate amount of traffic for each child as fair-based on the weighted fair output of the policy instance.

The fair/unfair forwarding control in the root parent policer is accomplished by implementing two different discard thresholds for the priority. The first threshold is discard-unfair and the second is discard-all for packet associated with the priority level. As the parent policer PIR bucket fills (due the aggregate forwarded rate being greater than the parent policers PIR decrement rate) and the bucket depth reaches the first threshold, all unfair packets within the priority are discarded. This leaves room in the bucket for the fair packets to be forwarded.

In the more complex case where one or more tiered arbiters are attached at the priority level, the policer control policy instance must consider more than just the child policer weights associated with the attached arbiter. If the arbiter is configured with an aggregate rate limit that its children cannot exceed, the policer control policy instance will switch to calculating the rate each child serviced by the arbiter should receive and enforces that rate using each child policers PIR leaky bucket.

When the child policer PIR leaky bucket is used to limit the bandwidth for the child policer and the child's PIR bucket discard threshold is reached, packets associated with the child policer are discarded. The child policer's discarded packets do not consume depth in the child policer's CIR or FIR buckets. The child policers discarded packets are also prevented from impacting the parent policer and will not consume the aggregate bandwidth managed by the parent policer.

Parent Policer Priority Level Thresholds

As stated above, each child policer is attached either to the root arbiter or explicitly to one of the tier 1 or tier 2 arbiters. When attached directly to the root arbiter, its priority relative to all other child policers is indicated by the parenting level parameter. When attached through one of the tiered arbiters, the parenting hierarchy of the arbiters must be traced through to the ultimate attachment to the root arbiter. The parenting level parameter of the arbiter parented to the root arbiter defines the child policer's priority level within the parent policer.

The priority level is important since it defines the parent policer discard thresholds that will be applied at the parent policer. The parent policer has 8 levels of strict priority and each priority level has its own discard-unfair and discard-all thresholds. Each priority's thresholds are larger than the thresholds of the lower priority levels. This ensures that when the parent policer is discarding, it will be priority sensitive.

To visualize the behavior of the parent policer, picture that when the aggregate forwarding rate of all child policers is currently above the decrement rate of the parent PIR leaky bucket, the bucket depth will increase over time. As the bucket depth increases, it will eventually cross the lowest priority's discard-unfair threshold. If this amount of discard sufficiently lowers the remaining aggregate child policer rate, the parent PIR bucket will hover around this bucket depth. If however, the remaining aggregate child rate is still greater than the decrement rate, the bucket will continue to rise and eventually reach the lowest priority's

discard-all threshold which will cause all packets associated with the priority level to be discarded (fair and unfair). Again, if the remaining aggregate child rate is less than or equal to the bucket decrement rate, the parent PIR bucket will hover around this higher bucket depth. If the remaining aggregate child rate is still higher than the decrement rate, the bucket will continue to rise through the remaining priority level discards until equilibrium is achieved.

As noted above, each child's rate feeding into the parent policer is governed by the child policer's PIR bucket decrement rate. The amount of bandwidth the child policer offers to the parent policer will not exceed the child policer's configured maximum rate.

Root Arbiter's Parent Policer's Priority Aggregate Thresholds

Each policer-control-policy root arbiter supports configurable aggregate priority thresholds which are used to control burst tolerance within each priority level. Two values are maintained per priority level; the shared-portion and the fair-portion. The shared-portion represents the amount of parent PIR bucket depth that is allowed to be consumed by both fair and unfair child packets at the priority level. The fair-portion represents the amount of parent PIR bucket depth that only the fair child policer packets may consume within the priority level. It should be noted that the fair and unfair child packets associated with a higher parent policer priority level may also consume the bucket depth set aside for this priority.

While the policy maintains a parent policer default or explicit configurable values for shared-portion and fair-portion within each priority level, it is possible that some priority levels will not be used within the parent policer. Most parent policer use cases require fewer than eight strict priority levels.

To derive the actual priority level discard-unfair and discard-all thresholds while only accounting for the actual in-use priority levels, the system maintains a child policer to parent policer association counter per priority level for each policer control policy instance. As a child policer is parented to either the root or a tiered arbiter, the system determines the parent policer priority level for the child policer and increments the association counter for that priority level on the parent policer instance.

The shared-portion for each priority level is affected by the parent policer global min-thresh-separation parameter that defines the minimum separation between any in-use discard thresholds. When more than one child policer is associated with a parent policer priority level, the shared-portion for that priority level will be the current value of min-thresh-separation. When only a single child policer is associated, the priority level's shared-portion is zero since all packets from the child will be marked fair and the discard-unfair threshold is meaningless. When the association counter is zero, both the shared-portion and the fair-portion for that priority level are zero since neither discard thresholds will be used. Whenever the association counter is greater than 0, the fair-portion for that priority level will be derived from the current value of the priority's mbs-contribution parameter and the global min-thresh-separation parameter.

Each priority level's discard-unfair and discard-all thresholds are calculated based on an accumulation of lower priorities shared-portions and fair-portions and the priority level's own shared-portion and fair-portion. The base threshold value for each priority level is equal to the sum of all lower priority level's shared-portions and fair-portions. The discard-unfair threshold is the priority level's base threshold plus the priority level's shared-portion. The discard-all threshold for the priority level is the priority level's base threshold plus both the shared-portion and fair-portion values of the priority. As can be seen, an in-use priority level's thresholds are always greater than the thresholds of lower priority levels.

Policer Control Policy Application

A policer-control-policy may be applied on any Ethernet ingress or egress SAP that is associated with a port (or ports in the case of LAG).

The **no** form of this command removes a non-associated policer control policy from the system. The command will not execute when policer-name is currently associated with any SAP or subscriber management sub-profile context.

Parameters

policy-name

Specifies the policy name. Each policer-control-policy must be created with a unique policy name. The name must be given as policy-name must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system will enter that policy's context for editing purposes. If policy-name does not exist, the system will attempt to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode.

create

The keyword is required when a new policy is being created and the system is configured for explicit object creation mode.

Platforms

7705 SAR Gen 2

policer-control-policy

Syntax

policer-control-policy *policy-name*

no policer-control-policy

Context

[Tree] (config>service>ies>if>sap>egress policer-control-policy)

[Tree] (config>service>ies>if>sap>ingress policer-control-policy)

Full Context

configure service ies interface sap egress policer-control-policy

configure service ies interface sap ingress policer-control-policy

Description

This command, within the qos CLI node, is used to create, delete or modify policer control policies. A policer control policy is very similar to the scheduler-policy which is used to manage a set of queues by defining a hierarchy of virtual schedulers and specifying how the virtual schedulers interact to provide an aggregate SLA. In a similar fashion, the policer-control-policy controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs. The policy may also be applied to the ingress or egress context of a sub-profile.

Policer Control Policy Instances

On the SAP side, an instance of a policy is created each time a policy is applied. When applied to a sub-profile, an instance of the policy is created each time a subscriber successfully maps one or more hosts to the profile per ingress SAP.

Each instance of the policer-control-policy manages the policers associated with the object that owns the policy instance (SAP or subscriber). If a policer on the object is parented to an appropriate arbiter name that exists within the policy, the policer will be managed by the instance. If a policer is not parented or is

parented to a non-existent arbiter, the policer will be orphaned and will not be subject to bandwidth control by the policy instance.

Maximum Rate and Root Arbiter

The policer-control-policy supports an overall maximum rate (max-rate) that defines the total amount of bandwidth that may be distributed to all associated child policers. By default, that rate is set to max which provides an unlimited amount of bandwidth to the policers. Once the policy is created, an actual rate should be configured in order for the policy instances to be effective. At the SAP level, the maximum rate may be overridden on a per instance basis. For subscribers, the maximum rate may only be overridden on the subscriber profile which will then be applied to all instances associated with the profile.

The maximum rate is defined within the context of the root arbiter which is always present in a policer-control-policy. The system creates a parent policer which polices the output of all child policers attached to the policy instance to the configured rate. Child policers may be parented directly to the root arbiter (parent root) or parented to one of the tiered arbiters (parent arbiter-name). Since each tiered arbiter must be parented to either another tiered arbiter or the root arbiter (default), every parented child policer is associated with the root arbiter and therefore the root arbiter's parent policer.

Parent Policer PIR Leaky Bucket Operation

The parent policer is a single leaky bucket that monitors the aggregate throughput rate of the associated child policers. Forwarded packets increment the bucket by the size of each packet. The rate of the parent policer is implemented as a bucket decrement function which attempts to drain the bucket. If the rate of the packets flowing through the bucket is less than the decrement rate, the bucket does not accumulate depth. Each packet that flows through the bucket is accompanied by a derived discard threshold. If the current depth of the bucket is less than the discard threshold, the packet is allowed to pass through, retaining the colors derived from the packet's child policer. If the current depth is equal to or greater than the threshold value, the packet is colored red and the bucket depth is not incremented by the packet size. Also, any increased bucket depths in the child policer are canceled making any discard event an atomic function between the child and the parent.

Due to the fact that multiple thresholds are supported by the parent policer, the policer control policy is able to protect the throughput of higher priority child policers from the throughput of the lower priority child policers within the aggregate rate.

Tier 1 and Tier 2 Arbiters

As stated above, each child is attached either to the always available root arbiter or to an explicitly created tier 1 or tier 2 arbiter. Unlike the hardware parent policer based root arbiter, the arbiters at tier 1 and tier 2 are only represented in software and are meant to provide an arbitrary hierarchical bandwidth distribution capability. An arbiter created on tier 2 must parent to either to an arbiter on tier 1 or to the root arbiter. Arbiters created on tier 1 always parent to the root arbiter. In this manner, every arbiter ultimately is parented or grand-parented by the root arbiter.

Each tiered arbiter supports an optional rate parameter that defines a rate limit for all child arbiters or child policers associated with the arbiter. Child arbiters and policers attached to the arbiter have a level attribute that defines the strict level at which the child is given bandwidth by the arbiter. Level 8 is the highest and 1 is the lowest. Also a weight attribute defines each child's weight at that strict level in order to determine how bandwidth is distributed to multiple children at that level when insufficient bandwidth is available to meet each child's required bandwidth.

Fair and Unfair Bandwidth Control

Each child policer supports three leaky buckets. The PIR bucket manages the policer's peak rate and maximum burst size, the CIR leaky bucket manages the policer's committed rate and committed burst size. The third leaky bucket is used by the policer control policy instance to manage the child policer's fair rate (FIR). When multiple child policers are attached to the root arbiter at the same priority level, the policy

instance uses each child's FIR bucket rate to control how much of the traffic forwarded by the policer is fair and how much is unfair.

In the simplest case where all the child policers in the same priority level are directly attached to the root arbiter, each child's FIR rate is set according to the child's weight divided by the sum of the active children's weights multiplied by the available bandwidth at the priority level. The result is that the FIR bucket will mark the appropriate amount of traffic for each child as fair-based on the weighted fair output of the policy instance.

The fair/unfair forwarding control in the root parent policer is accomplished by implementing two different discard thresholds for the priority. The first threshold is discard-unfair and the second is discard-all for packet associated with the priority level. As the parent policer PIR bucket fills (due the aggregate forwarded rate being greater than the parent policers PIR decrement rate) and the bucket depth reaches the first threshold, all unfair packets within the priority are discarded. This leaves room in the bucket for the fair packets to be forwarded.

In the more complex case where one or more tiered arbiters are attached at the priority level, the policer control policy instance must consider more than just the child policer weights associated with the attached arbiter. If the arbiter is configured with an aggregate rate limit that its children cannot exceed, the policer control policy instance will switch to calculating the rate each child serviced by the arbiter should receive and enforces that rate using each child policers PIR leaky bucket.

When the child policer PIR leaky bucket is used to limit the bandwidth for the child policer and the child's PIR bucket discard threshold is reached, packets associated with the child policer are discarded. The child policer's discarded packets do not consume depth in the child policer's CIR or FIR buckets. The child policers discarded packets are also prevented from impacting the parent policer and will not consume the aggregate bandwidth managed by the parent policer.

Parent Policer Priority Level Thresholds

As stated above, each child policer is attached either to the root arbiter or explicitly to one of the tier 1 or tier 2 arbiters. When attached directly to the root arbiter, its priority relative to all other child policers is indicated by the parenting level parameter. When attached through one of the tiered arbiters, the parenting hierarchy of the arbiters must be traced through to the ultimate attachment to the root arbiter. The parenting level parameter of the arbiter parented to the root arbiter defines the child policer's priority level within the parent policer.

The priority level is important since it defines the parent policer discard thresholds that will be applied at the parent policer. The parent policer has 8 levels of strict priority and each priority level has its own discard-unfair and discard-all thresholds. Each priority's thresholds are larger than the thresholds of the lower priority levels. This ensures that when the parent policer is discarding, it will be priority sensitive.

To visualize the behavior of the parent policer, picture that when the aggregate forwarding rate of all child policers is currently above the decrement rate of the parent PIR leaky bucket, the bucket depth will increase over time. As the bucket depth increases, it will eventually cross the lowest priority's discard-unfair threshold. If this amount of discard sufficiently lowers the remaining aggregate child policer rate, the parent PIR bucket will hover around this bucket depth. If however, the remaining aggregate child rate is still greater than the decrement rate, the bucket will continue to rise and eventually reach the lowest priority's discard-all threshold which will cause all packets associated with the priority level to be discarded (fair and unfair). Again, if the remaining aggregate child rate is less than or equal to the bucket decrement rate, the parent PIR bucket will hover around this higher bucket depth. If the remaining aggregate child rate is still higher than the decrement rate, the bucket will continue to rise through the remaining priority level discards until equilibrium is achieved.

As noted above, each child's rate feeding into the parent policer is governed by the child policer's PIR bucket decrement rate. The amount of bandwidth the child policer offers to the parent policer will not exceed the child policer's configured maximum rate.

Root Arbiter's Parent Policer's Priority Aggregate Thresholds

Each policer-control-policy root arbiter supports configurable aggregate priority thresholds which are used to control burst tolerance within each priority level. Two values are maintained per priority level; the shared-portion and the fair-portion. The shared-portion represents the amount of parent PIR bucket depth that is allowed to be consumed by both fair and unfair child packets at the priority level. The fair-portion represents the amount of parent PIR bucket depth that only the fair child policer packets may consume within the priority level. It should be noted that the fair and unfair child packets associated with a higher parent policer priority level may also consume the bucket depth set aside for this priority.

While the policy maintains a parent policer default or explicit configurable values for shared-portion and fair-portion within each priority level, it is possible that some priority levels will not be used within the parent policer. Most parent policer use cases require fewer than eight strict priority levels.

To derive the actual priority level discard-unfair and discard-all thresholds while only accounting for the actual in-use priority levels, the system maintains a child policer to parent policer association counter per priority level for each policer control policy instance. As a child policer is parented to either the root or a tiered arbiter, the system determines the parent policer priority level for the child policer and increments the association counter for that priority level on the parent policer instance.

The shared-portion for each priority level is affected by the parent policer global min-thresh-separation parameter that defines the minimum separation between any in-use discard thresholds. When more than one child policer is associated with a parent policer priority level, the shared-portion for that priority level will be the current value of min-thresh-separation. When only a single child policer is associated, the priority level's shared-portion is zero since all packets from the child will be marked fair and the discard-unfair threshold is meaningless. When the association counter is zero, both the shared-portion and the fair-portion for that priority level are zero since neither discard thresholds will be used. Whenever the association counter is greater than 0, the fair-portion for that priority level will be derived from the current value of the priority's mbs-contribution parameter and the global min-thresh-separation parameter.

Each priority level's discard-unfair and discard-all thresholds are calculated based on an accumulation of lower priorities shared-portions and fair-portions and the priority level's own shared-portion and fair-portion. The base threshold value for each priority level is equal to the sum of all lower priority level's shared-portions and fair-portions. The discard-unfair threshold is the priority level's base threshold plus the priority level's shared-portion. The discard-all threshold for the priority level is the priority level's base threshold plus both the shared-portion and fair-portion values of the priority. As can be seen, an in-use priority level's thresholds are always greater than the thresholds of lower priority levels.

Policer Control Policy Application

A policer-control-policy may be applied on any Ethernet ingress or egress SAP that is associated with a port (or ports in the case of LAG).

The **no** form of this command removes a non-associated policer control policy from the system. The command will not execute when policer-name is currently associated with any SAP or subscriber management sub-profile context.

Parameters

policy-name

Specifies the policy name. Each policer-control-policy must be created with a unique policy name. The name must be given as policy-name must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system will enter that policy's context for editing purposes. If policy-name does not exist, the system will attempt to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode.

Platforms

7705 SAR Gen 2

policer-control-policy

Syntax

policer-control-policy *policy-name*

no policer-control-policy

Context

[Tree] (config>service>vprn>if>sap>ingress policer-control-policy)

[Tree] (config>service>vprn>if>sap>egress policer-control-policy)

Full Context

configure service vprn interface sap ingress policer-control-policy

configure service vprn interface sap egress policer-control-policy

Description

This command, within the qos CLI node, is used to create, delete or modify policer control policies. A policer control policy is very similar to the scheduler-policy which is used to manage a set of queues by defining a hierarchy of virtual schedulers and specifying how the virtual schedulers interact to provide an aggregate SLA. In a similar fashion, the policer-control-policy controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs. The policy may also be applied to the ingress or egress context of a sub-profile.

Policer Control Policy Instances

On the SAP side, an instance of a policy is created each time a policy is applied. When applied to a sub-profile, an instance of the policy is created each time a subscriber successfully maps one or more hosts to the profile per ingress SAP.

Each instance of the policer-control-policy manages the policers associated with the object that owns the policy instance (SAP or subscriber). If a policer on the object is parented to an appropriate arbiter name that exists within the policy, the policer will be managed by the instance. If a policer is not parented or is parented to a non-existent arbiter, the policer will be orphaned and will not be subject to bandwidth control by the policy instance.

Maximum Rate and Root Arbiter

The policer-control-policy supports an overall maximum rate (max-rate) that defines the total amount of bandwidth that may be distributed to all associated child policers. By default, that rate is set to max which provides an unlimited amount of bandwidth to the policers. Once the policy is created, an actual rate should be configured in order for the policy instances to be effective. At the SAP level, the maximum rate may be overridden on a per instance basis. For subscribers, the maximum rate may only be overridden on the subscriber profile which will then be applied to all instances associated with the profile.

The maximum rate is defined within the context of the root arbiter which is always present in a policer-control-policy. The system creates a parent policer which polices the output of all child policers attached to the policy instance to the configured rate. Child policers may be parented directly to the root arbiter (parent root) or parented to one of the tiered arbiters (parent arbiter-name). Since each tiered arbiter must

be parented to either another tiered arbiter or the root arbiter (default), every parented child policer is associated with the root arbiter and therefore the root arbiter's parent policer.

Parent Policer PIR Leaky Bucket Operation

The parent policer is a single leaky bucket that monitors the aggregate throughput rate of the associated child policers. Forwarded packets increment the bucket by the size of each packet. The rate of the parent policer is implemented as a bucket decrement function which attempts to drain the bucket. If the rate of the packets flowing through the bucket is less than the decrement rate, the bucket does not accumulate depth. Each packet that flows through the bucket is accompanied by a derived discard threshold. If the current depth of the bucket is less than the discard threshold, the packet is allowed to pass through, retaining the colors derived from the packet's child policer. If the current depth is equal to or greater than the threshold value, the packet is colored red and the bucket depth is not incremented by the packet size. Also, any increased bucket depths in the child policer are canceled making any discard event an atomic function between the child and the parent.

Due to the fact that multiple thresholds are supported by the parent policer, the policer control policy is able to protect the throughput of higher priority child policers from the throughput of the lower priority child policers within the aggregate rate.

Tier 1 and Tier 2 Arbiters

As stated above, each child is attached either to the always available root arbiter or to an explicitly created tier 1 or tier 2 arbiter. Unlike the hardware parent policer based root arbiter, the arbiters at tier 1 and tier 2 are only represented in software and are meant to provide an arbitrary hierarchical bandwidth distribution capability. An arbiter created on tier 2 must parent to either to an arbiter on tier 1 or to the root arbiter. Arbiters created on tier 1 always parent to the root arbiter. In this manner, every arbiter ultimately is parented or grand-parented by the root arbiter.

Each tiered arbiter supports an optional rate parameter that defines a rate limit for all child arbiters or child policers associated with the arbiter. Child arbiters and policers attached to the arbiter have a level attribute that defines the strict level at which the child is given bandwidth by the arbiter. Level 8 is the highest and 1 is the lowest. Also a weight attribute defines each child's weight at that strict level in order to determine how bandwidth is distributed to multiple children at that level when insufficient bandwidth is available to meet each child's required bandwidth.

Fair and Unfair Bandwidth Control

Each child policer supports three leaky buckets. The PIR bucket manages the policer's peak rate and maximum burst size, the CIR leaky bucket manages the policer's committed rate and committed burst size. The third leaky bucket is used by the policer control policy instance to manage the child policer's fair rate (FIR). When multiple child policers are attached to the root arbiter at the same priority level, the policy instance uses each child's FIR bucket rate to control how much of the traffic forwarded by the policer is fair and how much is unfair.

In the simplest case where all the child policers in the same priority level are directly attached to the root arbiter, each child's FIR rate is set according to the child's weight divided by the sum of the active children's weights multiplied by the available bandwidth at the priority level. The result is that the FIR bucket will mark the appropriate amount of traffic for each child as fair-based on the weighted fair output of the policy instance.

The fair/unfair forwarding control in the root parent policer is accomplished by implementing two different discard thresholds for the priority. The first threshold is discard-unfair and the second is discard-all for packet associated with the priority level. As the parent policer PIR bucket fills (due the aggregate forwarded rate being greater than the parent policers PIR decrement rate) and the bucket depth reaches the first threshold, all unfair packets within the priority are discarded. This leaves room in the bucket for the fair packets to be forwarded.

In the more complex case where one or more tiered arbiters are attached at the priority level, the policer control policy instance must consider more than just the child policer weights associated with the attached arbiter. If the arbiter is configured with an aggregate rate limit that its children cannot exceed, the policer control policy instance will switch to calculating the rate each child serviced by the arbiter should receive and enforces that rate using each child policers PIR leaky bucket.

When the child policer PIR leaky bucket is used to limit the bandwidth for the child policer and the child's PIR bucket discard threshold is reached, packets associated with the child policer are discarded. The child policer's discarded packets do not consume depth in the child policer's CIR or FIR buckets. The child policers discarded packets are also prevented from impacting the parent policer and will not consume the aggregate bandwidth managed by the parent policer.

Parent Policer Priority Level Thresholds

As stated above, each child policer is attached either to the root arbiter or explicitly to one of the tier 1 or tier 2 arbiters. When attached directly to the root arbiter, its priority relative to all other child policers is indicated by the parenting level parameter. When attached through one of the tiered arbiters, the parenting hierarchy of the arbiters must be traced through to the ultimate attachment to the root arbiter. The parenting level parameter of the arbiter parented to the root arbiter defines the child policer's priority level within the parent policer.

The priority level is important since it defines the parent policer discard thresholds that will be applied at the parent policer. The parent policer has 8 levels of strict priority and each priority level has its own discard-unfair and discard-all thresholds. Each priority's thresholds are larger than the thresholds of the lower priority levels. This ensures that when the parent policer is discarding, it will be priority sensitive.

To visualize the behavior of the parent policer, picture that when the aggregate forwarding rate of all child policers is currently above the decrement rate of the parent PIR leaky bucket, the bucket depth will increase over time. As the bucket depth increases, it will eventually cross the lowest priority's discard-unfair threshold. If this amount of discard sufficiently lowers the remaining aggregate child policer rate, the parent PIR bucket will hover around this bucket depth. If however, the remaining aggregate child rate is still greater than the decrement rate, the bucket will continue to rise and eventually reach the lowest priority's discard-all threshold which will cause all packets associated with the priority level to be discarded (fair and unfair). Again, if the remaining aggregate child rate is less than or equal to the bucket decrement rate, the parent PIR bucket will hover around this higher bucket depth. If the remaining aggregate child rate is still higher than the decrement rate, the bucket will continue to rise through the remaining priority level discards until equilibrium is achieved.

As noted above, each child's rate feeding into the parent policer is governed by the child policer's PIR bucket decrement rate. The amount of bandwidth the child policer offers to the parent policer will not exceed the child policer's configured maximum rate.

Root Arbiter's Parent Policer's Priority Aggregate Thresholds

Each policer-control-policy root arbiter supports configurable aggregate priority thresholds which are used to control burst tolerance within each priority level. Two values are maintained per priority level; the shared-portion and the fair-portion. The shared-portion represents the amount of parent PIR bucket depth that is allowed to be consumed by both fair and unfair child packets at the priority level. The fair-portion represents the amount of parent PIR bucket depth that only the fair child policer packets may consume within the priority level. It should be noted that the fair and unfair child packets associated with a higher parent policer priority level may also consume the bucket depth set aside for this priority.

While the policy maintains a parent policer default or explicit configurable values for shared-portion and fair-portion within each priority level, it is possible that some priority levels will not be used within the parent policer. Most parent policer use cases require fewer than eight strict priority levels.

To derive the actual priority level discard-unfair and discard-all thresholds while only accounting for the actual in-use priority levels, the system maintains a child policer to parent policer association counter per priority level for each policer control policy instance. As a child policer is parented to either the root or a tiered arbiter, the system determines the parent policer priority level for the child policer and increments the association counter for that priority level on the parent policer instance.

The shared-portion for each priority level is affected by the parent policer global min-thresh-separation parameter that defines the minimum separation between any in-use discard thresholds. When more than one child policer is associated with a parent policer priority level, the shared-portion for that priority level will be the current value of min-thresh-separation. When only a single child policer is associated, the priority level's shared-portion is zero since all packets from the child will be marked fair and the discard-unfair threshold is meaningless. When the association counter is zero, both the shared-portion and the fair-portion for that priority level are zero since neither discard thresholds will be used. Whenever the association counter is greater than 0, the fair-portion for that priority level will be derived from the current value of the priority's mbs-contribution parameter and the global min-thresh-separation parameter.

Each priority level's discard-unfair and discard-all thresholds are calculated based on an accumulation of lower priorities shared-portions and fair-portions and the priority level's own shared-portion and fair-portion. The base threshold value for each priority level is equal to the sum of all lower priority level's shared-portions and fair-portions. The discard-unfair threshold is the priority level's base threshold plus the priority level's shared-portion. The discard-all threshold for the priority level is the priority level's base threshold plus both the shared-portion and fair-portion values of the priority. As can be seen, an in-use priority level's thresholds are always greater than the thresholds of lower priority levels.

Policer Control Policy Application

A policer-control-policy may be applied on any Ethernet ingress or egress SAP that is associated with a port (or ports in the case of LAG).

The **no** form of this command removes a non-associated policer control policy from the system. The command will not execute when policer-name is currently associated with any SAP or subscriber management sub-profile context.

Parameters

policy-name

Specifies the policy name. Each policer-control-policy must be created with a unique policy name. The name must be given as policy-name must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system will enter that policy's context for editing purposes. If policy-name does not exist, the system will attempt to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode.

Platforms

7705 SAR Gen 2

policer-control-policy

Syntax

policer-control-policy *policy-name* [**create**]

no policer-control-policy *policy-name*

Context

[Tree] (config>qos policer-control-policy)

Full Context

configure qos policer-control-policy

Description

This command is used to create, delete, or modify policer control policies. The **policer-control-policy** controls the aggregate bandwidth available to a set of child policers. When created, the policy can be applied to ingress or egress SAPs. The policy can also be applied to the ingress or egress context of a sub-profile.

Parameters

policy-name

Each policer-control-policy must be created with a unique policy name. The *policy-name* must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system enters that policy's context for editing purposes. If policy-name does not exist, the system attempts to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode.

create

The **create** keyword is required when a new policy is being created and the system is configured for explicit object creation mode.

Platforms

7705 SAR Gen 2

policer-control-policy

Syntax

policer-control-policy *policy-name*

no policer-control-policy

Context

[Tree] (config>service>cust>multi-service-site>ingress policer-control-policy)

[Tree] (config>service>cust>multi-service-site>egress policer-control-policy)

Full Context

configure service customer multi-service-site ingress policer-control-policy

configure service customer multi-service-site egress policer-control-policy

Description

This command, within the QoS CLI node, is used to create, delete or modify policer control policies. A policer control policy is very similar to the scheduler-policy which is used to manage a set of queues by defining a hierarchy of virtual schedulers and specifying how the virtual schedulers interact to provide an aggregate SLA. In a similar fashion, the policer-control-policy controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs.

Policer Control Policy Instances

On the SAP side, an instance of a policy is created each time a policy is applied. When applied to a sub-profile, an instance of the policy is created each time a subscriber successfully maps one or more hosts to the profile per ingress SAP.

Each instance of the policer-control-policy manages the policers associated with the object that owns the policy instance (SAP or subscriber). If a policer on the object is parented to an appropriate arbiter name that exists within the policy, the policer will be managed by the instance. If a policer is not parented or is parented to a non-existent arbiter, the policer will be orphaned and not subject to bandwidth control by the policy instance.

Maximum Rate and Root Arbiter

The policer-control-policy supports an overall maximum rate (max-rate) that defines the total amount of bandwidth that may be distributed to all associated child policers. By default, that rate is set to max which provides an unlimited amount of bandwidth to the policers. Once the policy is created, an actual rate should be configured in order for the policy instances to be effective. At the SAP level, the maximum rate may be overridden on a per instance basis. For subscribers, the maximum rate may only be overridden on the subscriber profile which will then be applied to all instances associated with the profile.

The maximum rate is defined within the context of the root arbiter which is always present in a policer-control-policy. The system creates a parent policer which polices the output of all child policers attached to the policy instance to the configured rate. Child policers may be parented directly to the root arbiter (parent root) or parented to one of the tiered arbiters (parent arbiter-name). Since each tiered arbiter must be parented to either another tiered arbiter or the root arbiter (default), every parented child policer is associated with the root arbiter and thus the root arbiter’s parent policer.

Parent Policer PIR Leaky Bucket Operation

The parent policer is a single leaky bucket that monitors the aggregate throughput rate of the associated child policers. Forwarded packets increment the bucket by the size of each packet. The rate of the parent policer is implemented as a bucket decrement function which attempts to drain the bucket. If the rate of the packets flowing through the bucket is less than the decrement rate, the bucket does not accumulate depth. Each packet that flows through the bucket is accompanied by a derived discard threshold. If the current depth of the bucket is less than the discard threshold, the packet is allowed to pass through, retaining the colors derived from the packet’s child policer. If the current depth is equal to or greater than the threshold value, the packet is colored red and the bucket depth is not incremented by the packet size. Also, any increased bucket depths in the child policer are canceled making any discard event an atomic function between the child and the parent.

Due to the fact that multiple thresholds are supported by the parent policer, the policer control policy is able to protect the throughput of higher priority child policers from the throughput of the lower priority child policers within the aggregate rate.

Tier 1 and Tier 2 Arbiters

As stated above, each child is attached either to the always available root arbiter or to an explicitly created tier 1 or tier 2 arbiter. Unlike the hardware parent policer based root arbiter, the arbiters at tier 1 and tier 2 are only represented in software and are meant to provide an arbitrary hierarchical bandwidth distribution

capability. An arbiter created on tier 2 must parent to either to an arbiter on tier 1 or to the root arbiter. Arbiters created on tier 1 always parent to the root arbiter. In this manner, every arbiter ultimately is parented or grand-parented by the root arbiter.

Each tiered arbiter supports an optional rate parameter that defines a rate limit for all child arbiters or child policers associated with the arbiter. Child arbiters and policers attached to the arbiter have a level attribute that defines the strict level at which the child is given bandwidth by the arbiter. Level 8 is the highest and 1 is the lowest. Also a weight attribute defines each child's weight at that strict level in order to determine how bandwidth is distributed to multiple children at that level when insufficient bandwidth is available to meet each child's required bandwidth.

Fair and Unfair Bandwidth Control

Each child policer supports three leaky buckets. The PIR bucket manages the policer's peak rate and maximum burst size, the CIR leaky bucket manages the policer's committed rate and committed burst size. The third leaky bucket is used by the policer control policy instance to manage the child policer's fair rate (FIR). When multiple child policers are attached to the root arbiter at the same priority level, the policy instance uses each child's FIR bucket rate to control how much of the traffic forwarded by the policer is fair and how much is unfair.

In the simplest case where all the child policers in the same priority level are directly attached to the root arbiter, each child's FIR rate is set according to the child's weight divided by the sum of the active children's weights multiplied by the available bandwidth at the priority level. The result is that the FIR bucket will mark the appropriate amount of traffic for each child as fair based on the weighted fair output of the policy instance.

The fair/unfair forwarding control in the root parent policer is accomplished by implementing two different discard thresholds for the priority. The first threshold is discard-unfair and the second is discard-all for packet associated with the priority level. As the parent policer PIR bucket fills (due the aggregate forwarded rate being greater than the parent policers PIR decrement rate) and the bucket depth reaches the first threshold, all unfair packets within the priority are discarded. This leaves room in the bucket for the fair packets to be forwarded.

In the more complex case where one or more tiered arbiters are attached at the priority level, the policer control policy instance must consider more than just the child policer weights associated with the attached arbiter. If the arbiter is configured with an aggregate rate limit that its children cannot exceed, the policer control policy instance will switch to calculating the rate each child serviced by the arbiter should receive and enforces that rate using each child policers PIR leaky bucket.

When the child policer PIR leaky bucket is used to limit the bandwidth for the child policer and the child's PIR bucket discard threshold is reached, packets associated with the child policer are discarded. The child policer's discarded packets do not consume depth in the child policer's CIR or FIR buckets. The child policers discarded packets are also prevented from impacting the parent policer and will not consume the aggregate bandwidth managed by the parent policer.

Parent Policer Priority Level Thresholds

As stated above, each child policer is attached either to the root arbiter or explicitly to one of the tier 1 or tier 2 arbiters. When attached directly to the root arbiter, its priority relative to all other child policers is indicated by the parenting level parameter. When attached through one of the tiered arbiters, the parenting hierarchy of the arbiters must be traced through to the ultimate attachment to the root arbiter. The parenting level parameter of the arbiter parented to the root arbiter defines the child policer's priority level within the parent policer.

The priority level is important since it defines the parent policer discard thresholds that will be applied at the parent policer. The parent policer has 8 levels of strict priority and each priority level has its own discard-

unfair and discard-all thresholds. Each priority's thresholds are larger than the thresholds of the lower priority levels. This ensures that when the parent policer is discarding, it will be priority sensitive.

To visualize the behavior of the parent policer, picture that when the aggregate forwarding rate of all child policers is currently above the decrement rate of the parent PIR leaky bucket, the bucket depth will increase over time. As the bucket depth increases, it will eventually cross the lowest priority's discard-unfair threshold. If this amount of discard sufficiently lowers the remaining aggregate child policer rate, the parent PIR bucket will hover around this bucket depth. If however, the remaining aggregate child rate is still greater than the decrement rate, the bucket will continue to rise and eventually reach the lowest priority's discard-all threshold which will cause all packets associated with the priority level to be discarded (fair and unfair). Again, if the remaining aggregate child rate is less than or equal to the bucket decrement rate, the parent PIR bucket will hover around this higher bucket depth. If the remaining aggregate child rate is still higher than the decrement rate, the bucket will continue to rise through the remaining priority level discards until equilibrium is achieved.

As noted above, each child's rate feeding into the parent policer is governed by the child policer's PIR bucket decrement rate. The amount of bandwidth the child policer offers to the parent policer will not exceed the child policer's configured maximum rate.

Root Arbiter's Parent Policer's Priority Aggregate Thresholds

Each policer-control-policy root arbiter supports configurable aggregate priority thresholds which are used to control burst tolerance within each priority level. Two values are maintained per priority level; the shared-portion and the fair-portion. The shared-portion represents the amount of parent PIR bucket depth that is allowed to be consumed by both fair and unfair child packets at the priority level. The fair-portion represents the amount of parent PIR bucket depth that only the fair child policer packets may consume within the priority level. It should be noted that the fair and unfair child packets associated with a higher parent policer priority level may also consume the bucket depth set aside for this priority.

While the policy maintains a parent policer default or explicit configurable values for shared-portion and fair-portion within each priority level, it is possible that some priority levels will not be used within the parent policer. Most parent policer use cases require fewer than eight strict priority levels.

In order to derive the actual priority level discard-unfair and discard-all thresholds while only accounting for the actual in-use priority levels, the system maintains a child policer to parent policer association counter per priority level for each policer control policy instance. As a child policer is parented to either the root or a tiered arbiter, the system determines the parent policer priority level for the child policer and increments the association counter for that priority level on the parent policer instance.

The shared-portion for each priority level is affected by the parent policer global min-thresh-separation parameter that defines the minimum separation between any in-use discard thresholds. When more than one child policer is associated with a parent policer priority level, the shared-portion for that priority level will be the current value of min-thresh-separation. When only a single child policer is associated, the priority level's shared-portion is zero since all packets from the child will be marked fair and the discard-unfair threshold is meaningless. When the association counter is zero, both the shared-portion and the fair-portion for that priority level are zero since neither discard thresholds will be used. Whenever the association counter is greater than 0, the fair-portion for that priority level will be derived from the current value of the priority's mbs-contribution parameter and the global min-thresh-separation parameter.

Each priority level's discard-unfair and discard-all thresholds are calculated based on an accumulation of lower priorities shared-portions and fair-portions and the priority level's own shared-portion and fair-portion. The base threshold value for each priority level is equal to the sum of all lower priority level's shared-portions and fair-portions. The discard-unfair threshold is the priority level's base threshold plus the priority level's shared-portion. The discard-all threshold for the priority level is the priority level's base threshold plus both the shared-portion and fair-portion values of the priority. As can be seen, an in-use priority level's thresholds are always greater than the thresholds of lower priority levels.

Policer Control Policy Application

A policer-control-policy may be applied on any Ethernet ingress or egress SAP that is associated with a port (or ports in the case of LAG).

The **no** form of the command removes a non-associated policer control policy from the system. The command will not execute when policer-name is currently associated with any SAP context.

Parameters

policy-name

Specifies the policy name up to 32 characters in length. Each policer-control-policy must be created with a unique policy name. The name must adhere to the system policy ASCII naming requirements. If the defined policy name already exists, the system will enter that policy's context for editing purposes. If policy name does not exist, the system will attempt to create a policy with the specified name.

Platforms

7705 SAR Gen 2

22.13 policer-override

policer-override

Syntax

[**no**] **policer-override**

Context

[Tree] (config>card>fp>ingress>access>queue-group policer-override)

[Tree] (config>card>fp>ingress>network>queue-group policer-override)

Full Context

configure card fp ingress access queue-group policer-override

configure card fp ingress network queue-group policer-override

Description

This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to one or more policers created on the SAP through the sap-ingress or sap-egress QoS policies.

The **no** form of this command removes any existing policer overrides.

Default

no policer-override

Platforms

7705 SAR Gen 2

policer-override

Syntax

[**no**] **policer-override**

Context

[Tree] (config>service>epipe>sap>ingress policer-override)

[Tree] (config>service>epipe>sap>egress policer-override)

Full Context

configure service epipe sap ingress policer-override

configure service epipe sap egress policer-override

Description

This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to one or more policers created on the SAP through the sap-ingress or sap-egress QoS policies.

The **no** form of this command is used to remove any existing policer overrides.

Default

no policer-overrides

Platforms

7705 SAR Gen 2

policer-override

Syntax

[**no**] **policer-override**

Context

[Tree] (config>service>vpls>sap>ingress policer-override)

[Tree] (config>service>vpls>sap>egress policer-override)

Full Context

configure service vpls sap ingress policer-override

configure service vpls sap egress policer-override

Description

This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to one or more policers created on the SAP through the sap-ingress or sap-egress QoS policies.

The **no** form of this command is used to remove any existing policer overrides.

Default

no policer-overrides

Platforms

7705 SAR Gen 2

policer-override

Syntax

[**no**] **policer-override**

Context

[Tree] (config>service>ies>if>sap>ingress policer-override)

[Tree] (config>service>ies>if>sap>egress policer-override)

Full Context

configure service ies interface sap ingress policer-override

configure service ies interface sap egress policer-override

Description

This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to one or more policers created on the SAP through the sap-ingress or sap-egress QoS policies.

The **no** form of this command is used to remove any existing policer overrides.

Default

no policer-override

Platforms

7705 SAR Gen 2

policer-override

Syntax

[**no**] **policer-override**

Context

[Tree] (config>service>vprn>if>sap>ingress policer-override)

[Tree] (config>service>vprn>if>sap>egress policer-override)

Full Context

configure service vprn interface sap ingress policer-override

configure service vprn interface sap egress policer-override

Description

This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to one or more policers created on the SAP through the sap-ingress or sap-egress QoS policies.

The **no** form of this command is used to remove any existing policer overrides.

Default

no policer-override

Platforms

7705 SAR Gen 2

22.14 policy

policy

Syntax

policy *msap-policy-name*

no policy

Context

[Tree] (config>service>vpls>sap>msap-defaults policy)

Full Context

configure service vpls sap msap-defaults policy

Description

This command sets default msap-policy for all subscribers created based on trigger packets received on the specified capture-sap in case the corresponding VSA is not included in the RADIUS authentication response. This command is applicable to capture SAP only.

Default

no policy

Platforms

7705 SAR Gen 2

policy

Syntax

policy *vrrp-policy-id*

no policy

Context

[Tree] (config>service>ies>if>ipv6>vrrp policy)

Full Context

configure service ies interface ipv6 vrrp policy

Description

This command creates VRRP control policies. The VRRP policy ID must be created by the policy command prior to association with the virtual router instance.

The policy command provides the ability to associate a VRRP priority control policy to a virtual router instance. The policy may be associated with more than one virtual router instance. The priority events within the policy either override or diminish the base-priority dynamically affecting the in-use priority. As priority events clear in the policy, the in-use priority may eventually be restored to the base-priority value.

The policy command is only available in the non-owner **vrrp** *virtual-router-id* nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed by VRRP priority control policies. For non-owner virtual router instances, if the policy command is not executed, the base-priority will be used as the in-use priority.

The **no** form of this command removes any existing VRRP priority control policy association from the virtual router instance. All such associations must be removed prior to the policy being deleted from the system.

Parameters

vrrp-policy-id

The vrrp-policy-id parameter associated the corresponding VRRP priority control policy-id with the virtual router instance. The vrrp-policy-id must already exist in the system for the policy command to be successful.

Values 1 to 9999

Platforms

7705 SAR Gen 2

policy

Syntax

policy *vrrp-policy-id*

no policy

Context

[Tree] (config>service>ies>if>vrrp policy)

Full Context

configure service ies interface vrrp policy

Description

This command creates VRRP control policies. The VRRP policy ID must be created by the policy command prior to association with the virtual router instance.

The policy command provides the ability to associate a VRRP priority control policy to a virtual router instance. The policy may be associated with more than one virtual router instance. The priority events within the policy either override or diminish the base-priority dynamically affecting the in-use priority. As priority events clear in the policy, the in-use priority may eventually be restored to the base-priority value.

The policy command is only available in the non-owner **vrrp** *virtual-router-id* nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed by VRRP priority control policies. For non-owner virtual router instances, if the policy command is not executed, the base-priority will be used as the in-use priority.

The **no** form of this command removes any existing VRRP priority control policy association from the virtual router instance. All such associations must be removed prior to the policy being deleted from the system.

Parameters

vrrp-policy-id

The vrrp-policy-id parameter associated the corresponding VRRP priority control policy-id with the virtual router instance. The vrrp-policy-id must already exist in the system for the policy command to be successful.

Values 1 to 9999

Platforms

7705 SAR Gen 2

policy

Syntax

policy *policy-name*

no policy

Context

[Tree] (config>service>vprn>bgp>next-hop-res policy)

Full Context

configure service vprn bgp next-hop-resolution policy

Description

This command specifies the name of a policy statement to use with the BGP next-hop resolution process. The policy controls which IP routes in RTM are eligible to resolve the BGP next-hop addresses of IPv4 and IPv6 routes. The policy has no effect on the resolution of BGP next-hops to MPLS tunnels. If a BGP next-hop of an IPv4 or IPv6 route R is resolved in RTM and the longest matching route for the next-hop address is an IP route N that is rejected by the policy then route R is unresolved; if the route N is accepted by the policy then it becomes the resolving route for R.

The default next-hop resolution policy (when the **no policy** command is configured) is to use the longest matching active route in RTM that is not a BGP route (unless **use-bgp-routes** is configured), an aggregate route or a subscriber management route.

Default

no policy

Parameters

policy-name

Specifies the route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes. Route policies are configured in the **config>router>policy-options** context.

Platforms

7705 SAR Gen 2

policy

Syntax

policy *vrrp-policy-id*

no policy

Context

[Tree] (config>service>vprn>if>vrrp policy)

[Tree] (config>service>vprn>if>ipv6>vrrp policy)

Full Context

configure service vprn interface vrrp policy

configure service vprn interface ipv6 vrrp policy

Description

This command associates a VRRP priority control policy with the virtual router instance (non-owner context only).

Parameters

vrrp-policy-id

Specifies a VRRP priority control policy.

Values 1 to 9999

Platforms

7705 SAR Gen 2

policy

Syntax

policy *vrrp-policy-id*

no policy

Context

[Tree] (config>router>if>ipv6>vrrp policy)

[Tree] (config>router>if>vrrp policy)

Full Context

configure router interface ipv6 vrrp policy

configure router interface vrrp policy

Description

This command adds a VRRP priority control policy association with the virtual router instance.

To further augment the virtual router instance base priority, VRRP priority control policies can be used to override or adjust the base priority value depending on events or conditions within the chassis.

The policy can be associated with more than one virtual router instance. The priority events within the policy either override or diminish the base priority set with the **priority** command dynamically affecting the in-use priority. As priority events clear in the policy, the in-use priority can eventually be restored to the base **priority** value.

The **policy** command is only available in the non-owner **vrrp** nodal context. The priority of **owner** virtual router instances is permanently set to 255 and cannot be changed by VRRP priority control policies. For non-owner virtual router instances, if the **policy** command is not executed, the base **priority** is used as the in-use priority.

The **no** form of the command removes existing VRRP priority control policy associations from the virtual router instance. All associations must be removed prior to deleting the policy from the system.

Default

no policy — No VRRP priority control policy is associated with the virtual router instance.

Parameters

vrrp-policy-id

The policy ID of the VRRP priority control expressed as a decimal integer. The *vrrp-policy-id* must already exist for the command to function.

Values 1 to 9999

Platforms

7705 SAR Gen 2

policy

Syntax

policy *policy-id* **context** *context-value*

policy *policy-id* **context name** *name*

no policy *policy-id*

Context

[Tree] (config>vrrp policy)

Full Context

configure vrrp policy

Description

This command creates the context to configure a VRRP priority control policy which is used to control the VRRP in-use priority based on priority control events. It is a parental node for the various VRRP priority control policy commands that define the policy parameters and priority event conditions.

The virtual router instance **priority** command defines the initial or base value to be used by non-owner virtual routers. This value can be modified by assigning a VRRP priority control policy to the virtual router instance. The VRRP priority control policy can override or diminish the base priority setting to establish the actual in-use priority of the virtual router instance.

The **policy** *policy-id* command must be created first, before it can be associated with a virtual router instance.

Because VRRP priority control policies define conditions and events that must be maintained, they can be resource intensive. The number of policies is limited to 1000.

The *policy-id* do not have to be consecutive integers. The range of available policy identifiers is from 1 to 9999.

The **no** form of the command deletes the specific *policy-id* from the system. The *policy-id* must be removed first from all virtual router instances before the **no policy** command can be issued. If the *policy-id* is associated with a virtual router instance, the command will fail.

Parameters

policy-id

The VRRP priority control ID expressed as a decimal integer that uniquely identifies this policy from any other VRRP priority control policy defined on the system. Up to 1000 policies can be defined.

Values 1 to 9999

context-value

Specifies the service ID to which this policy applies. A value of zero (0) means that this policy does not apply to a service but applies to the base router instance.

Values 1 to 2147483647

Platforms

7705 SAR Gen 2

policy

Syntax

policy *policy-name* [**create**] [**type** {**access-network | port**}]

no policy *policy-name*

Context

[Tree] (config>sys>security>dist-cpu-protection policy)

Full Context

configure system security dist-cpu-protection policy

Description

This command configures one of the maximum 18 Distributed CPU Protection (DCP) policies. These policies can be applied to objects such as SAPs, network interfaces or ports.

Parameters

policy-name

Specifies the name of the policy, up to 32 characters.

create

Keyword used to create a new policy.

type

Specifies the Distributed CPU protection type for the policy.

Values access-network — Specifies this is a distributed CPU protection policy for access or network interfaces.

port — Specifies this is a distributed CPU protection policy for ports.

Default access-network

Platforms

7705 SAR Gen 2

policy

Syntax

policy *policy-name*

no policy

Context

[Tree] (config>router>bgp>next-hop-resolution policy)

Full Context

configure router bgp next-hop-resolution policy

Description

This command specifies the policy statement name to use with the BGP next-hop resolution process. The policy determines the eligibility of IP routes in the RTM to resolve the BGP next-hop addresses of IPv4 and IPv6 routes. The policy has no effect on the resolution of BGP next hops to MPLS tunnels.

For example, if a BGP next hop of an IPv4 or IPv6 route R is resolved in RTM and the longest matching route for the next-hop address is an IP route N that is rejected by the policy then route R is unresolved. If the route N is accepted by the policy, it becomes the resolving route for R.

The **no** form of this command reverts to the default next-hop resolution policy, which uses the longest matching active route in RTM that is not a BGP route (unless **use-bgp-routes** is configured), an aggregate route or a subscriber management route.

Default

no policy

Parameters

policy-name

Specifies the route policy name. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

Platforms

7705 SAR Gen 2

policy

Syntax

policy *plcy-or-long-expr*

no policy

Context

[Tree] (config>router>policy-options>policy-statement>entry>from policy)

Full Context

configure router policy-options policy-statement entry from policy

Description

This command is used to call another policy by name and evaluate it as a subroutine, or to evaluate a logical expression of subroutine policies.

If the result of the subroutine evaluation is an 'accept', then the route is considered to match the entry in the parent policy that called the subroutine. If the result of the subroutine evaluation is a 'reject', then the route is considered a non-match of the entry in the parent policy that called the subroutine.

Up to 3 levels of subroutine calls are supported. If a subroutine at maximum depth has this command, it is automatically considered a non-match of all routes.

The **no** form of this command removes the policy statement as a match criterion.

Default

no policy

Parameters

plcy-or-long-expr

Specifies the name of a single **policy-statement** (up to 64 characters in length) or a policy logical expression (up to 255 characters in length) consisting of **policy-statement** names (enclosed in square brackets), logical operations ('and', 'or', 'not'), and parentheses for grouping.

Platforms

7705 SAR Gen 2

policy

Syntax

[no] policy *association-name*

Context

[Tree] (config>router>pcep>pcc>pce-assoc policy)

Full Context

configure router pcep pcc pce-associations policy

Description

This command creates a named PCE policy association from which the parameters for specified policy association are configured.

The **no** form of the command deletes the specified policy association.

Parameters

association-name

Specifies the name of the policy association, up to 32 characters.

Platforms

7705 SAR Gen 2

policy

Syntax

[no] policy *policy-assoc-name*

Context

[Tree] (config>router>mpls>lsp-template>pce-assoc policy)

[Tree] (config>router>mpls>lsp>pce-assoc policy)

Full Context

configure router mpls lsp-template pce-associations policy

configure router mpls lsp pce-associations policy

Description

This command binds the LSP to a named policy association. The policy association name must exist under the PCC. Up to five policy associations can be configured per LSP.

The **no** form of the command removes the LSP binding from the specified policy association.

Parameters

policy-assoc-name

Specifies the name of an existing policy association, up to 32 characters.

Platforms

7705 SAR Gen 2

22.15 policy-options

policy-options

Syntax

[**no**] **policy-options**

Context

[Tree] (config>router policy-options)

Full Context

configure router policy-options

Description

Commands in this context configure route policies. Route policies are applied to the routing protocol.

The **no** form of this command deletes the route policy configuration.

Platforms

7705 SAR Gen 2

22.16 policy-reference-checks

policy-reference-checks

Syntax

[**no**] **policy-reference-checks**

Context

[Tree] (config>router policy-reference-checks)

Full Context

configure router policy-reference-checks

Description

This command checks policy references to ensure that a policy exists and displays a CLI error if the policy does not exist. Enabling this option protects against accidentally referencing a missing or misspelled policy, that can lead to unexpected results when the policy is evaluated.

The **no** version of this command disables policy reference checks and allows policies that do not exist to be referenced.

Default

no policy-reference-checks

Platforms

7705 SAR Gen 2

22.17 policy-statement

policy-statement

Syntax

[**no**] **policy-statement** *name*

Context

[Tree] (config>router>policy-options policy-statement)

Full Context

configure router policy-options policy-statement

Description

This command creates the context to configure a route policy statement.

Route policy statements control the flow of routing information to and from a specific protocol, set of protocols, or to a specific BGP neighbor.

The **policy-statement** is a logical grouping of match and action criteria. A single **policy-statement** can affect routing in one or more protocols and/or one or more protocols peers/neighbors. A single **policy-statement** can also affect both the import and export of routing information.

The **no** form of this command deletes the policy statement.

Default

no policy-statement

Parameters

name

Specifies the route policy statement name. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

22.18 policy-variables

policy-variables

Syntax

policy-variables

Context

[Tree] (config>router>policy-options>policy-statement>entry>from policy-variables)

Full Context

configure router policy-options policy-statement entry from policy-variables

Description

Commands in this context configure **policy-variables** parameters.

The **no** form of this command removes all policy variables.

Platforms

7705 SAR Gen 2

22.19 poll

poll

Syntax

poll ca *ca-profile-name*

Context

[Tree] (admin>certificate>cmpv2 poll)

Full Context

admin certificate cmpv2 poll

Description

This command polls the status of the pending CMPv2 request toward the specified CA.

If the response is ready, this command will resume the CMPv2 protocol exchange with server as the original command would do. The requests could be also still be pending as a result, then this command could be used again to poll the status.

SR OS allows only one pending CMP request per CA, which means no new request is allowed when a pending request is present.

Parameters

ca-profile-name

Specifies a ca-profile name up to 32 characters.

Platforms

7705 SAR Gen 2

22.20 poll-interval

poll-interval

Syntax

poll-interval *seconds*

no poll-interval

Context

[Tree] (config>service>vprn>ospf>area>if poll-interval)

[Tree] (config>service>vprn>ospf3>area>if poll-interval)

Full Context

configure service vprn ospf area interface poll-interval

configure service vprn ospf3 area interface poll-interval

Description

This command configures the poll interval, in seconds. The poll interval is the time between two Hello packets to a dead (non-adjacent) OSPF NBMA neighbor. The default value of the poll interval timer is higher than the hello interval timer to avoid wasting bandwidth on non-broadcast networks, since OSPF messages are unicast to each configured neighbor. The poll interval timer is used only on **non-broadcast** interface types and has no effect if configured on other interface types.

The **no** form of this command removes the **poll-interval** configuration.

Default

120

Parameters

seconds

Specifies the poll interval, in seconds.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

poll-interval

Syntax

poll-interval *seconds*

no poll-interval

Context

[Tree] (config>router>ospf3>area>interface poll-interval)

[Tree] (config>router>ospf>area>interface poll-interval)

Full Context

configure router ospf3 area interface poll-interval

configure router ospf area interface poll-interval

Description

This command configures the poll interval, in seconds. The poll interval is the time between two Hello packets to a dead (non-adjacent) OSPF NBMA neighbor. The default value of the poll interval timer is higher than the hello interval timer to avoid wasting bandwidth on non-broadcast networks, since OSPF messages are unicast to each configured neighbor. The poll interval timer is used only on **non-broadcast** interface types and has no effect if configured on other interface types.

The **no** form of this command removes the **poll-interval** configuration.

Default

120

Parameters

seconds

Specifies the poll interval, in seconds.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

22.21 pool

pool

Syntax

pool *pool-name* [**create**]

no pool *pool-name*

Context

[Tree] (config>router>dhcp6>server pool)

[Tree] (config>service>vprn>dhcp6>server pool)

[Tree] (config>router>dhcp>server pool)

[Tree] (config>service>vprn>dhcp>server pool)

Full Context

configure router dhcp6 local-dhcp-server pool

configure service vprn dhcp6 local-dhcp-server pool

configure router dhcp local-dhcp-server pool

configure service vprn dhcp local-dhcp-server pool

Description

This command configures a DHCP address pool on the router.

The **no** form of this command removes the pool name from the configuration.

Parameters

pool name

Specifies the name of this IP address pool. Allowed values are any string, up to 32 characters.

create

Keyword used to create the pool. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

pool

Syntax

pool [*name*]

Context

[Tree] (config>port>access>ingress pool)

[Tree] (config>port>network>egress pool)

[Tree] (config>port>access>egress pool)

Full Context

configure port access ingress pool

configure port network egress pool

configure port access egress pool

Description

This command configures pool policies.

On the MDA level, access and network egress and access ingress pools are only allocated on channelized MDAs. Network ingress pools are allocated on the FP level for non-channelized MDAs.

Default

pool default

Parameters

name

If specified, the name must be **default**.

Platforms

7705 SAR Gen 2

pool

Syntax

pool [*name*]

Context

[Tree] (config>card>fp>ingress>network pool)

Full Context

configure card fp ingress network pool

Description

This command configures the per-FP network ingress pool.

Default

pool default

Parameters

name

If specified, the name must be **default**.

Platforms

7705 SAR Gen 2

pool

Syntax

pool *nat-pool-name* [**nat-group** *nat-group-id* **type** *pool-type* [**applications** *applications*] [**create**]

no pool *nat-pool-name*

Context

[Tree] (config>service>vprn>nat>outside pool)

Full Context

configure service vprn nat outside pool

Description

This command configures a NAT pool.

Parameters

nat-pool-name

Specifies the NAT pool name.

Values 32 chars max

nat-group-id

Specifies the NAT group ID.

Values 1 to 4

create

This parameter must be specified to create the instance.

pool-type

Specifies the pool type.

Values large-scale, l2-aware, wlan-gw-anchor

applications

Specifies the application.

Values agnostic

flexible-port-allocation

create

Keyword used to create the pool.

Platforms

7705 SAR Gen 2

pool

Syntax

pool *nat-pool-name* **nat-group** *nat-group-id* **type** *pool-type* [**applications** *applications*] [**create**]

no pool *nat-pool-name*

Context

[Tree] (config>router>nat>outside pool)

Full Context

configure router nat outside pool

Description

This command creates a NAT pool in the outside routing context. The NAT pool defines the parameters that will be used for IP address and port translation within the pool.

Parameters

nat-pool-name

Specifies the NAT pool name, up to 32 characters.

nat-group-id

Specifies the NAT group ID.

Values 1 to 4

create

Creates the instance.

pool-type

Species the pool type.

Values large-scale, l2-aware, wlan-gw-anchor

applications

This creation-time parameter configures the NAT pool for protocol agnostic operation. The IP addresses are translated in 1:1 fashion regardless of the protocol. No ports are translated for TCP or UDP traffic. Traffic through the pool can be initiated from inside or outside. When nat-pool is configured in agnostic mode, certain parameters in the pool are pre-set and cannot be changed:

- mode one-to-one

- port-forwarding-range 0
- port-reservation blocks 1
- subscriber-limit 1
- deterministic port-reservation 65536.

This pool is used to configure static 1:1 NAT, where the operator have the control of the mapping between the inside and outside IP addresses. The static IP address mapping is using CLI constructs used in deterministic NAT (prefix and map deterministic NAT commands in the inside routing context).

ALG for TCP/UDP is supported in the protocol agnostic pool.

Values agnostic

Platforms

7705 SAR Gen 2

pool

Syntax

pool *nat-pool-name* **service-name** *service-name*

pool *nat-pool-name* **router** *router-instance*

no pool

Context

[Tree] (config>service>nat>nat-policy pool)

Full Context

configure service nat nat-policy pool

Description

This command configures the NAT pool of this policy.

Parameters

nat-pool-name

Specifies the name of the NAT pool, up to 32 characters.

router-instance

Specifies the router instance the pool belongs to, either by router name or service ID.

Values 1 to 2147483647 svc-name — a string up to 64 characters.

Values *router-name*: "Base" | "management"

Default Base

service-name

Specifies the name of the service, up to 64 characters.

Platforms

7705 SAR Gen 2

22.22 pool-name

pool-name

Syntax

[**no**] **pool-name**

Context

[Tree] (config>service>ies>if>dhcp>option>vendor pool-name)

[Tree] (config>service>vprn>if>dhcp>option>vendor pool-name)

Full Context

configure service ies interface dhcp option vendor-specific-option pool-name

configure service vprn interface dhcp option vendor-specific-option pool-name

Description

This command sends the pool name in the Nokia vendor specific sub-option of the DHCP relay packet.

The **no** form of this command reverts to the default.

Platforms

7705 SAR Gen 2

pool-name

Syntax

[**no**] **pool-name**

Context

[Tree] (config>router>if>dhcp>option>vendor-specific-option pool-name)

Full Context

configure router interface dhcp option vendor-specific-option pool-name

Description

This command enables the sending of the pool name in the Nokia vendor-specific suboption of the DHCP relay packet.

The **no** form of this command disables the feature.

Default

no pool-name

Platforms

7705 SAR Gen 2

22.23 pop

pop

Syntax

[**no**] **pop**

Context

[Tree] (config>router>mpls>if>label-map pop)

Full Context

configure router mpls interface label-map pop

Description

This command specifies that the incoming label must be popped (removed). No label stacking is supported for a static LSP. The service header follows the top label. Once the label is popped, the packet is forwarded based on the service header.

The **no** form of this command removes the **pop** action for the *in-label*.

Platforms

7705 SAR Gen 2

22.24 populate

populate

Syntax

populate {**static** | **dynamic** | **evpn**} [**route-tag** [*1..255*]]

no populate {**static** | **dynamic** | **evpn**}

Context

[Tree] (config>service>vprn>if>ipv6>nd-host-route populate)

[Tree] (config>service>vprn>if>arp-host-route populate)

[Tree] (config>service>ies>if>arp-host-route populate)

Full Context

configure service vprn interface ipv6 nd-host-route populate

configure service vprn interface arp-host-route populate

configure service ies interface arp-host-route populate

Description

This command enables the creation of ARP/ND host-route entries in the route-table out of a certain ARP/ND entry type.

The **no** form of this command reverts to the default.

Default

no populate

Parameters

evpn

Enables the creation of ARP-ND host routes in the route table out of EVPN ARP/ND entries (entries learned from EVPN MAC/IP routes).

dynamic

Enables the creation of ARP-ND host routes in the route table out of dynamic ARP/ND entries (learned from received ARP/ND messages from the hosts).

static

Enables the creation of ARP-ND host routes in the route table out of configured static ARP/ND entries.

route-tag [*1..255*]

Specifies the route tag that is added in the route table for ARP-ND host routes of type **evpn**, **dynamic**, or **static**. This tag can be matched on BGP VRF export and BGP peer export policies.

Platforms

7705 SAR Gen 2

22.25 port

port

Syntax

port *port-id* [**sync-tag** *sync-tag*] [**create**]

no port *port-id*]

Context

[Tree] (config>redundancy>multi-chassis>peer>sync port)

Full Context

configure redundancy multi-chassis peer sync port

Description

This command specifies the port to be synchronized with the multi-chassis peer and a synchronization tag to be used while synchronizing this port with the multi-chassis peer.

Parameters

port-id

Specifies the port to be synchronized with the multi-chassis peer.

Values

port-id	slot/mda/port	
lag-id	lag-id	
	lag	keyword
	id	1 to 200
pw-id	pw-*id*	
	pw	keyword
	id	1 to 10239

sync-tag

Specifies a synchronization tag, up to 32 characters in length, to be used while synchronizing this port with the multi-chassis peer.

create

Creates an entry; mandatory while creating an entry.

Platforms

7705 SAR Gen 2

port

Syntax

[**no**] **port** {*port-id* | *aps-id* | *connector-port-id*}

Context

[Tree] (config port)

Full Context

configure port

Description

This command enables access to the context to configure ports, multilink bundles, and bundle protection groups (BPGs). Before a port can be configured, the chassis slot must be provisioned with a valid card type and the MDA parameter must be provisioned with a valid MDA type.

Default

No ports are configured. All ports must be explicitly configured and enabled.

Parameters

port-id

Specifies the physical port ID in the following format:

Values *slot*/*mda*/*port* [.*channel*]

for GNSS RF ports:

A/gnss or **B/gnss**

aps-id

This option configures APS on unbundled SONET/SDH ports. All SONET-SDH port parameters, with certain exceptions, for the working and protection circuit ports must be configured in the **config>port>aps-id** context. The working and protection circuit ports inherit all those parameters configured. The exception parameters for the working and protect circuits can be configured in the **config>port>sonet-sdh** context. Exception list commands include:

- clock-source
- [no] loopback
- [no] report-alarm
- section-trace
- [no] threshold

When an **configure port aps-id** is created all applicable parameters under the port CLI tree (including parameters under any submenus) assume **aps-id** defaults, or when those are not explicitly specified, default to SONET/SDH port defaults for any SONET port.

All but a few exception SONET/SDH parameters for the working channel port must be configured in the **configure port sonet-sdh** context. The protection channel inherits all the configured parameters. The exception parameters for the protection channel can be configured in the **configure port sonet-sdh** context.

Signal failure (SF) and signal degrade (SD) alarms are not enabled by default on POS interfaces. It is recommended to change the default alarm notification configuration for POS ports that belong to APS groups in order to be notified of SF/SD occurrences to be able to interpret the cause for an APS group to switch the active line.

For path alarms, modify the logical line aps-id in the **configure port aps-***id* **<sonet-sdh>path report-alarm** context. For example:

configure port **aps-1** sonet-sdh path report-alarm p-ais

For line alarms, separately, modify the 2 physical ports that are members of the logical aps-id port (the working and protect lines). APS reacts only to line alarms, not path alarms. For example:

configure port **1/2/3** sonet-sdh report-alarm lb2er-sd

configure port **4/5/6** sonet-sdh report-alarm lb2er-sd

If the SD and SF threshold rates must be modified, the changes must be performed at the line level on both the working and protect APS port member.

The **no** form of this command deletes an aps-*group-id* or bundle-aps-group-id. In order for an aps- *group-id* to be deleted,

The same rules apply for physical ports, bundles deletions apply to APS ports/bundles deletions (for example an aps-group-id must be shutdown, have no service configuration on it, and no path configuration on it). In addition working and protection circuits must be removed before an aps-group-id may be removed.

Values **port aps**-*group-id* **aps**: keyword where *group-id*: 1 to 64

Example: **port aps**-64

connector-port-id

Specifies the physical port of a connector in the following format.

Values *slot*/*mda*/*connector*/*port*

Platforms

7705 SAR Gen 2

port

Syntax

port *port-id*

no port

Context

[Tree] (config>port-xc>pxc port)

Full Context

configure port-xc pxc port

Description

This command configures the referenced Ethernet port as a loopback or a cross-connect port (PXC). When this command is executed, the system automatically creates two PXC subports under this Ethernet port.

The physical PXC port does not require any external connectivity or optical transceivers to function properly. Consequently, all optic-related alarms are disabled on the port.

The physical PXC port is automatically configured as a hybrid port. The MTU is preset to 9212 bytes, the encapsulation type is set to dot1q, and dot1x tunneling is turned on.

A single physical port can be associated with more than one PXC. In other words, multiple PXCs are supported per physical port. Because PXC subports use a single physical port to transmit traffic in both directions, the nominal port bandwidth is asymmetrically divided between the two directions. For example, a 10 Gb/s Ethernet port in PXC mode can accommodate 9 Gb/s of traffic in one direction and 1 Gb/s in the other. Any other ratio can be achieved as long as the sum of the bandwidth of the two PXC subports does not exceed the bandwidth capacity of the physical port (10 Gb/s in this case).

Since the PXC uses a single physical port to transmit traffic in both directions, the nominal port bandwidth is asymmetrically divided between the two directions. For example, a 10 Gb/s Ethernet port in PXC mode can accommodate 9 Gb/s of traffic in one direction and 1 Gb/s in the other. Any other ratio can be achieved as long as the sum of the bandwidth of the two PXC subports does not exceed the bandwidth capacity of the physical port (10 Gb/s in this case).

The following rules apply to PXC port configurations:

- Only unused physical ports (not associated with an interface or SAP) can be referenced inside of a PXC ID configuration.
- The physical port cannot be removed from a PXC ID configuration if the corresponding PXC subports are currently in use.
- A physical port cannot be used outside the configured PXC context. For example, a regular IP interface cannot use this physical port, or a SAP on that port cannot be associated with a service.

The **no** form of this command removes the port ID from the configuration.

Parameters

port-id

Specifies the physical port in the *slot*/*mda*/*port* format.

Platforms

7705 SAR Gen 2

port

Syntax

port *port-id* [*port-id*] [**priority** *priority*] [**sub-group** *sub-group-id*] [**hash-weight** *weight*]

no port *port-id* [*port-id*]

Context

[Tree] (config>lag port)

Full Context

configure lag port

Description

This command adds ports to a Link Aggregation Group (LAG).

The port configuration of the first port added to the LAG is used as a basis to compare to subsequently added ports. If a discrepancy is found with a newly added port, that port will not be added to the LAG.

Multiple (space separated) ports can be added or removed from the LAG link assuming the maximum of number of ports is not exceeded.

Ports that are part of a LAG must be configured with auto-negotiate limited or disabled.

The **no** form of this command removes ports from the LAG.

Default

No ports are defined as members of a LAG.

Parameters

port-id

Specifies the port ID.

The maximum number of ports in a LAG depends on the platform type, the hardware deployment, and the SR OS software release. Adding a port over the maximum allowed per given router or switch is blocked. Some platforms support double port scale for specific port types on LAGs with LAG ID in the range of 1 to 64 inclusive. Up to 16 ports can be specified in a single statement, up to 64 ports total.

priority

Specifies the port priority used by LACP. The port priority is also used to determine the primary port. The port with the lowest priority is the primary port. In the event of a tie, the smallest port ID becomes the primary port.

Values 1 to 65535

sub-group-id

Identifies a LAG subgroup. When using subgroups in a LAG, they should only be configured on one side of the LAG, not both. Only having one side perform the active/standby selection guarantees a consistent selection and fast convergence. The active or

standby selection is signaled through LACP to the other side. The hold time should be configured when using subgroups to prevent the LAG going down when switching between active and standby subgroup since momentarily all ports are down in a LAG (break-before-make).

Values 1 to 8 identifies a LAG subgroup The **auto-iom** subgroup is defined based on the IOM (all ports of the same IOM are assigned to the same subgroup). The **auto-mda** subgroup is defined based on the MDA. (all ports of the same MDA are assigned to the same subgroup).

weight

Specifies the flow hashing distribution between LAG ports.

Values 1 to 100000, port-speed

Platforms

7705 SAR Gen 2

port

Syntax

port *port*

no port

Context

[Tree] (config>service>vprn>aaa>rmt-srv>radius port)

Full Context

configure service vprn aaa remote-servers radius port

Description

This command configures the UDP port number to contact the RADIUS server.

The **no** form of this command reverts to the default value.

Default

port 1812 (as specified in RFC 2865, *Remote Authentication Dial In User Service* (*RADIUS*))

Parameters

port

Specifies the UDP port number to contact the RADIUS server.

Values 1 to 65535

Platforms

7705 SAR Gen 2

port

Syntax

port *value*

no port

Context

[Tree] (config>service>vprn>log>syslog port)

Full Context

configure service vprn log syslog port

Description

This command configures the UDP port that will be used to send syslog messages to the syslog target host.

The port configuration is needed if the syslog target host uses a port other than the standard UDP syslog port 514.

Only one port can be configured. If multiple **port** commands are entered, the last entered port overwrites the previously entered ports.

The **no** form of this command reverts to default value.

Default

no port

Parameters

value

The value is the configured UDP port number used when sending syslog messages.

Values 1 to 65535

Platforms

7705 SAR Gen 2

port

Syntax

port {*port-id* | **lag** *lag-id*} [**egress**] [**ingress**]

no port {*port-id* | **lag** *lag-id*} {[**egress**] [**ingress**]}

Context

[Tree] (config>mirror>mirror-source port)

Full Context

configure mirror mirror-source port

Description

This command enables mirroring of traffic ingressing or egressing a port (Ethernet port, SONET/SDH channel, TDM channel, or Link Aggregation Group (LAG)).

The **port** command associates a port or LAG to a mirror source. The port is identified by the *port-id*. The defined port may be Ethernet, Access or network, SONET/SDH, or TDM channel access. A network port may be a single port or a Link Aggregation Group (LAG) ID. When a LAG ID is given as the *port-id*, mirroring is enabled on all ports making up the LAG. If the port is a SONET/SDH interface, the *channel-id* must be specified to identify which channel is being mirrored. Either a LAG port member *or* the LAG port can be mirrored.

The port is only referenced in the mirror source for mirroring purposes. The mirror source association does not need to be removed before deleting the card to which the port belongs. If the port is removed from the system, the mirroring association will be removed from the mirror source.

The same port may not be associated with multiple mirror source definitions with the **ingress** parameter defined. The same port may not be associated with multiple mirror source definitions with the **egress** parameter defined.

If a SAP is mirrored on an access port, the SAP mirroring will have precedence over the access port mirroring when a packet matches the SAP mirroring criteria. Filter and label mirroring destinations will also precedence over a port-mirroring destination.

If the port is not associated with a **mirror-source**, packets on that port will not be mirrored. Mirroring may still be defined for a SAP, label or filter entry, which will mirror based on a more specific criteria.

The encapsulation type on an access port or channel cannot be changed to Frame Relay if it is being mirrored.

The **no port** command disables port mirroring for the specified port. Mirroring of packets on the port may continue due to more specific mirror criteria. If the **egress** or **ingress** parameter keywords are specified in the **no** command, only the ingress or egress mirroring condition will be removed.

Parameters

port-id

Specifies the port ID.

port-id	*slot*/*mda*/*port* [.*channel*]		
	eth-sat-id	esat-*id*/*slot*/*port*	
		esat	keyword
		id	1 to 20
	pxc-id	pxc-*id*.*sub-port*	
		pxc	keyword
		id	1 to 64
		sub-port	a, b

bgrp-id	bpgrp-*type-bpgrp-num*	
	bgrp	keyword
	type	ima, ppp
	bgrp-num	1 to 2000
ccag-id	ccag-*id.path-id cc-type*:*cc-id*	
	ccag	keyword
	id	1 to 8
	path-id	a, b
	cc-type	sap-net, .net-sap
	cc-id	0 to 4094

lag-id

The LAG identifier, expressed as a decimal integer.

Values 1 to 800

egress

Specifies that packets egressing the port should be mirrored. Egress packets are mirrored to the mirror destination after egress packet modification.

ingress

Specifies that packets ingressing the port should be mirrored. Ingress packets are mirrored to the mirror destination prior to ingress packet modification.

Platforms

7705 SAR Gen 2

port

Syntax

port {*port-id* | **lag** *lag-id*} {[**egress**] [**ingress**]}

no port {*port-id* | **lag** *lag-id*} [**egress**] [**ingress**]

Context

[Tree] (debug>mirror-source port)

Full Context

debug mirror-source port

Description

This command enables mirroring of traffic ingressing or egressing a port (Ethernet port, SONET/SDH channel, TDM channel, or Link Aggregation Group (LAG)).

The **port** command associates a port or LAG to a mirror source. The port is identified by the *port-id*. The defined port may be Ethernet, Access or network, SONET/SDH, or TDM channel access. A network port may be a single port or a Link Aggregation Group (LAG) ID. When a LAG ID is given as the *port-id*, mirroring is enabled on all ports making up the LAG. If the port is a SONET/SDH interface, the *channel-id* must be specified to identify which channel is being mirrored. Either a LAG port member *or* the LAG port can be mirrored.

The port is only referenced in the mirror source for mirroring purposes. The mirror source association does not need to be removed before deleting the card to which the port belongs. If the port is removed from the system, the mirroring association will be removed from the mirror source.

The same port may not be associated with multiple mirror source definitions with the **ingress** parameter defined. The same port may not be associated with multiple mirror source definitions with the **egress** parameter defined.

If a SAP is mirrored on an access port, the SAP mirroring will have precedence over the access port mirroring when a packet matches the SAP mirroring criteria. Filter and label mirroring destinations will also precedence over a port-mirroring destination.

If the port is not associated with a **mirror-source**, packets on that port will not be mirrored. Mirroring may still be defined for a SAP, label or filter entry, which will mirror based on a more specific criteria.

The encapsulation type on an access port or channel cannot be changed to Frame Relay if it is being mirrored.

The **no port** command disables port mirroring for the specified port. Mirroring of packets on the port may continue due to more specific mirror criteria. If the **egress** or **ingress** parameter keywords are specified in the **no** command, only the ingress or egress mirroring condition will be removed.

Parameters

port-id

Specifies the port ID.

port-id	*slot*/*mda*/*port* [.*channel*]		
	eth-sat-id	esat-*id*/*slot*/*port*	
		esat	keyword
		id	1 to 20
	pxc-id	pxc-*id*.*sub-port*	
		pxc	keyword
		id	1 to 64
		sub-port	a, b
	ccag-id	ccag-*id*.*path-id cc-type*:*cc-id*	
		ccag	keyword
		id	1 to 8

path-id	a,b
cc-type	sap-net, net-sap
cc-id	0 to 4094

lag-id

Specifies the LAG identifier, expressed as a decimal integer.

Values 1 to 800

egress

Specifies that packets egressing the port should be mirrored. Egress packets are mirrored to the mirror destination after egress packet modification.

ingress

Specifies that packets ingressing the port should be mirrored. Ingress packets are mirrored to the mirror destination prior to ingress packet modification.

Platforms

7705 SAR Gen 2

port

Syntax

port {**lt** | **gt** | **eq**} *port-number*

port port-list *port-list-name*

port range *port-number port-number*

no port

Context

[Tree] (config>filter>ipv6-filter>entry>match port)

[Tree] (config>filter>ipv6-exception>entry>match port)

[Tree] (config>filter>ip-filter>entry>match port)

Full Context

configure filter ipv6-filter entry match port

configure filter ipv6-exception entry match port

configure filter ip-filter entry match port

Description

This command configures a TCP/UDP/SCTP source or destination port match criterion in IPv4 and IPv6 CPM (SCTP not supported) and/or ACL filter policies. A packet matches this criterion if the packet TCP/UDP/SCTP (as configured by protocol/next-header match) source OR destination port matches either the specified port value or a port in the specified port range or port-list.

Operational Note: This command is mutually exclusive with src-port and dst-port commands. Configuring "port eq 0", may match non-initial fragments where the source/destination port values are not present in a packet fragment if other match criteria are also met.

The **no** form of this command deletes the specified port match criterion.

Default

no port

Parameters

lt | gt | eq

Specifies the operator to use relative to *port-number* for specifying the port number match criteria.

lt

Specifies that all port numbers less than *port-number* match.

gt

Specifies that all port numbers greater than *port-number* match.

eq

Specifies that the *port-number* must be an exact match.

port-number

Specifies a source or destination port to be used as a match criterion. The port number can be expressed as a decimal integer, as well as in hexadecimal or binary format. The following value shows a decimal integer only.

Values 0 to 65535

port-list *port-list-name*

Specifies an inclusive range of source or destination port values to be used as match criteria.

range *port-number port-number*

Specifies an inclusive range of source or destination port values to be used as match criteria.

Platforms

7705 SAR Gen 2

port

Syntax

[**no**] **port** *port-number*

[**no**] **port range** *start end*

Context

[Tree] (config>filter>match-list>port-list port)

Full Context

configure filter match-list port-list port

Description

This command adds a port or a range of ports to an existing port match list. The **no** form of this command deletes the specified port or range of ports form the list.

Parameters

port-number

Specifies the port number to add to the list. The port number can be expressed as a decimal integer, as well as in hexadecimal or binary format. Below shows decimal integer only.

Values 0 to 65535

start end

Specifies an inclusive port range between two port numbers values. The *start* of the range and *end* of the range can be expressed as decimal integers, as well as in hexadecimal or binary format. The following value shows decimal integer only.

Values 0 to 65535

Platforms

7705 SAR Gen 2

port

Syntax

port *port-id*

no port

Context

[Tree] (config>router>origin-validation>rpki-session port)

Full Context

configure router origin-validation rpki-session port

Description

This command configures the destination port number to use when contacting the cache server. The default port number is 323. The port cannot be changed without first shutting down the session.

Default

no port

Parameters

port-id

Specifies a port ID.

Values 0 to 65535

Platforms

7705 SAR Gen 2

port

Syntax

port *port-name*

no port

Context

[Tree] (config>router>if port)

Full Context

configure router interface port

Description

This command creates an association with a logical IP interface and a physical port.

An interface can also be associated with the system (loopback address).

The command returns an error if the interface is already associated with another port or the system. In this case, the association must be deleted before the command is re-attempted. The *port-id or port-id for* Ethernet ports can be in one of the following forms:

Ethernet interfaces

If the card in the slot has MDAs/XMAs, *port-id* is in the *slot_number*/MDA or XMA *_number*/*port_number* format; for example, **1/1/3** specifies port 3 of the MDA/XMA installed in MDA/XMA slot 1 on the card installed in chassis slot 1.

SONET/SDH interfaces

When the *port-id* represents a POS interface, the *port-id* must include the *channel-id.* The POS interface must be configured as a **network** port.

The **no** form of this command deletes the association with the port. The **no** form of this command can only be performed when the interface is administratively down.

Default

no port

Parameters

port-name

The physical port identifier to associate with the IP interface.

Values

Table 74: Port Names

port-name	port-id[:encap-val]	
	encap-val	0 for null
		[0 to 4094] for dot1q
		[0 to 4094].* [1 to 4094].[0to 4094] for qinq
port-id	slot/mda/port[.channel]	
	aps-id	aps-<group-id>[.channel]
	aps	keyword
	group-id	1 to 128
	ccag-id	ccag-<id>.<path-id>[cc-type]
	ccag	keyword
	id	1 to 8
	path-id	a, b
	cc-type	[.sap-net| .net-sap]
	eth-tunnel-id	eth-tunnel-<id>
	eth-tunnel	keyword
	id	1 to 1024
	lag-id	lag-<id>
	lag	keyword
	id	1 to 800
	id	1 to 1024
	eth-sat-id	esat-<id>/<slot>/[u]<port>
	esat	keyword
	id	1 to 20

	u	keyword for up-link port

Platforms

7705 SAR Gen 2

port

Syntax

port *value*

no port

Context

[Tree] (config>log>syslog port)

Full Context

configure log syslog port

Description

This command configures the UDP port that will be used to send syslog messages to the syslog target host.

The port configuration is needed if the syslog target host uses a port other than the standard UDP syslog port 514.

Only one port can be configured. If multiple **port** commands are entered, the last entered port overwrites the previously entered ports.

The **no** form of this command removes the value from the configuration.

Parameters

value

Specifies the value that is the configured UDP port number used when sending syslog messages.

Values 1 **to** 65535

Platforms

7705 SAR Gen 2

port

Syntax

port *port*

no port

Context

[Tree] (config>system>netconf>listen port)

Full Context

configure system netconf listen port

Description

This command specifies the port on which the SR OS NETCONF server listens for new connections. Only one port can be configured for NETCONF management.

The configured port applies to both non-VPRN and VPRN management. New NETCONF connections are able to use the configured port. The SR OS NETCONF server errors if a port, different from the configured port, is used to SSH to the SR OS NETCONF server. For NETCONF connections not using VPRN management, active NETCONF connections are not disconnected if the port used to establish the connections is changed. For NETCONF connections using VPRN management, active NETCONF connections are disconnected if the port used to establish the connections is changed.

The **no** form of this command resets the port on which the SR OS NETCONF server listens to the default port of 830.

Parameters

port

Specifies the port on which NETCONF listens for new connections.

Values 22, 830

Default 830

Platforms

7705 SAR Gen 2

port

Syntax

port *port*

no port

Context

[Tree] (config>system>security>radius port)

Full Context

configure system security radius port

Description

This command configures the TCP port number to contact the RADIUS server.

The **no** form of this command reverts to the default value.

Default

port 1812 (as specified in RFC 2865, *Remote Authentication Dial In User Service* (*RADIUS*))

Parameters

port

Specifies the TCP port number to contact the RADIUS server.

Values 1 to 65535

Platforms

7705 SAR Gen 2

port

Syntax

port *port*

no port

Context

[Tree] (config>system>grpc-tunnel>tunnel>handler port)

Full Context

configure system grpc-tunnel tunnel handler port

Description

This command assigns the TCP port number that the handler listens to internally.

The **no** form of this command disables the handler from listening to a TCP port.

Default

no port

Parameters

port

Specifies the TCP port number.

Values 1 to 65535

Platforms

7705 SAR Gen 2

22.26 port-control

port-control

Syntax

port-control [**auto** | **force-auth** | **force-unauth**]

Context

[Tree] (config>port>ethernet>dot1x port-control)

Full Context

configure port ethernet dot1x port-control

Description

This command configures the 802.1x authentication mode.

The **no** form of this command returns the value to the default.

Default

port-control force-auth

Parameters

force-auth

Disables 802.1x authentication and causes the port to transition to the authorized state without requiring any authentication exchange. The port transmits and receives normal traffic without requiring 802.1x-based host authentication.

force-unauth

Causes the port to remain in the unauthorized state, ignoring all attempts by the hosts to authenticate. The switch cannot provide authentication services to the host through the interface.

auto

Enables 802.1x authentication. The port starts in the unauthorized state, allowing only EAPoL frames to be sent and received through the port. Both the router and the host can initiate an authentication procedure. The port will remain in unauthorized state (no traffic except EAPoL frames is allowed) until the first client is authenticated successfully. After this, traffic is allowed on the port for all connected hosts.

Platforms

7705 SAR Gen 2

22.27 port-down

port-down

Syntax

[**no**] **port-down** *port-id*

Context

[Tree] (config>vrrp>policy>priority-event port-down)

Full Context

configure vrrp policy priority-event port-down

Description

This command configures a port down priority control event that monitors the operational state of a port or SONET/SDH channel. When the port or channel enters the operational down state, the event is considered set. When the port or channel enters the operational up state, the event is considered cleared.

Multiple unique **port-down** event nodes can be configured within the **priority-event** context up to the overall limit of 32 events. Up to 32 events can be defined in any combination of types.

The **port-down** command can reference an arbitrary port or channel. The port or channel does not need to be preprovisioned or populated within the system. The operational state of the **port-down** event is set as follows:

- Set – non-provisioned
- Set – not populated
- Set – down
- Cleared – up

When the port or channel is provisioned, populated, or enters the operationally up or down state, the event operational state is updated appropriately.

When the event enters the operationally down, non-provisioned, or non-populated state, the event is considered to be set. When an event transitions from clear to set, the set is processed immediately and must be reflected in the associated virtual router instances in-use priority value. As the event transitions from cleared to set, a hold-set timer is loaded with the value configured by the events **hold-set** command. This timer prevents the event from clearing until it expires, damping the effect of event flapping. If the event clears and becomes set again before the hold-set timer expires, the timer is reset to the **hold-set** value, extending the time before another clear can take effect.

When the event enters the operationally up state, the event is considered to be cleared. Once the events **hold-set** expires, the effects of the events **priority** value are immediately removed from the in-use priority of all associated virtual router instances.

The actual effect on the virtual router instance in-use priority value depends on the defined event priority and its delta or explicit nature.

The **no** form of the command deletes the specific port or channel monitoring event. The event may be removed at anytime. When the event is removed, the in-use priority of all associated virtual router instances will be re-evaluated. The events **hold-set** timer has no effect on the removal procedure.

Default

no port-down — No port down priority control events are defined.

Parameters

port-id

The port ID of the port monitored by the VRRP priority control event.

The *port-id* can only be monitored by a single event in this policy. The port can be monitored by multiple VRRP priority control policies. A port and a specific channel on the port are considered to be separate entities. A port and a channel on the port can be monitored by separate events in the same policy.

Values

slot/*mda*/*port*[.*channel*]		
eth-sat-id	esat-*id*/*slot*/*port*	
	esat	keyword
	id	1 to 20
pxc-id	pxc-*id*.*sub-port*	
	pxc	keyword
	id	1 to 64
	sub-port	a, b
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 to 64
ccag-id	ccag-*id*. *path-id*[*cc-type*]	
	ccag	keyword
	id	1 to 8
	path-id	a, b
	cc-type	.sap-net, .net-sap

Platforms

7705 SAR Gen 2

22.28 port-forwarding

port-forwarding

Syntax

port-forwarding

Context

[Tree] (config>service>nat port-forwarding)

Full Context

configure service nat port-forwarding

Description

Commands in this context configure NAT port forwarding parameters.

Platforms

7705 SAR Gen 2

22.29 port-forwarding-dyn-block-reservation

port-forwarding-dyn-block-reservation

Syntax

[**no**] **port-forwarding-dyn-block-reservation**

Context

[Tree] (config>router>nat>outside>pool port-forwarding-dyn-block-reservation)

[Tree] (config>service>vprn>nat>outside>pool port-forwarding-dyn-block-reservation)

Full Context

configure router nat outside pool port-forwarding-dyn-block-reservation

configure service vprn nat outside pool port-forwarding-dyn-block-reservation

Description

This command will enable the reservation of the dynamic port blocks when the first port forward for the subscriber is created. The dynamic port bloc allocation is logged only if the block is being utilized (mapping are created). In other words, dynamic port block reservation due to the port forward creation but without any dynamic mapping, will not be logged.

The reserved port block will be released only when the last mapping in the block expires and there is not port forward associated with the subscriber. The de-allocation log (syslog or Radius) will be generated when the dynamic port block is completely released.

Dynamic port block reservation can be enabled only if the configured maximum number of subscriber per outside IP address is less or equal then the maximum number of configured port blocks per outside IP address.

Default

no port-forwarding-dyn-block-reservation

Platforms

7705 SAR Gen 2

22.30 port-forwarding-range

port-forwarding-range

Syntax

port-forwarding-range [*range-start*] *range-end*

no port-forwarding-range

Context

[Tree] (config>router>nat>outside>pool port-forwarding-range)

[Tree] (config>service>vprn>nat>outside>pool port-forwarding-range)

Full Context

configure router nat outside pool port-forwarding-range

configure service vprn nat outside pool port-forwarding-range

Description

This command configures the lower and upper limit for port forwards in the ephemeral port space (wildcard port space) of all IP addresses in a NAT pool. A well-known port range (ports 1 to 1023) is always enabled for port forwards, and it cannot be disabled for pools in NAPT mode.

Pools in 1:1 mode do not support configured port forwards. These pools do not perform port translation and they automatically forward traffic initiated on the outside toward the inside.

Port 0 is always excluded from the port forwarding range.

The upper bound of the wildcard port range is reserved for port forwards. If the value for the *range-start* is not provided, the wildcard port range implicitly starts at 1024.

range-start 0 cannot be configured by an operator because it is reserved for 1:1 pools that do not support configured port forwards.

If you configure *port-forwarding-range 3000*, configures ports 1 to 3000 as port forwards. This implies that the well-known ports and wildcard ports are contiguous. If you configure *port-forwarding-range 2000 3000*, the router implicitly includes ports 1 to 1023, plus enables the wildcard port range 2000 to 3000, which is now disjoined from the well-known ports.

The *range-start* parameter has additional values that are configurable in the CLI. 0 is reserved for pools that do not support configured port forwards (those are 1:1 pools).

range-start 1 means that well-known ports and wildcard port forwards are contiguous. This is configured by omitting the *range-start* parameter and only configuring the *range-end* parameter.

The **no** form of this command disables the port forwards capability in the wildcard port range of all IP addresses in a NAT pool.

Default ranges in the *range-start* and *range-end* parameters in the MIB for the NAT pools that support port forwarding ranges are set to include only well-known ports, *range-start 1* and *range-end 1023*.

Parameters

range-start

Specifies the lower boundary of the wildcard port range reserved for port forwards. When configured, the value must be less than the *range-end* value.

Values 0, 1, 1025 to 65535

Default 1

range-end

Specifies the upper boundary of the wildcard port range reserved for port forwards.

Values 0, 1023 to 65535

Default 1023

Platforms

7705 SAR Gen 2

22.31 port-id

port-id

Syntax

[**no**] **port-id**

Context

[Tree] (config>router>if>dhcp>option>vendor-specific-option port-id)

Full Context

configure router interface dhcp option vendor-specific-option port-id

Description

This command enables sending of the port-id in the Nokia vendor specific suboption of the DHCP relay packet

The **no** form of this command disables the sending.

Default

no port-id

Platforms

7705 SAR Gen 2

22.32 port-id-subtype

port-id-subtype

Syntax

port-id-subtype {**tx-if-alias** | **tx-if-name** | **tx-local**}

Context

[Tree] (config>port>ethernet>lldp>dstmac port-id-subtype)

Full Context

configure port ethernet lldp dest-mac port-id-subtype

Description

This command specifies how to encode the PortID TLV transmit to the peer. The default setting **tx-local** (ifindex value) is required by some versions of the NSP NSM-P to properly build the Layer 2 topology map using LLDP. Changing this value to transmit the ifname (**tx-if-name**) or ifAlias (**tx-if-alias**) in place of the ifindex (**tx-local**) may affect the ability of the NSP NFM-P to build the Layer 2 topology map using LLDP.

Default

port-id-subtype tx-local

Parameters

tx-if-alias

Transmits the ifAlias String (subtype 1) that describes the port as stored in the IF-MIB, either user configured or the default entry (i.e. 10/100/Gig Ethernet SFP).

tx-if-name

Transmits the ifName string (subtype 5) that describes the port as stored in the IF-MIB ifName info.

tx-local

The interface ifIndex value (subtype 7) as the PortID.

Platforms

7705 SAR Gen 2

port-id-subtype

Syntax

port-id-subtype {**tx-if-alias** | **tx-if-name** | **tx-local**}

Context

[Tree] (config>lag>lldp-member-template>dstmac port-id-subtype)

Full Context

configure lag lldp-member-template dest-mac port-id-subtype

Description

This command configures the encoding of the PortID TLV that is transmitted to the peer. Some versions of the NSP NFM-P require the default setting **tx-local** (ifIndex value) to properly build the Layer 2 topology map using LLDP. Changing this value to transmit the ifName (**tx-if-name**) or ifAlias (**tx-if-alias**) in place of the ifIndex (**tx-local**) may affect the ability of the NSP NFM-P to build the Layer 2 topology map using LLDP.

Default

port-id-subtype tx-local

Parameters

tx-if-alias

Keyword to transmit the ifAlias String (subtype 1), which describes the port as stored in the IF-MIB, either user configured or the default entry (for example, 10/100/Gig Ethernet SFP).

tx-if-name

Keyword to transmit the ifName string (subtype 5), which describes the port as stored in the IF-MIB ifName information.

tx-local

Keyword to transmit the interface ifIndex value (subtype 7) as the port ID.

Platforms

7705 SAR Gen 2

22.33 port-limits

port-limits

Syntax

port-limits

Context

[Tree] (config>service>nat>nat-policy port-limits)

Full Context

configure service nat nat-policy port-limits

Description

This command configures the port limits of this policy.

Platforms

7705 SAR Gen 2

22.34 port-list

port-list

Syntax

port-list *port-list-name* [**create**]

no port-list *port-list-name*

Context

[Tree] (config>filter>match-list port-list)

Full Context

configure filter match-list port-list

Description

This command creates a list of TCP/UDP/SCTP port values or ranges for match criteria in IPv4 and IPv6 ACL and CPM filter policies.

The **no** form of this command deletes the specified list.

Operational notes:

SCTP port match is supported in ACL filter policies only.

A port-list must contain only TCP/UDP/SCTP port values or ranges.

A TCP/UDP/SCTP port match list cannot be deleted if it is referenced by a filter policy.

See general description related to match-list usage in filter policies.

Parameters

port-list-name

Specifies a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

22.35 port-num

port-num

Syntax

port-num *virtual-port-number*

no port-num [*virtual-port-number*]

Context

[Tree] (config>service>vpls>sap>stp port-num)

[Tree] (config>service>vpls>spoke-sdp>stp port-num)

Full Context

configure service vpls sap stp port-num

configure service vpls spoke-sdp stp port-num

Description

This command configures the virtual port number which uniquely identifies a SAP within configuration bridge protocol data units (BPDUs). The internal representation of a SAP is unique to a system and has a reference space much bigger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a SAP and identifies it with its own virtual port number that is unique to every other SAP defined on the TLS. The virtual port number is assigned at the time that the SAP is added to the TLS. Since the order that the SAP was added to the TLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance.

The virtual port number cannot be administratively modified.

Platforms

7705 SAR Gen 2

22.36 port-parent

port-parent

Syntax

port-parent [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]

no port-parent

Context

[Tree] (config>qos>sap-egress>queue port-parent)

Full Context

configure qos sap-egress queue port-parent

Description

This command specifies whether this queue feeds off a port-level scheduler. When configured, this SAP egress queue is parented by a port-level scheduler. This object is mutually exclusive with SAP egress queue parent. Only one kind of parent is allowed.

The **port-parent** command defines a child/parent association between an egress queue and a port-based scheduler or between an intermediate service scheduler and a port-based scheduler. The **port-parent** command allows for a set of within-CIR and above-CIR parameters that define the port priority levels and weights for the queue or scheduler. If the **port-parent** command is executed without any parameters, the default parameters are assumed.

In this context, the **port-parent** command is mutually exclusive to the **parent** command (used to create a parent/child association between a queue and an intermediate scheduler). Executing a **port-parent** command when a parent definition is in place causes the current intermediate scheduler association to be removed and replaced by the defined port-parent association. Executing a **parent** command when a port-parent definition exists causes the port scheduler association to be removed and replaced by the defined intermediate scheduler name.

Changing the parent context on a SAP egress policy queue may cause a SAP or subscriber or multiservice site context of the queue (policy associated with a SAP or subscriber profile or multiservice site) to enter an orphaned state. If an instance of a queue is created on a port or channel that does not have a port scheduler enabled and the sap-egress policy creating the queue has a port-parent association, the queue will be allowed to run according to its own rate parameters and will not be controlled by a virtual scheduling context. If an instance of a queue is on a port or channel that has a port scheduler configured and the sap-egress policy defines the queue as having a non-existent intermediate scheduler parent, the queue will be treated as an orphan and will be handled according to the current orphan behavior on the port scheduler.

The **no** form of this command removes a port scheduler parent association for the queue or scheduler. If a port scheduler is defined on the port on which the queue or scheduler instance exists, the queue or scheduler will become orphaned if an port scheduler is configured on the egress port of the queue or scheduler.

Default

no port-parent

Parameters

weight *weight*

Specifies the weight the queue or scheduler will use at the above-CIR port priority level (defined by the level parameter).

Values 0 to 100

Default 1

level *level*

Specifies the port priority the queue or scheduler will use to receive bandwidth for its above-CIR offered-load.

Values 1 to 8 (8 is the highest priority)

Default 1

cir-weight *cir-weight*

Specifies the weight the queue or scheduler will use at the within-CIR port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-CIR pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 100

Default 0

cir-level *cir-level*

Specifies the port priority the queue or scheduler will use to receive bandwidth for its within-CIR offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-CIR pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 8 (8 is the highest priority)

Default 0

Platforms

7705 SAR Gen 2

port-parent

Syntax

port-parent [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]

no port-parent

Context

[Tree] (config>qos>network-queue>queue port-parent)

Full Context

configure qos network-queue queue port-parent

Description

This command specifies whether this queue feeds off a port-level scheduler. For the network-queue policy context, only the port-parent command is supported. When a port scheduler exists on the port, network queues without a port-parent association will be treated as an orphan queue on the port scheduler and treated according to the current orphan behavior on the port scheduler. If the port-parent command is defined for a network queue on a port without a port scheduler defined, the network queue will operate as if a parent association does not exist. When a port scheduler policy is associated with the egress port, the port-parent command will come into effect.

When a network-queue policy is associated with an FP for ingress queue definition, the port-parent association of the queues is ignored.

The **no** form of this command removes a port scheduler parent association for the queue or scheduler. If a port scheduler is defined on the port then the queue or scheduler instance exists, the queue or scheduler will become orphaned.

Default

no port-parent

Parameters

weight *weight*

Specifies the weight the queue or scheduler will use at the above-CIR port priority level (defined by the level parameter).

Values 0 to 100

Default 1

level *level*

Specifies the port priority the queue or scheduler will use to receive bandwidth for its above-CIR offered-load.

Values 1 to 8 (8 is the highest priority)

Default 1

cir-weight ***cir-weight***

Specifies the weight the queue or scheduler will use at the within-CIR port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0, the queue or scheduler does not receive bandwidth during the port scheduler's within-CIR pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter is used.

Values 0 to 100

Default 0

cir-level ***cir-level***

Specifies the port priority the queue or scheduler will use to receive bandwidth for its within-CIR offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port scheduler's within-CIR pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 8 (8 is the highest priority)

Default 0

Platforms

7705 SAR Gen 2

port-parent

Syntax

port-parent [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]

no port-parent

Context

[Tree] (config>qos>qgrps>egr>qgrp>queue port-parent)

Full Context

configure qos queue-group-templates egress queue-group queue port-parent

Description

This command defines the port scheduling parameters used to control the queue's behavior when a virtual egress port scheduling is enabled where the egress queue group template is applied. The **port-parent** command follows the same behavior and provisioning characteristics as the **parent** command in the SAP egress QoS policy. The **port-parent** command and the **parent** command are mutually exclusive.

The **no** form of this command removes the values from the configuration.

Parameters

weight *weight*

Specifies the weight the queue or scheduler will use at the above-CIR port priority level (defined by the level parameter).

Values 0 to 100

Default 1

level *level*

Specifies the port priority the queue or scheduler will use to receive bandwidth for its above-CIR offered-load.

Values 1 to 8 (8 is the highest priority)

Default 1

cir-weight *cir-weight*

Specifies the weight the queue or scheduler will use at the within-CIR port priority level (defined by the **cir-level** parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the **cir-weight** parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-CIR pass and the **cir-level** parameter is ignored. If the **cir-weight** parameter is 1 or greater, the **cir-level** parameter is used.

Values 0 to 100

Default 0

cir-level *cir-level*

Specifies the port priority the queue or scheduler will use to receive bandwidth for its within-CIR offered-load. If the **cir-weight** parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-CIR pass and the **cir-level** parameter is ignored. If the **cir-weight** parameter is 1 or greater, the **cir-level** parameter is used.

Values 0 to 8 (8 is the highest priority)

Default 0

Platforms

7705 SAR Gen 2

port-parent

Syntax

port-parent [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]

no port-parent

Context

[Tree] (config>qos>scheduler-policy>tier>scheduler port-parent)

Full Context

configure qos scheduler-policy tier scheduler port-parent

Description

The **port-parent** command defines a child/parent association between an egress scheduler and a port-based scheduler, or between an intermediate service scheduler and a port-based scheduler. The **port-parent** command allows for a set of within-CIR and above-CIR parameters that define the port priority levels and weights for the scheduler. If the **port-parent** command is executed without any parameters, the default parameters are assumed.

In this context, the **port-parent** command and the **parent** command (used to create a parent/child association to an intermediate scheduler) are mutually exclusive. Executing a **port-parent** command when a parent definition is in place causes the current intermediate scheduler association to be removed and replaced by the defined port-parent association. Executing a **parent** command when a port-parent definition exists causes the port scheduler association to be removed and replaced by the defined intermediate scheduler name.

Changing the parent context on a SAP egress policy policer or queue may cause a SAP or subscriber context of the policer or queue (policy associated with a SAP or subscriber profile) to enter an orphaned state. If an instance of a policer or queue is created on a port or channel that does not have a port scheduler enabled and the sap-egress policy creating the policer queue has a port-parent association, the policer or queue will be allowed to run according to its own rate parameters and will not be controlled by a virtual scheduling context. If an instance of a policer or queue is on a port or channel that has a port scheduler configured and the sap-egress policy defines the policer or queue as having a non-existent intermediate scheduler parent, the policer or queue will be treated as an orphan and will be handled according to the current orphan behavior on the port scheduler.

The **no** form of this command removes a port scheduler parent association for the scheduler. If a port scheduler is defined on the port that the scheduler instance exists, the scheduler will become orphaned if an port scheduler is configured on the egress port of the queue or scheduler.

Default

no port-parent

Parameters

weight *weight*

Specifies the weight the queue or scheduler will use at the above-CIR port priority level (defined by the **level** parameter).

Values 0 to 100

Default 1

level *level*

Specifies the port priority the queue or scheduler will use to receive bandwidth for its above-CIR offered-load.

Values 1 to 8 (8 is the highest priority)

Default 1

cir-weight ***cir-weight***

Specifies the weight the queue or scheduler will use at the within-CIR port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the **cir-weight** parameter is set to a value of 0, the queue or scheduler does not receive bandwidth during the port scheduler's within-CIR pass and the **cir-level** parameter is ignored. If the **cir-weight** parameter is 1 or greater, the **cir-level** parameter comes into play.

Values 0 to 100

Default 0

cir-level ***cir-level***

Specifies the port priority the queue or scheduler will use to receive bandwidth for its within-CIR offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port scheduler's within-CIR pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

Values 0 to 8 (8 is the highest priority)

Default 0

Platforms

7705 SAR Gen 2

22.37 port-redirect-group

port-redirect-group

Syntax

port-redirect-group {**queue** *queue-id* | **policer** *policer-id* [**queue** *queue-id*]}

no port-redirect-group

Context

[Tree] (config>qos>network>egress>fc port-redirect-group)

Full Context

configure qos network egress fc port-redirect-group

Description

This command is used to redirect the FC of a packet of a pseudowire (PW) or network IP interface to an egress port queue group.

It defines the mapping of an FC to a queue ID or a policer ID and a queue ID and redirects the lookup of the queue or policer of the same ID in some egress port queue-group instance. However, the queue-group name and instance are explicitly provided only at the time the network QoS policy is applied to egress context of a spoke-sdp or a network IP interface.

The **no** version of this command removes the redirection of the FC.

Parameters

queue-id

This parameter must be specified when executing the **port-redirect-group** command. The specified *queue-id* must exist within the egress port queue group on each IP interface where the network QoS policy is applied.

Values 1 to 8

policer id

The specified policer-id must exist within the queue-group template applied to the ingress context of the forwarding plane.

Values 1 to 8

Platforms

7705 SAR Gen 2

22.38 port-reservation

port-reservation

Syntax

port-reservation blocks *num-blocks*

port-reservation ports *num-ports*

no port-reservation

Context

[Tree] (config>router>nat>outside>pool port-reservation)

[Tree] (config>service>vprn>nat>outside>pool port-reservation)

Full Context

configure router nat outside pool port-reservation

configure service vprn nat outside pool port-reservation

Description

This command configures the size of the port block that will be assigned to a host that is served by this pool. The number of ports configured are available to UDP, TCP, and ICMP (as identifiers).

Parameters

num-blocks

Specifies the number of port blocks per IP address. Setting this parameter to one (1) for large scale NAT enables 1:1 NAT for IP addresses in this pool.

Values 1 to 64512

num-ports

Specifies the number of ports per block.

Values 0 to 64512 (for deterministic pools)

1 to 64512 (for non-deterministic pools)

Platforms

7705 SAR Gen 2

22.39 port-role

port-role

Syntax

[**no**] **port-role**

Context

[Tree] (debug>service>id>stp port-role)

Full Context

debug service id stp port-role

Description

This command enables STP debugging for changes in port roles.

Platforms

7705 SAR Gen 2

22.40 port-scheduler-policy

port-scheduler-policy

Syntax

port-scheduler-policy *port-scheduler-name* [**create**]

no port-scheduler-policy *port-scheduler-name*

Context

[Tree] (config>qos port-scheduler-policy)

Full Context

configure qos port-scheduler-policy

Description

When a port scheduler has been associated with an egress port, it is possible to override the following parameters:

- The max-rate allowed for the scheduler
- The maximum rate for each priority level (1 to 8)
- The cir associated with each priority level (1 to 8)

The orphan priority level (level 0) has no configuration parameters and cannot be overridden.

The **no** form of this command removes a port scheduler policy from the system. If the port scheduler policy is associated with an egress port or channel, the command will fail.

Parameters

port-scheduler-name

Specifies an existing port scheduler name. Each port scheduler must be uniquely named within the system and can be up to 32 ASCII characters.

Platforms

7705 SAR Gen 2

22.41 port-state

port-state

Syntax

[**no**] **port-state**

Context

[Tree] (debug>service>id>stp port-state)

Full Context

debug service id stp port-state

Description

This command enables STP debugging for port states.

The **no** form of the command disables debugging.

Platforms

7705 SAR Gen 2

22.42 port-threshold

port-threshold

Syntax

port-threshold *value* [**action** { **dynamic-cost** | **static-cost** | **down**}] [**cost** *static-cost*]

no port-threshold

Context

[Tree] (config>lag port-threshold)

Full Context

configure lag port-threshold

Description

This command configures the behavior for the Link Aggregation Group (LAG) if the number of operational links is equal to or below a threshold level.

Nokia recommends that operators use the **weight-threshold** or **hash-weight-threshold** command instead of the **port-threshold** command to control LAG operational status. For example, when 1GE and 10GE ports are mixed in a LAG, each 1GE port will have a weight of 1, while each 10GE port will have a weight of 10.

The **weight-threshold** or **hash-weight-threshold** command can also be used for LAGs with all ports of equal speed to allow a common operational model. For example, each port has a weight of 1 to mimic **port-threshold** and its related configuration.

The **no** form of this command reverts to the default values.

Default

port-threshold 0 action down

Parameters

value

Specifies the decimal integer threshold number of operational links for the LAG at or below which the configured action is invoked. If the number of operational links exceeds the **port-threshold** value, any action taken for being below the threshold value will cease.

Values 0 to 63

action

Specifies the action to take if the number of active links in the LAG is at or below the threshold value.

dynamic-cost

Specifies that dynamic costing is activated. As a result, the LAG remains operationally up with a cost relative to the number of operational links. The link is only regarded as operationally down when all links in the LAG are down.

static-cost

Specifies that static costing is activated. As a result, the LAG remains operationally up with the configured cost, regardless of the number of operational links. The link is only regarded as operationally down when all links in the LAG are down.

down

Specifies that LAG is brought operationally down if the number of operational links is equal to or less than the configured threshold value. The LAG is only regarded as up once the number of operational links exceeds the configured threshold value.

static-cost

Specifies decimal integer static cost of the LAG.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

22.43 port-xc

port-xc

Syntax

port-xc

Context

[Tree] (config port-xc)

Full Context

configure port-xc

Description

Commands in this context configure port-cross connect functionality.

Platforms

7705 SAR Gen 2

22.44 ppk

ppk

Syntax

ppk list *ppk-list-name* **id** *ppk-id*

Context

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>dyn ppk)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel>dyn ppk)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>dyn ppk)

[Tree] (config>ipsec>trans-mode-prof>dyn ppk)

[Tree] (config>router>if>ipsec>ipsec-tunnel>dyn ppk)

Full Context

configure service ies interface ipsec ipsec-tunnel dynamic-keying ppk

configure service vprn interface sap ipsec-tunnel dynamic-keying ppk

configure service vprn interface ipsec ipsec-tunnel dynamic-keying ppk

configure ipsec ipsec-transport-mode-profile dynamic-keying ppk

configure router interface ipsec ipsec-tunnel dynamic-keying ppk

Description

This command specifies the PPK to use for dynamic keying of the IPsec tunnel.

The **no** form of this command removes the PPK.

Default

no ppk

Parameters

ppk-list-name

Specifies the name of the PPK list, up to 32 characters.

ppk-id

Specifies the ID of a PPK entry in the list, up to 64 characters.

Platforms

7705 SAR Gen 2

22.45 ppk-id

ppk-id

Syntax

ppk-id *ppk-id* **value** *value-string* **format** {**ascii** | **hex**} [**hash** | **hash2** | **custom**]

no ppk-id *ppk-id*

Context

[Tree] (config>ipsec>ppk-list ppk-id)

Full Context

configure ipsec ppk-list ppk-id

Description

This command configures the attributes for a PPK entry within the list.

The **no** form of this command deletes the PPK entry from the list.

Parameters

ppk-id

Specifies a unique ID for the PPK, up to 64 characters.

value-string

Specifies the PPK value.

ascii

Keyword to specify that the PPK value is formatted as an ASCII string, up to 64 characters.

hex

Keyword that specifies the PPK value is formatted as a hexadecimal string, up to 128 hex nibbles.

Values 0x0 to 0xFFFFFFFF...

hash

Keyword that specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Keyword that specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Keyword that specifies the custom encryption to the management interface.

Platforms

7705 SAR Gen 2

22.46 ppk-list

ppk-list

Syntax

ppk-list *ppk-list-name* [**create**]

no ppk-list *ppk-list-name*

Context

[Tree] (config>ipsec ppk-list)

Full Context

configure ipsec ppk-list

Description

Commands in this context configure a list of Post-quantum Preshared Keys (PPKs) to use for IKEv2 authentication, as described in RFC 8784.

The **no** form of this command deletes the PPK list.

Parameters

ppk-list-name

Specifies the name of the PPK list, up to 32 characters.

create

Keyword to create the PPK list.

Platforms

7705 SAR Gen 2

ppk-list

Syntax

ppk-list *ppk-list-name*

no ppk-list

Context

[Tree] (config>ipsec>tnl-temp ppk-list)

Full Context

configure ipsec tunnel-template ppk-list

Description

This command specifies a PPK list to use in the tunnel template, which represents a list of PPKs available for the IPsec gateway. The actual PPK to use depends on the tunnel initiator.

The **no** form of this command removes the PPK list from the tunnel template.

Default

no ppk-list

Parameters

ppk-list-name

Specifies the name of the PPK list, up to 32 characters.

Platforms

7705 SAR Gen 2

22.47 ppk-required

ppk-required

Syntax

[**no**] **ppk-required**

Context

[Tree] (config>ipsec>ike-policy ppk-required)

Full Context

configure ipsec ike-policy ppk-required

Description

This command configures the mandatory use of PPKs for the IKEv2 key derivation process in the IKE policy.

The **no** form of this command configures the use of PPKs for IKEv2 as optional. The router can fall back to derive keys without PPK.

Default

no ppk-required

Platforms

7705 SAR Gen 2

22.48 pre-login-message

pre-login-message

Syntax

pre-login-message *login-text-string* [**name**]

no pre-login-message

Context

[Tree] (config>system>login-control pre-login-message)

Full Context

configure system login-control pre-login-message

Description

This command configures a message to display before logging in to the router using Telnet, SSH, or the console port.

Only one message can be configured. If a new pre-login message is configured, the new message overwrites the previous message.

Note: The pre-login message is displayed on both active and standby systems.

The **no** form of this command removes the pre-login message.

Default

no pre-login-message

Parameters

login-text-string

Specifies the pre-login message text, up to 900 characters. Any printable, 7-bit ASCII characters can be used. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Some special characters can be used to format the message text. Use the newline (\n) character to create multiline messages. A newline (\n) character in the message moves to the beginning of the next line by sending ASCII/UTF-8 characters 0xA (LF) and 0xD (CR) to the client terminal. A carriage return (\r) character in the message sends the ASCII/UTF-8 character 0xD (CR) to the client terminal.

name

Displays the configured system name before the pre-login message. To remove the system name from the pre-login message, remove the current message and configure a new message without using the **name** parameter.

Platforms

7705 SAR Gen 2

22.49 pre-shared-key

pre-shared-key

Syntax

pre-shared-key *pre-shared-key-index* [**encryption-type** *encryption-type*] [**create**]

no pre-shared-key *pre-shared-key-index*

Context

[Tree] (config>macsec>conn-assoc>static-cak pre-shared-key)

Full Context

configure macsec connectivity-association static-cak pre-shared-key

Description

This command specifies the pre-shared key used to enable MACsec using static connectivity association key (CAK) security mode. This command also specifies the encryption algorithm used for encrypting the SAK.

A pre-shared key includes a connectivity association key name (CKN) and a connectivity association key (CAK). The pre-shared key-the CKN and CAK-must match on both ends of a link.

A pre-shared key is configured on both devices at each end of point-to-point link to enable MACsec using static CAK security mode. The MACsec Key Agreement (MKA) protocol is enabled after the successful MKA liveliness negotiation.

The encryption-type is used for encrypting the SAK and authenticating the MKA packet. The symmetric encryption key SAK (Security Association Key) needs to be encrypted (wrapped) via the MKA protocols. The AES key is derived via pre-shared-key.

The **no** form of this command removes the index.

Parameters

pre-shared-key-index

Specifies the index of this pre-shared-key.

Values 1, 2

encryption-type

Specifies the type of encryption.

Values **aes-128-cmac**, **aes-256-cmac**

create

Mandatory to create an entry.

Platforms

7705 SAR Gen 2

pre-shared-key

Syntax

pre-shared-key *key* [**hash** | **hash2** | **custom**]

no pre-shared-key

Context

[Tree] (config>ipsec>client-db>client>credential pre-shared-key)

Full Context

configure ipsec client-db client credential pre-shared-key

Description

This command specifies a pre-shared key used to authenticate peers.

The **no** form of this command reverts to the default.

Default

no pre-shared-key

Parameters

key

An ASCII string to use as the pre-shared key for dynamic keying. When the **hash** or **hash2** parameters are not used, the key is a clear text key; otherwise, the key text is encrypted.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

pre-shared-key

Syntax

pre-shared-key *key* [**hash** | **hash2** | **custom**]

no pre-shared-key

Context

[Tree] (config>ipsec>trans-mode-prof>dyn pre-shared-key)

[Tree] (config>service>ies>if>sap>ipsec-gw pre-shared-key)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying pre-shared-key)

[Tree] (config>router>if>ipsec>ipsec-tunnel>dyn pre-shared-key)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>dyn pre-shared-key)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>dyn pre-shared-key)

[Tree] (config>service>vprn>if>sap>ipsec-gw pre-shared-key)

Full Context

configure ipsec ipsec-transport-mode-profile dynamic-keying pre-shared-key

configure service ies interface sap ipsec-gw pre-shared-key

configure service vprn interface sap ipsec-tunnel dynamic-keying pre-shared-key

configure router interface ipsec ipsec-tunnel dynamic-keying pre-shared-key

configure service vprn interface ipsec ipsec-tunnel dynamic-keying pre-shared-key

configure service ies interface ipsec ipsec-tunnel dynamic-keying pre-shared-key

configure service vprn interface sap ipsec-gw pre-shared-key

Description

This command configures the pre-shared key for authentication.

The **no** form of this command reverts to the default.

Default

no pre-shared-key

Parameters

key

Specifies an ASCII string to use as the pre-shared key for dynamic keying. When the **hash** or **hash2** parameters are not used, the key is a clear text key; otherwise, the key text is encrypted.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

22.50 pre-update-time

pre-update-time

Syntax

pre-update-time [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

Context

[Tree] (config>system>security>pki>ca-prof>auto-crl-update pre-update-time)

Full Context

configure system security pki ca-profile auto-crl-update pre-update-time

Description

This command specifies the pre-download time for next-update-based update.

Default

pre-update-time hrs 1

Parameters

days

Specifies the time period, in days, prior to the next update time of the current CRL.

Values 0 to 366

hours

Specifies the time period, in hours, prior to the next update time of the current CRL.

Values 0 to 23

minutes

Specifies the time period, in minutes, prior to the next update time of the current CRL.

Values 0 to 59

seconds

Specifies the time period, in seconds, prior to the next update time of the current CRL.

Values 0 to 59

Platforms

7705 SAR Gen 2

22.51 prec

prec

Syntax

prec *ip-prec-value* [**fc** *fc-name*] [**priority** {**high** | **low**}]

no prec *ip-prec-value*

Context

[Tree] (config>qos>sap-ingress prec)

Full Context

configure qos sap-ingress prec

Description

This command explicitly sets the forwarding class or enqueuing priority when a packet is marked with an IP precedence value (*ip-prec-value)*. Adding an IP precedence rule on the policy forces packets that match the specified *ip-prec-value* to override the forwarding class and enqueuing priority based on the parameters included in the IP precedence rule.

When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy.

When the enqueuing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueuing priority derived from earlier matches in the classification hierarchy.

The *ip-prec-value* is derived from the most significant three bits in the IP header ToS byte field (precedence bits). The three precedence bits define eight Class-of-Service (CoS) values commonly used to map packets to per-hop Quality of Service (QoS) behavior. The precedence bits are also part of the DiffServ Code Point (DSCP) method of mapping packets to QoS behavior. The DSCP uses the most significant six bits in the IP header ToS byte and so overlaps with the precedence bits. Both IP precedence and DSCP classification rules are supported. DSCP rules have a higher match priority than IP precedence rules and where a *dscp-name* DSCP value overlaps an *ip-prec-value*, the DSCP rule takes precedence.

The **no** form of this command removes the explicit IP precedence classification rule from the SAP ingress policy. Removing the rule on the policy immediately removes the rule on all ingress SAPs using the policy.

Parameters

ip-prec-value

The *ip-prec-value* is a required parameter that specifies the unique IP header ToS byte precedence bits value that will match the IP precedence rule. If the command is executed more than once with the same *ip-prec-value*, the previous forwarding class and enqueuing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueuing priority parameter is missing.

A maximum of eight IP precedence rules are allowed on a single policy.

The precedence is evaluated from the lowest to highest value.

Values 0 to 7

fc *fc-name*

The value given for the *fc-name* parameter must be one of the predefined forwarding classes in the system. Specifying the *fc-name* is optional. When a packet matches the rule, the forwarding class is only overridden when the **fc** *fc-name* parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

The *subclass-name* parameter is optional and used with the *fc-name* parameter to define a pre-existing subclass. The *fc-name* and *subclass-name* parameters must be separated by a period (.). If *subclass-name* does not exist in the context of *fc-name*, an error will occur. If *subclass-name* is removed using the **no fc** *fc-name.subclass-name* **force** command, the

default-fc command will automatically drop the *subclass-name* and only use *fc-name* (the parent forwarding class for the subclass) as the forwarding class.

Values fc: *class*[.*subclass*]

class: be, l2, af, l1, h2, ef, h1, nc

subclass: 29 characters max

Default Inherit (When **fc** is not defined, the rule preserves the previous forwarding class of the packet.)

priority

The priority parameter overrides the default enqueuing priority for all packets received on an ingress SAP using this policy that match this rule. Specifying the priority is optional. When a packet matches the rule, the enqueuing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and priority is not explicitly defined in the rule, the enqueuing priority is inherited based on previous rule matches.

Values high, low

Default Inherits the priority defined by the default-priority statement.

high

This parameter is used in conjunction with the **priority** parameter. Setting the enqueuing parameter to **high** for a packet increases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

low

This parameter is used in conjunction with the **priority** parameter. Setting the enqueuing parameter to **low** for a packet decreases the likelihood of enqueuing the packet when the ingress queue is congested. Ingress enqueuing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueuing priority is lost.

Platforms

7705 SAR Gen 2

prec

Syntax

prec {*ip-prec-value* | **in-profile** *ip-prec-value* **out-profile** *ip-prec-value* [**exceed-profile** *ip-prec-value*]}

no prec

Context

[Tree] (config>qos>sap-egress>fc prec)

Full Context

configure qos sap-egress fc prec

Description

This command defines a value to be used for remarking packets for the specified FC. If the optional in/out/exceed-profile is specified, the command will remark different IP precedence values depending on whether the packet was classified to be in, exceed, or out-of-profile. All inplus-profile traffic is marked with the same value as in-profile traffic.

Parameters

ip-prec-value

This parameter specifies the IP precedence to be used to remark all traffic.

Values 0 to 7

exceed-profile *ip-prec-value*

This optional parameter specifies the IP precedence to be used to remark traffic that is exceed-profile. If not specified, this defaults to the same value configured for the **out-profile** parameter.

Values 0 to 7

in-profile *ip-prec-value*

This parameter specifies the IP precedence to be used to remark traffic that is in-profile.

Values 0 to 7

out-profile *ip-prec-value*

This parameter specifies the IP precedence to be used to remark traffic that is out-of-profile.

Values 0 to 7

Platforms

7705 SAR Gen 2

prec

Syntax

prec *ip-prec-value* [**fc** *fc-name*] [**profile** {**in** | **out** | **exceed** | **inplus**}]

no prec *ip-prec-value*

Context

[Tree] (config>qos>sap-egress prec)

Full Context

configure qos sap-egress prec

Description

This command defines a specific IP precedence value that must be matched to perform the associated reclassification actions. If an egress packet on the SAP matches the specified IP precedence value, the forwarding class, or profile behavior may be overridden. By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions.

The IP precedence bits used to match against precedence reclassification rules come from the Type of Service (ToS) field within the IPv4 header. If the packet does not have an IPv4 header, precedence-based matching is not performed.

The reclassification actions from a precedence reclassification rule may be overridden by a DSCP or IP flow matching event.

The **fc** keyword is optional. When specified, the egress classification rule will overwrite the forwarding class derived from ingress. The new forwarding class is used for egress remarking and queue mapping decisions. If a DSCP, ipv6-criteria, or ip-criteria match occurs after the IP precedence match, the new forwarding class may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new fc, the fc from the IP precedence match will be used.

The **profile** keyword is optional. When specified, the egress classification rule will overwrite the profile of the packet derived from ingress. The new profile value is used for egress remarking and queue congestion behavior. If a DSCP, IPv6 criteria, or IP criteria match occurs after the IP precedence match, the new profile may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new profile, the profile from the IP precedence match will be used.

The **no** form of this command removes the reclassification rule from the SAP egress QoS policy.

Parameters

fc *fc-name*

This keyword is optional. When specified, packets matching the IP precedence value will be explicitly reclassified to the forwarding class specified as *fc-name* regardless of the ingress classification decision. The explicit forwarding class reclassification may be overwritten by a higher priority DSCP, IPv6 criteria, or IP criteria reclassification match. The FC name defined must be one of the eight forwarding classes supported by the system. To remove the forwarding class reclassification action for the specified precedence value, the **prec** command must be re-executed without the **fc** parameter defined.

Values be, l1, af, l2, h1, ef, h2 or nc

profile {in | out | exceed | inplus}

This keyword is optional. When specified, packets matching the IP precedence value will be explicitly reclassified to the specified profile regardless of the ingress profiling decision. The explicit profile reclassification may be overwritten by a higher priority DSCP, IPv6 criteria, or IP criteria reclassification match. To remove the profile reclassification action for the specified precedence value, the **prec** command must be re-executed without the **profile** parameter defined.

in

Specifies that any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.

out

Specifies that any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.

exceed

Specifies that any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.

inplus

Specifies that any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

Platforms

7705 SAR Gen 2

prec

Syntax

prec *ip-prec-value* **fc** *fc-name* **profile** {**in** | **out** | **exceed** | **inplus**}

no prec *ip-prec-value*

Context

[Tree] (config>qos>network>egress prec)

Full Context

configure qos network egress prec

Description

This command defines a specific IP precedence value that must be matched in order to perform the associated reclassification actions. If an egress packet on an IES/VPRN interface spoke SDP, on a CSC network interface in a VPRN, or network interface that the network QoS policy is applied to, matches the specified IP precedence value, the forwarding class and profile may be overridden.

By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions.

The IP precedence bits used to match against the reclassification rules come from the Type of Service (ToS) field within the IPv4 header or the Traffic Class field from the IPv6 header. If the packet does not have an IP header, IP precedence-based matching is not performed.

The configuration of egress prec classification and the configuration of an egress IP criteria or IPv6 criteria entry statement within a network QoS policy are mutually exclusive.

The IP precedence-based and DSCP-based reclassification are supported on a network interface, on a CSC network interface in a VPRN, and on a PW used in an IES or VPRN spoke interface.

This command will block the application of a network QoS policy with the egress reclassification commands to a spoke SDP part of a Layer 2 service. Conversely, this command will not allow the user to add the egress reclassification commands to a network QoS policy if it is being used by a Layer 2 spoke SDP.

The egress reclassification commands will only take effect if the redirection of the spoke SDP or CSC interface to use an egress port queue-group succeeds. For example, the following commands will succeed:

```
    -
 config>service>vprn>if>
spoke-sdp>egress>qos network-policy-id port-redirect-
group
queue-group-name instance instance-id
    - config>service>ies>if>spoke-
sdp>
egress>qos network-policy-id port-redirect-group queue-group-
name
instance instance-id
    - config>service>vprn>nw-if> qos network-policy-id port-redirect-
group
queue-group-name instance instance-id
```

When the redirection command fails in CLI, the PW will use the network QoS policy assigned to the network IP interface; however, any reclassification in the network QoS policy applied to the network interface will be ignored.

The **no** form of this command removes the egress reclassification rule.

Parameters

ip-prec-value

0 to 7

fc fc-name

be, l2, af, l1, h2, ef, h1, nc

profile {in | out | exceed | inplus}

The profile reclassification action is mandatory. When specified, packets matching the IP precedence value will be explicitly reclassified to the profile specified regardless of the ingress profiling decision. To remove the profile reclassification action for the specified IP precedence value, the **no prec** command must be executed.

This value may be overwritten by an explicit profile action in an DSCP reclassification match.

in - Specifies that any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.

out - Specifies that any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.

exceed - Specifies that any packets matching the reclassification rule will be treated as exceed-profile by the egress forwarding plane.

inplus - Specifies that any packets matching the reclassification rule will be treated as inplus-profile by the egress forwarding plane.

Platforms

7705 SAR Gen 2

22.52 precedence

precedence

Syntax

precedence [*precedence-value* | **primary**]

no precedence

Context

[Tree] (config>service>epipe>spoke-sdp precedence)

Full Context

configure service epipe spoke-sdp precedence

Description

This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic.

The **no** form of this command returns the precedence value to the default.

Default

precedence 4

Parameters

precedence-value

Specifies the spoke SDP precedence.

Values 1 to 4

primary

Assigns primary precedence to the spoke SDP.

Platforms

7705 SAR Gen 2

precedence

Syntax

precedence *prec-value*

precedence primary

no precedence

Context

[Tree] (config>service>epipe>spoke-sdp-fec precedence)

Full Context

configure service epipe spoke-sdp-fec precedence

Description

This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic.

The **no** form of this command returns the precedence value to the default.

Default

precedence 42

Parameters

prec-value

Specifies the spoke SDP precedence.

Values 1 to 4

primary

Assigns primary precedence to this spoke SDP.

Platforms

7705 SAR Gen 2

precedence

Syntax

precedence [*precedence-value* | **primary**]

no precedence

Context

[Tree] (config>service>vpls>spoke-sdp precedence)

Full Context

configure service vpls spoke-sdp precedence

Description

This command configures the precedence of this SDP bind when there are multiple SDP binds attached to one service endpoint. When an SDP bind goes down, the next highest precedence SDP bind begins forwarding traffic.

Parameters

precedence-value

Specifies the precedence of this SDP bind

Values 1 to 4

primary

Assigns this as the primary spoke-SDP

Platforms

7705 SAR Gen 2

precedence

Syntax

precedence {*precedence-value* | **primary**}

no precedence

Context

[Tree] (config>mirror>mirror-dest>spoke-sdp precedence)

Full Context

configure mirror mirror-dest spoke-sdp precedence

Description

This command indicates that the SDP is of type secondary with a specific precedence value or of type primary.

The mirror or LI service always uses the primary type as the active pseudowire and only switches to a secondary pseudowire when the primary is down. The mirror service switches the path back to the primary pseudowire when it is back up. The user can configure a timer to delay reverting back to primary or to never revert back.

If the active pseudowire goes down, the mirror service switches the path to a secondary sdp with the lowest precedence value. That is, secondary SDPs which are operationally up are considered in the order of their precedence value, 1 being the lowest value and 4 being the highest value. If the precedence value is the same, then the SDP with the lowest SDP ID is selected.

An explicitly named endpoint can have a maximum of one SAP and one ICB. Once a SAP is added to the endpoint, only one more object of type ICB SDP is allowed. An explicitly named endpoint, which does not have a SAP object, can have a maximum of four SDPs, which can include any of the following: a single primary SDP, one or many secondary SDPs with precedence, and a single ICB SDP.

An SDP is created with type secondary and with the lowest precedence value of 4.

Parameters

precedence-value

Specifies the precedence of the SDP.

Values 1 to 4

primary

Specified that a special value of the precedence which assigns the SDP the lowest precedence and enables the revertive behavior.

Platforms

7705 SAR Gen 2

22.53 preempt

preempt

Syntax

[**no**] **preempt**

Context

[Tree] (config>service>ies>if>ipv6>vrrp preempt)

Full Context

configure service ies interface ipv6 vrrp preempt

Description

The preempt mode value controls whether a specific backup virtual router preempts a lower priority master.

When preempt is enabled, the virtual router instance overrides any non-owner master with an "in use" message priority value less than the virtual router instance in-use priority value. If preempt is disabled, the virtual router only becomes master if the master down timer expires before a VRRP advertisement message is received from another virtual router.

The IP address owner will always become master when available. Preempt mode cannot be disabled on the owner virtual router.

The **no** form of this command disables preempt mode.

Default

preempt

Platforms

7705 SAR Gen 2

preempt

Syntax

[**no**] **preempt**

Context

[Tree] (config>service>ies>if>vrrp preempt)

Full Context

configure service ies interface vrrp preempt

Description

The preempt command provides the ability of overriding an existing non-owner master to the virtual router instance. Enabling preempt mode is almost required for proper operation of the base-priority and vrrp-policy-id definitions on the virtual router instance. If the virtual router cannot preempt an existing non-owner master, the effect of the dynamic changing of the in-use priority is greatly diminished.

The preempt command is only available in the non-owner vrrp virtual-router-id nodal context. The owner may not be preempted due to the fact that the priority of non-owners can never be higher than the owner. The owner will always preempt all other virtual routers when it is available.

Non-owner virtual router instances will only preempt when preempt is set and the current master has an in-use message priority value less than the virtual router instances in-use priority.

A master non-owner virtual router will only allow itself to be preempted when the incoming VRRP Advertisement message Priority field value is one of the following:

- Greater than the virtual router in-use priority value
- Equal to the in-use priority value and the source IP address (primary IP address) is greater than the virtual router instance primary IP address

The **no** form of this command prevents a non-owner virtual router instance from preempting another, less desirable virtual router. Use the preempt command to restore the default mode.

Default

preempt

Platforms

7705 SAR Gen 2

preempt

Syntax

[**no**] **preempt**

Context

[Tree] (config>service>vprn>if>ipv6>vrrp preempt)

Full Context

configure service vprn interface ipv6 vrrp preempt

Description

The preempt mode value controls whether a specific backup virtual router preempts a lower priority master.

When preempt is enabled, the virtual router instance overrides any non-owner master with an "in use" message priority value less than the virtual router instance in-use priority value. If preempt is disabled, the virtual router only becomes master if the master down timer expires before a VRRP advertisement message is received from another virtual router.

The IP address owner will always become master when available. Preempt mode cannot be disabled on the owner virtual router.

The default value for preempt mode is enabled.

Default

preempt

Platforms

7705 SAR Gen 2

preempt

Syntax

[**no**] **preempt**

Context

[Tree] (config>router>if>ipv6>vrrp preempt)

[Tree] (config>router>if>vrrp preempt)

Full Context

configure router interface ipv6 vrrp preempt

configure router interface vrrp preempt

Description

The preempt mode value controls whether a specific backup virtual router preempts a lower priority master.

When preempt is enabled, the virtual router instance overrides any non-owner master with an "in use" message priority value less than the virtual router instance in-use priority value. If preempt is disabled, the virtual router only becomes master if the master down timer expires before a VRRP advertisement message is received from another virtual router.

The IP address owner will always become master when available. Preempt mode cannot be disabled on the owner virtual router.

The default value for preempt mode is enabled.

Default

preempt

Platforms

7705 SAR Gen 2

22.54 preemption-timer

preemption-timer

Syntax

preemption-timer *seconds*

no preemption-timer

Context

[Tree] (config>router>rsvp preemption-timer)

Full Context

configure router rsvp preemption-timer

Description

This parameter configures the time in seconds a node holds to a reservation for which it triggered the soft preemption procedure.

The preempting node starts a separate preemption timer for each preempted LSP path. While this timer is on, the node should continue to refresh the Path and Resv for the preempted LSP paths. When the preemption timer expires, the node tears down the reservation if the head-end node has not already done so.

A value of zero means the LSP should be preempted immediately; hard preempted.

The **no** form of this command reverts to the default value.

Default

preemption-timer 300

Parameters

seconds

Specifies the time (in s), of the preemption timer.

Values 0 to 1800 seconds

Platforms

7705 SAR Gen 2

22.55 prefer-local-time

prefer-local-time

Syntax

[**no**] **prefer-local-time**

Context

[Tree] (config>system>time prefer-local-time)

Full Context

configure system time prefer-local-time

Description

This command sets the preference to use local or UTC time in the system. This preference is applied to objects such as log file names, created and completed times reported in log files, NETCONF and gRPC date-and-time leafs, and rollback times displayed in **show** routines.

Note:
The operator may force the timezone used for **show** outputs during a CLI session using an environment variable in the **environment**>**time-display {utc | local}** command.

Note:
The preference for CLI output is set with the **environment time-display** command.

Note:
The format used for the date-time strings may change when the **prefer-local-time** option is enabled. For example, when enabled, all date-time strings include a suffix of three to five characters that indicates the timezone used for the presentation. This suffix may not be present if the option in not enabled.

Note:
The time format for timestamps on log events is controlled on a per-log basis using the **config**>**log**>**log-id**>**time-format {utc | local}** CLI command and not via **prefer-local-time**.

The **no** form of this command indicates preference for UTC time.

Default

no prefer-local-time

Platforms

7705 SAR Gen 2

22.56 prefer-protocol-stitching

prefer-protocol-stitching

Syntax

[**no**] **prefer-protocol-stitching**

Context

[Tree] (config>router>ldp prefer-protocol-stitching)

Full Context

configure router ldp prefer-protocol-stitching

Description

This command stitches an LDP ILM to an SR NHLFE rather than to an LDP NHLFE when both LDP and SR NHLFEs exist.

The **no** form of this command stitches an LDP ILM to an LDP NHLFE by preference over an SR NHLFE.

Default

no prefer-protocol-stitching

Platforms

7705 SAR Gen 2

23 p Commands – Part III

23.1 prefer-tunnel-in-tunnel

prefer-tunnel-in-tunnel

Syntax

[**no**] **prefer-tunnel-in-tunnel**

Context

[Tree] (config>router>ldp prefer-tunnel-in-tunnel)

Full Context

configure router ldp prefer-tunnel-in-tunnel

Description

This command specifies to use tunnel-in-tunnel over a simple LDP tunnel. Specifically, the user packets for LDP FECs learned over this targeted LDP session can be sent inside an RSVP LSP which terminates on the same egress router as the destination of the targeted LDP session. The user can specify an explicit list of RSVP LSP tunnels under the Targeted LDP session or LDP will perform a lookup in the Tunnel Table Manager (TTM) for the best RSVP LSP. In the former case, only the specified LSPs will be considered to tunnel LDP user packets. In the latter case, all LSPs available to the TTM and which terminate on the same egress router as this target ed LDP session will be considered. In both cases, the metric specified under the LSP configuration is used to control this selection.

The lookup in the TTM will prefer a LDP tunnel over an LDP-over-RSVP tunnel if both are available. Also, the tunneling operates on the dataplane only. Control packets of this targeted LDP session are sent over the IGP path.

Platforms

7705 SAR Gen 2

23.2 preference

preference

Syntax

[**no**] **preference** *preference*

Context

[Tree] (config>service>vprn>bgp preference)

[Tree] (config>service>vprn>bgp>group preference)

Full Context

configure service vprn bgp preference

configure service vprn bgp group preference

Description

This command configures the route preference for routes learned from the configured peer(s).

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The lower the preference the higher the chance of the route being the active route. The OS assigns BGP routes highest default preference compared to routes that are direct, static or learned via MPLS or OSPF.

The **no** form of this command, if used at the global level, reverts to default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

preference 170

Parameters

preference

Specifies the route preference, expressed as a decimal integer.

Values 1 to 255

Platforms

7705 SAR Gen 2

preference

Syntax

preference *preference-value*

no preference

Context

[Tree] (config>service>vprn>static-route-entry>next-hop preference)

[Tree] (config>service>vprn>static-route-entry>black-hole preference)

[Tree] (config>service>vprn>static-route-entry>indirect preference)

[Tree] (config>service>vprn>static-route-entry>ipsec-tunnel preference)

[Tree] (config>service>vprn>static-route-entry>grt preference)

Full Context

configure service vprn static-route-entry next-hop preference

configure service vprn static-route-entry black-hole preference

configure service vprn static-route-entry indirect preference

configure service vprn static-route-entry ipsec-tunnel preference

configure service vprn static-route-entry grt preference

Description

This command specifies the route preference to be assigned to the associated static route. The lower the preference value the more preferred the route is considered.

Table 75: Default Route Preference lists the default route preference based on the route source.

Table 75: Default Route Preference

Label	Preference	Configurable
Direct attached	0	No
Static route	5	Yes
OSPF Internal routes	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
Aggregate	130	No
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

The **no** form of this command returns the returns the associated static route preference to its default value.

Default

preference 5

Parameters

preference-value

Specifies the route preference value.

Values 1 to 255

Platforms

7705 SAR Gen 2

preference

Syntax

preference *preference*

no preference

Context

[Tree] (config>service>vprn>isis>level preference)

Full Context

configure service vprn isis level preference

Description

This command configures the preference level of either IS-IS Level 1 or IS-IS Level 2 internal routes. By default, the preferences are listed in the table below.

A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide to which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the table below. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision what route to use is determined by the configuration of the **ecmp** in the config>router context.

Default

Default preferences are listed in Table 76: Default Preferences.

Table 76: Default Preferences

Route Type	Preference	Configurable
Direct attached	0	No
Static route	5	Yes
MPLS	7	—
OSPF internal routes	10	No
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes

Route Type	Preference	Configurable
OSPF external	150	Yes
IS-IS level 1 external	160	Yes[5]
IS-IS level 2 external	165	Yes[5]
BGP	170	Yes

Parameters

preference

The preference for external routes at this level expressed as a decimal integer.

Values 1 to 255

Platforms

7705 SAR Gen 2

preference

Syntax

preference *preference*

no preference

Context

[Tree] (config>service>vprn>ospf3 preference)

[Tree] (config>service>vprn>ospf preference)

Full Context

configure service vprn ospf3 preference

configure service vprn ospf preference

Description

This command configures the preference for OSPF internal routes.

A route can be learned by the router from different protocols in which case the costs are not comparable, when this occurs the preference is used to decide to which route will be used.

Different protocols should not be configured with the same preference. If the same preference is configured, the tiebreaker is per the default preference table as defined in Table 77: Default Route Preferences . If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

5 External preferences are changed using the **external-preference** command in the **config>router>isis>level** *level-number* context.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the **config**>**router** context.

The **no** form of this command reverts to the default value.

Table 77: Default Route Preferences

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes[6]
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes

Default

preference 10 — OSPF internal routes have a preference of 10.

Parameters

preference

The preference for internal routes expressed as a decimal integer. Defaults for different route types are listed in the following table.

Values 1 to 255

Platforms

7705 SAR Gen 2

preference

Syntax

preference *preference*

no preference

6 Preference for OSPF internal routes is configured with the **preference** command.

Context

[Tree] (config>service>vprn>ripng preference)

[Tree] (config>service>vprn>rip>group>neighbor preference)

[Tree] (config>service>vprn>ripng>group preference)

[Tree] (config>service>vprn>ripng>group>neighbor preference)

[Tree] (config>service>vprn>rip preference)

[Tree] (config>service>vprn>rip>group preference)

Full Context

configure service vprn ripng preference

configure service vprn rip group neighbor preference

configure service vprn ripng group preference

configure service vprn ripng group neighbor preference

configure service vprn rip preference

configure service vprn rip group preference

Description

This command sets the route preference assigned to RIP routes. This value can be overridden by route policies.

The **no** form of this command resets the *preference* to the default.

Default

no preference

Parameters

preference

Specifies the preference value.

Values 1 to 255

Default 100

Platforms

7705 SAR Gen 2

preference

Syntax

preference *preference-value*

no preference

Context

[Tree] (config>router>mpls>fwd-policies>fwd-policy preference)

Full Context

configure router mpls forwarding-policies forwarding-policy preference

Description

This command configures the preference of an MPLS forwarding policy.

The **no** form of this command removes the preference parameter from the MPLS forwarding policy.

Default

preference 255

Parameters

preference-value

Specifies the preference value.

The *preference-value* parameter allows the user to configure multiple label-binding forwarding policies with the same binding label or multiple endpoint policies with the same endpoint address. This provides the capability to achieve a 1:N backup strategy for the forwarding policy. Only the most preferred, lowest numerically preference value, policy is activated in data path.

Values 1 to 255

Platforms

7705 SAR Gen 2

preference

Syntax

preference *preference*

no preference

Context

[Tree] (config>router>static-route-entry>indirect preference)

[Tree] (config>router>static-route-entry>next-hop preference)

[Tree] (config>router>static-route-entry>black-hole preference)

Full Context

configure router static-route-entry indirect preference

configure router static-route-entry next-hop preference

configure router static-route-entry black-hole preference

Description

This command specifies the route preference to be assigned to the associated static route. The lower the preference value the more preferred the route is considered.

Table 78: Default Route Preference shows the default route preference based on the route source.

Table 78: Default Route Preference

Label	Preference	Configurable
Direct attached	0	No
Static route	5	Yes
OSPF Internal routes	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
Aggregate	130	No
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

The **no** form of this command returns the returns the associated static route preference to its default value.

Default

preference 5

Parameters

preference

Specifies the route preference value.

Values 1 to 255

Platforms

7705 SAR Gen 2

preference

Syntax

[**no**] **preference** *preference*

Context

[Tree] (config>router>bgp>group preference)

[Tree] (config>router>bgp preference)

[Tree] (config>router>bgp>group>neighbor preference)

Full Context

configure router bgp group preference

configure router bgp preference

configure router bgp group neighbor preference

Description

This command configures the route preference for routes learned from the configured peers.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The lower the preference the higher the chance of the route being the active route. The router assigns BGP routes highest default preference compared to routes that are direct, static or learned via MPLS or OSPF.

The **no** form of this command used at the global level reverts to default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

preference 170

Parameters

preference

Specifies the route preference expressed as a decimal integer.

Values 1 to 255

Platforms

7705 SAR Gen 2

preference

Syntax

preference *preference*

no preference

Context

[Tree] (config>router>isis>level preference)

Full Context

configure router isis level preference

Description

This command configures the preference level of either IS-IS Level 1 or IS-IS Level 2 internal routes. By default, the preferences are listed in the table below.

A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide to which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision what route to use is determined by the configuration of the **ecmp** in the **config>router** context.

Default

preference (Level 1) — 15

preference (Level 2) — 18

Parameters

preference

Specifies the preference for external routes at this level expressed as a decimal integer. The default preferences are listed in Table 79: Default Internal Route Preferences .

Table 79: Default Internal Route Preferences

Route Type	Preference	Configurable
Direct attached	0	—
Static-route	5	Yes
OSPF internal routes	10	—
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes[7]
IS-IS level 2 external	165	Yes[7]

[7] External preferences are changed using the external-preference command in the **config>router>isis>level** *level-number* context.

Route Type	Preference	Configurable
BGP	170	Yes

Values 1 to 255

Platforms

7705 SAR Gen 2

preference

Syntax

preference *preference*

no preference

Context

[Tree] (config>router>ospf preference)

[Tree] (config>router>ospf3 preference)

Full Context

configure router ospf preference

configure router ospf3 preference

Description

This command configures the preference for OSPF internal routes.

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in Table 80: Route Preference Defaults by Route Type . If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the **config>router** context.

The **no** form of this command reverts to the default value.

Default

preference 10

Parameters

preference

Specifies the preference for internal routes expressed as a decimal integer. Defaults for different route types are listed in Table 80: Route Preference Defaults by Route Type .

Table 80: Route Preference Defaults by Route Type

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes[8]
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

Values 1 to 255

Platforms

7705 SAR Gen 2

preference

Syntax

preference {**none** | **all**}

no preference

Context

[Tree] (config>router>ospf>lfa>mhp preference)

[Tree] (config>router>isis>lfa>mhp preference)

Full Context

configure router ospf loopfree-alternates multi-homed-prefix preference

configure router isis loopfree-alternates multi-homed-prefix preference

8 Preference for OSPF internal routes is configured with the **preference** command.

Description

This command configures the preference for the multihomed prefix LFA backup path. This knob can be enabled at a LFA computing node to force the programming of the multihomed prefix LFA backup path which, in some topologies, can avoid transiting using the best ABR or ASBR.

The **no** form of this command reverts to the default value.

Default

preference none

Parameters

none

Specifies the preference for an LFA, TI-LFA, or RLFA backup path over the multihomed prefix LFA backup path. The multihomed prefix LFA is only programmed in cases where the prefix is not protected by LFA, RLFA, or TI-LFA.

all

Specifies the forced programming of the multihomed prefix LFA backup path regardless of the outcome of the LFA, TI-LFA, or RLFA backup path computation.

Platforms

7705 SAR Gen 2

preference

Syntax

preference *preference*

no preference

Context

[Tree] (config>router>rip>group preference)

[Tree] (config>router>ripng>group>neighbor preference)

[Tree] (config>router>rip>group>neighbor preference)

[Tree] (config>router>ripng preference)

[Tree] (config>router>rip preference)

[Tree] (config>router>ripng>group preference)

Full Context

configure router rip group preference

configure router ripng group neighbor preference

configure router rip group neighbor preference

configure router ripng preference

configure router rip preference

configure router ripng group preference

Description

This command configures the preference for RIP routes.

A route can be learned by the router from different protocols in which case the costs are not comparable. When this occurs, the preference is used to decide which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in Table 81: Route Preference Defaults by Route Type . If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the **config>router** context.

The **no** form of the command reverts to the default value.

Default

preference 100

Parameters

preference

Specifies the preference for RIP routes expressed as a decimal integer. Defaults for different route types are listed in Table 81: Route Preference Defaults by Route Type .

Table 81: Route Preference Defaults by Route Type

Route Type	Preference	Configurable
Direct attached	0	—
Static routes	5	Yes
OSPF internal	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

Values 0 to 255

Platforms

7705 SAR Gen 2

preference

Syntax

preference *preference*

Context

[Tree] (conf>router>segment-routing>sr-policies>policy preference)

Full Context

configure router segment-routing sr-policies static-policy preference

Description

This command associates a preference value with a statically defined-segment routing policy. This is an optional parameter.

When there are multiple policies for the same (color, endpoint) combination that are targeted for local installation, only one is selected as the active path for the (color, endpoint). In this selection process (which considers both static local policies and BGP signaled policies), the policy with the highest preference value is preferred over all policies with a lower preference value.

The **no** form of this command reverts to the default value.

Default

preference 100

Parameters

preference

Specifies the preference ID.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

preference

Syntax

preference *preference*

no preference

Context

[Tree] (config>router>policy-options>policy-statement>entry>action preference)

Full Context

configure router policy-options policy-statement entry action preference

Description

This command assigns a route preference to routes matching the route policy statement entry.

If no preference is specified, the default Route Table Manager (RTM) preference for the protocol is used.

The **no** form of this command disables setting an RTM preference in the route policy entry.

Note:
This command is supported with the following protocols: RIP import, BGP import, VPRN VRF import (**vrf-import**), and VPRN GRT lookup export (**export-grt**).

Default

no preference

Parameters

preference

Specifies the route preference expressed as a decimal integer.

Values 1 to 255 (0 represents unset - MIB only)

name — The preference parameter variable name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@”.

Platforms

7705 SAR Gen 2

preference

Syntax

preference

Context

[Tree] (config>ipsec>tnl-temp>rev-route preference)

Full Context

configure ipsec tunnel-template reverse-route preference

Description

This command configures the route preference assigned to the DL2L tunnel reverse routes. The system uses this preference when selecting a route to install in the route table.

Default

preference 0

Parameters

preference

Specifies the preference value for reverse routes.

Values 0 to 255

Platforms

7705 SAR Gen 2

23.3 preferred-lifetime

preferred-lifetime

Syntax

preferred-lifetime [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no preferred-lifetime

Context

[Tree] (config>service>vprn>dhcp6>server>pool>prefix preferred-lifetime)

[Tree] (config>router>dhcp6>server>pool>prefix preferred-lifetime)

Full Context

configure service vprn dhcp6 local-dhcp-server pool prefix preferred-lifetime

configure router dhcp6 local-dhcp-server pool prefix preferred-lifetime

Description

This command configures the preferred lifetime.

The **no** form of this command reverts to the default value.

Default

preferred-lifetime hrs 1

Parameters

preferred-lifetime

Specifies the preferred time for a prefix.

Values

days:	0 to 3650
hours:	0 to 23
minutes:	0 to 59
seconds	0 to 59

Platforms

7705 SAR Gen 2

preferred-lifetime

Syntax

[**no**] **preferred-lifetime** {*seconds* | **infinite**}

Context

[Tree] (config>service>vprn>router-advert>if>prefix preferred-lifetime)

[Tree] (config>router>router-advert>if>prefix preferred-lifetime)

Full Context

configure service vprn router-advertisement interface prefix preferred-lifetime

configure router router-advertisement interface prefix preferred-lifetime

Description

This command configures the remaining length of time in seconds that this prefix will continue to be preferred, such as, time until deprecation. The address generated from a deprecated prefix should not be used as a source address in new communications, but packets received on such an interface are processed as expected.

Default

preferred-lifetime 604800

Parameters

seconds

Specifies the remaining length of time in seconds that this prefix will continue to be preferred.

Values 0 to 4294967294

infinite

Specifies that the prefix will always be preferred. A value of 4,294,967,295 represents infinity.

Platforms

7705 SAR Gen 2

23.4 prefix

prefix

Syntax

prefix *ipv6-addr/prefix-length* [**failover** {**local** | **remote** | **access-driven**}] [**pd**] [**wan-host**] [**create**]

no prefix *ipv6-addr/prefix-length*

Context

[Tree] (config>router>dhcp6>server>pool prefix)

[Tree] (config>service>vprn>dhcp6>server>pool prefix)

Full Context

configure router dhcp6 local-dhcp-server pool prefix

configure service vprn dhcp6 local-dhcp-server pool prefix

Description

This command allocates a prefix to a pool from which Prefix Delegation prefixes and or WAN addresses can be assigned for DHCP6.

The **no** form of this command removes the prefix parameters from the configuration.

Default

prefix failover local

Parameters

prefix ipv6-addr/prefix-length

Specifies the prefix.

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x [0 to FFFF]H
	d [0 to 255]D

prefix-length 1 to 128

failover {local | remote | access-driven}

This command designates a prefix as local, remote, or access-driven. This is used when multi-chassis synchronization is enabled.

Values **local** — An IPv6 prefix designated as local is used for new lease grants or to renew the existing lease grants. Local prefix designation should be always paired with the remote designation of the same prefix on the peering node.

The IPv6 prefix configured as local on one node can only be configured as remote on the other node. No other combination is allowed between the two nodes for an IPv6 prefix that is configured as local.

The DHCPv6 relay could point to both IPv6 DHCP server addresses — the one hosting the local IPv6 prefix and the one hosting the corresponding remote IPv6 prefix. Under normal circumstances the new lease will always be allocated from the local IPv6 prefix while the leases can be renewed from either IPv6 prefix (local or remote). Under network failure, the remote IPv6 prefix can be taken over according to the intercommunication link state transitions and associated timers.

remote — A prefix designated as remote is used only to renew the existing DHCP leases. The new leases are assigned from it only after the **maximum-client-lead-time** and **partner-down-delay** *time* elapses.

To ensure faster takeover, the partner-down-delay can be set to 0 and the MCLT time can be ignored. Extra caution should be exercised when enabling this mode of operation, as described in the configuration guides.

The IPv6 prefix configured as remote on one node can only be configured as local on the other node. No other combination is allowed between the two nodes for an IP address ranges that is configured as remote.

access-driven — A prefix designated as access-driven is like local (a new prefix assignment as well as a renewal). However, as the prefix is shared between the redundant server pair, the following additional conditions should be met to avoid duplicate address allocations:

- A dual home access protection mechanism such as SRRP or MC-LAG must ensure a single active path from the

  DHCP client to the server.
- The DHCP relay should point to the local server only.

pd

Specifies that this aggregate is used by IPv6 ESM hosts for DHCPv6 prefix-delegation.

wan-host

Specifies that this aggregate is used by IPv6 ESM hosts for local addressing or by a routing gateway's WAN interface.

create

Keyword used to create the prefix configuration. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

prefix

Syntax

prefix *low-order-vsi-id*

no prefix

Context

[Tree] (config>service>vpls>bgp-ad>vsi-id prefix)

Full Context

configure service vpls bgp-ad vsi-id prefix

Description

This command specifies the low-order 4 bytes used to compose the Virtual Switch Instance Identifier (VSI-ID) to use for NLRI in BGP auto-discovery in this VPLS service.

If no value is set, the system IP address will be used.

Default

no prefix

Parameters

low-order-vsi-id

Specifies a unique VSI ID

Values 0— 4294967295

Platforms

7705 SAR Gen 2

prefix

Syntax

[**no**] **prefix** *ip-prefix*/*prefix-length*

Context

[Tree] (config>service>vprn>bgp>group>dynamic-neighbor>match prefix)

Full Context

configure service vprn bgp group dynamic-neighbor match prefix

Description

This command configures a prefix to accept dynamic BGP sessions (sessions from source IP addresses not matching any configured neighbor addresses). A dynamic session is associated with the group having the longest match prefix entry for the source IP address of the peer. The group association determines local parameters that apply to the session, including the local AS, the local IP address, the MP-BGP families, the import and export policies, and so on.

The **no** form of this command removes a prefix entry.

Parameters

ip-prefix/prefix-length

Specifies a prefix from which to accept dynamic BGP sessions.

Values *ipv4-prefix* — a.b.c.d (host bits must be 0)

ipv4-prefix-length — 0 to 32

ipv6-prefix — x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x — [0 to FFFF]H

d — [0 to 255]D

ipv6-prefix-length — 0 to 128

Platforms

7705 SAR Gen 2

prefix

Syntax

[**no**] **prefix** *ipv6-prefixIprefix-length*

Context

[Tree] (config>service>vprn>router-advert>if prefix)

Full Context

configure service vprn router-advertisement interface prefix

Description

This command configures an IPv6 prefix in the router advertisement messages. To support multiple IPv6 prefixes, use multiple prefix statements. No prefix is advertised until explicitly configured using prefix statements.

Parameters

ipv6-prefix

Specifies the IP prefix for prefix list entry in dotted decimal notation.

Values

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 32
ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6-prefix-length	0 to 128

prefix-length

Specifies a route must match the most significant bits and have a prefix length.

Values 1 to 128

Platforms

7705 SAR Gen 2

prefix

Syntax

prefix *ip-prefix/prefix-length* [**create**]

no prefix *ip-prefix/prefix-length*

Context

[Tree] (config>test-oam>twamp>server prefix)

Full Context

configure test-oam twamp server prefix

Description

This command configures an IP address prefix containing one or more TWAMP clients. For a TWAMP client to connect to the TWAMP server (and subsequently conduct tests) it must establish the control connection using an IP address that is part of a configured prefix.

Parameters

ip-prefix/prefix-length

Specifies an IPv4 or IPv6 address prefix.

Values

ipv4-prefix:	a.b.c.d (host bits must be 0)
ipv4-prefix-le:	0 to 32
ipv6-prefix:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6-prefix-le:	0 to 128

prefix length

Specifies the prefix length.

Values 0 to128

create

Creates a prefix instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

prefix

Syntax

prefix *ip-prefix/prefix-length* [**create**]

no prefix *ip-prefix/prefix-length*

Context

[Tree] (config>service>vprn>twamp-light>reflector prefix)

[Tree] (config>router>twamp-light>reflector prefix)

Full Context

configure service vprn twamp-light reflector prefix

configure router twamp-light reflector prefix

Description

This command defines which TWAMP Light packet prefixes the reflector processes.

The **no** form of this command with the specific prefix removes the accepted source.

Parameters

ip-prefix/prefix-length

Specifies the IPv4 or IPv6 address and length.

Values

ipv4-prefix:	a.b.c.d (host bits must be 0)
ipv4-prefix-le:	0 to 32
ipv6-prefix:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6-prefix-le:	0 to 128

create

Creates a prefix instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

prefix

Syntax

[**no**] **prefix** *ip-prefix/prefix-length*

Context

[Tree] (config>qos>match-list>ip-prefix-list prefix)

Full Context

configure qos match-list ip-prefix-list prefix

Description

This command adds an IPv4 address prefix to an existing IPv4 address prefix match list.

To add a set of unique prefixes, execute the command with all unique prefixes. The prefixes are allowed to overlap IPv4 address space.

An IPv4 prefix addition will be blocked, if resource exhaustion is detected anywhere in the system because of QoS Policies that use this IPv4 address prefix list.

The **no** form of this command deletes the specified prefix from the list.

Parameters

ip-prefix

A valid IPv4 address prefix in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255 (host bit must be 0)

prefix-length

Length of the entered IP prefix

Values 1 to 32

Platforms

7705 SAR Gen 2

prefix

Syntax

[**no**] **prefix** *ipv6-prefix/prefix-length*

Context

[Tree] (config>qos>match-list>ipv6-prefix-list prefix)

Full Context

configure qos match-list ipv6-prefix-list prefix

Description

This command adds an IPv6 address prefix to an existing IPv6 address prefix match list.

To add set of unique prefixes, execute the command with all unique prefixes. The prefixes are allowed to overlap IPv6 address space.

An IPv6 prefix addition will be blocked if resource exhaustion is detected anywhere in the system because of QoS Policies that use this IPv6 address prefix list.

The **no** form of this command deletes the specified prefix from the list.

Parameters

ipv6-prefix

Specifies the IPv6 prefix for the IP match criterion in hex digits.

Values ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

prefix-length

Specifies the IPv6 prefix length for the IPv6 address expressed as a decimal integer.

Values 1 to 128

Platforms

7705 SAR Gen 2

prefix

Syntax

[**no**] **prefix** *ip-prefix/prefix-length*

Context

[Tree] (config>filter>match-list>ip-prefix-list prefix)

Full Context

configure filter match-list ip-prefix-list prefix

Description

This command adds an IPv4 address prefix to an existing IPv4 address prefix match list.

The **no** form of this command deletes the specified prefix from the list.

Operational Notes:

To add set of unique prefixes, execute the command with all unique prefixes. The prefixes are allowed to overlap IPv4 address space.

An IPv4 prefix addition will be blocked, if resource exhaustion is detected anywhere in the system because of filter policies that use this IPv4 address prefix list.

Parameters

ip-prefix

Specifies a valid IPv4 address prefix in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255 (host bit must be 0)

prefix-length

Specifies the length of the entered IPv4 prefix.

Values 0 to 32

Platforms

7705 SAR Gen 2

prefix

Syntax

[**no**] **prefix** *ipv6-prefix/prefix-length*

Context

[Tree] (config>filter>match-list>ipv6-prefix-list prefix)

Full Context

configure filter match-list ipv6-prefix-list prefix

Description

This command adds an IPv6 address prefix to an existing IPv6 address prefix match list.

The **no** form of this command deletes the specified prefix from the list.

Operational Notes:

To add set of different prefixes, execute the command with all unique prefixes. The prefixes are allowed to overlap IPv6 address space.

An IPv6 prefix addition will be blocked, if resource exhaustion is detected anywhere in the system because of filter policies that use this IPv6 address prefix list.

Parameters

ipv6-prefix/prefix-length

Specifies an IPv6 address prefix written as hexadecimal numbers separated by colons with host bits set to 0. One string of zeros can be omitted, so 2001:db8::700:0:217A is equivalent to 2001:db8:0:0:0:700:0:217A.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0..FFFF]H

d: [0..255]D

prefix-length

Specifies the length of the entered IPv6 prefix.

Values 1 to 128

Platforms

7705 SAR Gen 2

prefix

Syntax

[**no**] **prefix** *ipv6-prefix*/*prefix-length*

Context

[Tree] (config>router>router-advert>if prefix)

Full Context

configure router router-advertisement interface prefix

Description

This command configures an IPv6 prefix in the router advertisement messages. To support multiple IPv6 prefixes, use multiple prefix statements. No prefix is advertised until explicitly configured using prefix statements.

Parameters

ipv6-prefix

The IP prefix for prefix list entry in dotted decimal notation.

Values

ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 to FFFF]H
	d:	[0 to 255]D
ipv6-prefix-length	0 to 128	

prefix-length

Specifies a route must match the most significant bits and have a prefix length.

Values 1 to 128

Platforms

7705 SAR Gen 2

prefix

Syntax

[**no**] **prefix** *ip-prefix/ip-prefix-length*

Context

[Tree] (config>router>bgp>group>dynamic-neighbor>match prefix)

Full Context

configure router bgp group dynamic-neighbor match prefix

Description

This command configures a prefix to accept dynamic BGP sessions (sessions from source IP addresses not matching any configured neighbor addresses). A dynamic session is associated with the group having the longest match prefix entry for the source IP address of the peer. The group association determines local parameters that apply to the session, including the local AS, the local IP address, the MP-BGP families, the import and export policies, and so on.

The **no** form of this command removes a prefix entry.

Parameters

ip-prefix/ip-prefix-length

Specifies a prefix from which to accept dynamic BGP sessions.

Values *ipv4-prefix* — a.b.c.d (host bits must be 0)

ipv4-prefix-length — 0 to 32

ipv6-prefix — x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x — [0 to FFFF]H

d — [0 to 255]D

ipv6-prefix-length — 0 to 128

Platforms

7705 SAR Gen 2

prefix

Syntax

[**no**] **prefix** *ip-prefix/prefix-length* [**exact** | **longer** | **through** *length* | **prefix-length-range** *length1-length2* | **to** *ip-prefix/prefix-length* | **address-mask** *mask-pattern*]

Context

[Tree] (config>router>policy-options>prefix-list prefix)

Full Context

configure router policy-options prefix-list prefix

Description

This command creates a prefix entry in the route policy prefix list.

The **no** form of this command deletes the prefix entry from the prefix list.

Parameters

ip-prefix/prefix-length

Specifies the IP prefix and length for the prefix list entry in dotted decimal notation.

Values ipv4-prefix:

- a.b.c.d (host bits must be 0)

ipv4-prefix-length: [0 to 32]

ipv6-prefix:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

ipv6-prefix-length: [0 to 128]

exact

Specifies the prefix list entry only matches the route with the specified *ip-prefix* and prefix *mask* (length) values.

longer

Specifies the prefix list entry matches any route that matches the specified *ip-prefix* and prefix *mask* length values equal to or greater than the specified mask.

through *length*

Specifies the prefix list entry matches any route that matches the specified ip-prefix and has a prefix length between the specified *length* values inclusive.

Values 0 to 32

prefix-length-range *length1 - length2*

Specifies a route must match the most significant bits and have a prefix length with the given range. The range is inclusive of start and end values.

Values 0 to 32, *length2* > *length1*

to *ip-prefix/prefix-length*

Specifies a second IP prefix and length used in route policy prefix lists. A route matches prefix1 to prefix2 if it matches prefix1 and prefix2 according to their respective prefix lengths and if the route's own prefix length is between the prefix lengths of prefix1 and prefix2. It could take many individual 'exact' match prefix entries to reproduce the same logic.

mask-pattern

Specifies the address mask to use for matching entries to this prefix entry. A route matches a prefix and address mask combination if the bitwise logical AND of this prefix and the

mask equals the bitwise logical AND of the route's address and the same mask and, additionally, the prefix length of the route matches the prefix length of the prefix entry.

Values ipv4-address:

- a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

prefix

Syntax

prefix *ip-prefix/prefix-length*

no prefix

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls>sr-ospf prefix)

[Tree] (config>oam-pm>session>ip>tunnel>mpls>sr-isis prefix)

Full Context

configure oam-pm session ip tunnel mpls sr-ospf prefix

configure oam-pm session ip tunnel mpls sr-isis prefix

Description

This command configures the IP prefix used with the IGP instance to tunnel IP packets for the session tests.

The **no** form of this command deletes the prefix from the configuration.

Default

no prefix

Parameters

ip-prefix/prefix-length

Specifies an IPv4 or IPv6 address prefix.

Values ipv4-prefix: a.b.c.d (host bits must be 0)

ipv4-prefix-le:	0 to 32
ipv6-prefix:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6-prefix-le:	0 to 128

Platforms

7705 SAR Gen 2

23.5 prefix-attributes-tlv

prefix-attributes-tlv

Syntax

[**no**] **prefix-attributes-tlv**

Context

[Tree] (config>service>vprn>isis prefix-attributes-tlv)

Full Context

configure service vprn isis prefix-attributes-tlv

Description

This command enables IS-IS Prefix Attributes TLV support to exchange extended IPv4 and IPv6 reachability information. Extended reachability information is required for traffic engineering features using path computation element (PCE) or optimal route reflection.

The **no** form of this command removes the **prefix-attributes-tlv** configuration.

Default

no prefix-attributes-tlv

Platforms

7705 SAR Gen 2

prefix-attributes-tlv

Syntax

[**no**] **prefix-attributes-tlv**

Context

[Tree] (config>router>isis prefix-attributes-tlv)

Full Context

configure router isis prefix-attributes-tlv

Description

This command enables IS-IS Prefix Attributes TLV support to exchange extended IPv4 and IPv6 reachability information. Extended reachability information is required for traffic engineering features using path computation element (PCE) or optimal route reflection.

The **no** form of this command removes the **prefix-attributes-tlv** configuration.

Default

no prefix-attributes-tlv

Platforms

7705 SAR Gen 2

23.6 prefix-exclude

prefix-exclude

Syntax

prefix-exclude *policy-name* [*policy-name*]

no prefix-exclude

Context

[Tree] (config>router>ldp>aggregate-prefix-match prefix-exclude)

Full Context

configure router ldp aggregate-prefix-match prefix-exclude

Description

This command specifies the policy name containing the prefixes to be excluded from the aggregate prefix match procedures. In this case, LDP will perform an exact match of a specific FEC element prefix as

opposed to a longest match of one or more LDP FEC element prefixes, against this prefix when it receives a FEC-label binding or when a change to this prefix occurs in the routing table.

The **no** form of this command removes all policies from the configuration.

Default

no prefix-exclude

Parameters

policy-name

Specifies the route policy name, up to five. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

Platforms

7705 SAR Gen 2

prefix-exclude

Syntax

[**no**] **prefix-exclude** *ip-prefix/prefix-length*

Context

[Tree] (config>filter>match-list>ip-pfx-list prefix-exclude)

Full Context

configure filter match-list ip-prefix-list prefix-exclude

Description

This command excludes IPv4 prefix(es) from an **ip-prefix-list.** The **prefix-exclude** command is mutually exclusive with **apply-path**.

The **no** form of this command deletes the specified excluded prefixes from the **ip-prefix-list**.

Parameters

ip-prefix

Specifies a valid IPv4 address prefix in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255 (host bit must be 0)

prefix-length

Specifies the length of the entered IPv4 prefix.

Values 0 to 32

Platforms

7705 SAR Gen 2

prefix-exclude

Syntax

[**no**] **prefix** *ipv6-prefix/prefix-length*

Context

[Tree] (config>filter>match-list>ipv6-pfx-list prefix-exclude)

Full Context

configure filter match-list ipv6-prefix-list prefix-exclude

Description

This command excludes IPv6 prefix(es) from an **ipv6-prefix-list**.The **prefix-exclude** command is mutually exclusive with **apply-path**.

The **no** form of this command deletes the specified excluded prefixes from the **ipv6-prefix-list**.

Parameters

ipv6-prefix/prefix-length

Specifies an IPv6 address prefix written as hexadecimal numbers separated by colons with host bits set to 0. One string of zeros can be omitted, so 2001:db8::700:0:217A is equivalent to 2001:db8:0:0:0:700:0:217A.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0..FFFF]H

d: [0..255]D

prefix-length

Specifies the length of the entered IPv6 prefix.

Values 1 to 128

Platforms

7705 SAR Gen 2

23.7 prefix-ipv4

prefix-ipv4

Syntax

prefix-ipv4 {**enable** | **disable**}

Context

[Tree] (config>router>ldp>if-params>if>ipv6>fec-type-capability prefix-ipv4)

[Tree] (config>router>ldp>if-params>if>ipv4>fec-type-capability prefix-ipv4)

[Tree] (config>router>ldp>session-params>peer>fec-type-capability prefix-ipv4)

Full Context

configure router ldp interface-parameters interface ipv6 fec-type-capability prefix-ipv4

configure router ldp interface-parameters interface ipv4 fec-type-capability prefix-ipv4

configure router ldp session-parameters peer fec-type-capability prefix-ipv4

Description

This command enables or disables IPv4 prefix FEC capability on the session or interface.

Platforms

7705 SAR Gen 2

23.8 prefix-ipv6

prefix-ipv6

Syntax

prefix-ipv6 {**enable** | **disable**}

Context

[Tree] (config>router>ldp>if-params>if>ipv6>fec-type-capability prefix-ipv6)

[Tree] (config>router>ldp>if-params>if>ipv4>fec-type-capability prefix-ipv6)

[Tree] (config>router>ldp>session-params>peer>fec-type-capability prefix-ipv6)

Full Context

configure router ldp interface-parameters interface ipv6 fec-type-capability prefix-ipv6

configure router ldp interface-parameters interface ipv4 fec-type-capability prefix-ipv6

configure router ldp session-parameters peer fec-type-capability prefix-ipv6

Description

This command enables or disables IPv6 prefix FEC capability on the session or interface.

Platforms

7705 SAR Gen 2

23.9 prefix-limit

prefix-limit

Syntax

prefix-limit *family limit* [**threshold** *percentage*] [**idle-timeout** {*minutes* | **forever**} | **log-only** | **hold-excess** *percentage*] [**post-import**]

no prefix-limit *family*

Context

[Tree] (config>service>vprn>bgp>group>neighbor prefix-limit)

[Tree] (config>service>vprn>bgp>group prefix-limit)

Full Context

configure service vprn bgp group neighbor prefix-limit

configure service vprn bgp group prefix-limit

Description

This command configures the maximum number of BGP routes received from a peer before administrative action is taken. The administrative action can include generating a log or taking the session down. If a session is taken down, configure the **idle-timeout** parameter to bring it back up automatically after a specific duration. Alternatively, it can be configured to stay down indefinitely, until the user performs a reset.

No prefix limits for any address family are configured by default.

This command allows the user to apply a separate limit to each address family. A set of address family limits can be applied to one neighbor or to all neighbors in a group.

The **no** form of this command removes the **prefix-limit**.

Parameters

threshold *percentage*

Specifies the threshold value (as a percentage) that triggers a warning message to be sent.

Values 1 to 100

family

Specifies the address family to which the limit applies.

Values ipv4, label-ipv4, ipv6, mcast-ipv4, flow-ipv4, flow-ipv6, mcast-ipv6

limit

Specifies the number of routes that can be learned from a peer expressed as a decimal integer.

Values 1 to 4294967295

idle-timeout *minutes*

Specifies the duration in minutes before automatically re-establishing a session.

Values 1 to 1024

idle-timeout forever

Specifies that the session is re-established only after the **clear router bgp** command is executed.

log-only

Enables the warning message to be sent at the specified threshold percentage, and also when the limit is reached. However, the BGP session is not taken down.

post-import

Specifies that the limit should be applied only to the number of routes that are accepted by import policies.

hold-excess *percentage*

Specifies the percentage of maximum routes that are allowed to be installed in the route table. If a peer within scope of the configuration exceeds the limit, the overflow routes are held in the BGP RIB as inactive routes and are ineligible for forwarding or advertisement to other peers. If the **post-import** parameter is configured, only routes not rejected by import policies count toward the limit. A BGP route in the overflow state is reconsidered for activation and reinstallation when an UPDATE message is received for the route. This parameter is mutually exclusive with the **idle-timeout** and **log-only** parameters.

Platforms

7705 SAR Gen 2

prefix-limit

Syntax

prefix-limit *limit* [**log-only**] [**threshold** *percent*] [**overload-timeout** { *seconds* | **forever**}]

no prefix-limit

Context

[Tree] (config>service>vprn>isis prefix-limit)

Full Context

configure service vprn isis prefix-limit

Description

This command configures the maximum number of prefixes that IS-IS can learn, and use to protect the system from a router that has accidentally advertised a large number of prefixes. If the number of prefixes reaches the configured percentage of this limit, an SNMP trap is sent. If the limit is exceeded, IS-IS will go into overload.

The **overload-timeout** option controls the length of time that IS-IS is in the overload state when the prefix limit is reached. The system automatically attempts to restart IS-IS at the end of this duration. If the **overload-timeout forever** option is used, IS-IS is not restarted automatically and stays in overload until the condition is manually cleared by the administrator. This is also the default behavior when the **overload-timeout** option is not configured.

The **no** form of this command removes the **prefix-limit**.

Default

prefix-limit overload-timeout forever

Parameters

limit

Specifies the number of prefixes that can be learned, expressed as a decimal integer.

Values 1 to 4294967296

log-only

Enables a warning message to be sent at the specified threshold percentage and also when the limit is exceeded. However, overload is not set when this parameter is configured.

percent

Specifies the threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

overload-timeout

Keyword used to control the length of time that IS-IS is in the overload state when the prefix limit is reached.

seconds

Specifies the time in minutes before IS-IS is restarted.

Values 1 to 1800

forever

Specifies that IS-IS should be restarted only after the execution of the **clear router isis overload prefix-limit** command.

Platforms

7705 SAR Gen 2

prefix-limit

Syntax

prefix-limit *family limit* [**threshold** *percentage*] [**idle-timeout** {*minutes* | **forever**} | **log-only** | **hold-excess** *percentage*] [**post-import**]

no prefix-limit *family*

Context

[Tree] (config>router>bgp>group prefix-limit)

[Tree] (config>router>bgp>group>neighbor prefix-limit)

Full Context

configure router bgp group prefix-limit

configure router bgp group neighbor prefix-limit

Description

This command configures the maximum number of BGP routes received from a peer before administrative action is taken. The administrative action can include generating a log or taking the session down. If a session is taken down, configure the **idle-timeout** parameter to bring it back up automatically after a specific duration. Alternatively, it can be configured to stay down indefinitely, until the user performs a reset.

No prefix limits for any address family are configured by default.

This command allows the user to apply a separate limit to each address family. A set of address family limits can be applied to one neighbor or to all neighbors in a group.

The **no** form of this command removes the **prefix-limit**.

Parameters

log-only

Enables the warning message to be sent at the specified threshold percentage, and also when the limit is reached. However, the BGP session is not taken down.

threshold *percentage*

Specifies the threshold value (as a percentage) that triggers a warning message to be sent.

Values 1 to 100

family

Specifies the address family to which the limit applies.

Values ipv4, label-ipv4, vpn-ipv4, ipv6, label-ipv6, vpn-ipv6, mcast-ipv4, l2-vpn, mvpn-ipv4, mdt-safi, ms-pw, flow-ipv4, route-target, mcast-vpn-ipv4, mvpn-ipv6, flow-ipv6, evpn, mcast-ipv6, bgp-ls, sr-policy-ipv4, sr-policy-ipv6, mcast-vpn-ipv6, flow-vpn-ipv4, flow-vpn-ipv6

limit

Specifies the number of routes that can be learned from a peer expressed as a decimal integer.

Values 1 to 4294967295

idle-timeout *minutes*

Specifies the duration in minutes before automatically re-establishing a session.

Values 1 to 1024

idle-timeout forever

Specifies that the session is re-established only after the **clear router bgp** command is executed.

post-import

Specifies that the limit applies only to the number of routes that are accepted by import policies.

hold-excess *percentage*

Specifies the percentage of maximum routes that are allowed to be installed in the route table. If a peer within scope of the configuration exceeds the limit, the overflow routes are held in the BGP RIB as inactive routes and are ineligible for forwarding or advertisement to other peers. If the **post-import** parameter is configured, only routes not rejected by import policies count toward the limit. A BGP route in the overflow state is reconsidered for activation and reinstallation when an UPDATE message is received for the route. This parameter is mutually exclusive with the **idle-timeout** and **log-only** parameters.

Platforms

7705 SAR Gen 2

prefix-limit

Syntax

prefix-limit *limit* [**log-only**] [**threshold** *percent*] [**overload-timeout** { *seconds* | **forever**}]

no prefix-limit

Context

[Tree] (config>router>isis prefix-limit)

Full Context

configure router isis prefix-limit

Description

This command configures the maximum number of prefixes that IS-IS can learn, and use to protect the system from a router that has accidentally advertised a large number of prefixes. If the number of prefixes reaches the configured percentage of this limit, an SNMP trap is sent. If the limit is exceeded, IS-IS will go into overload.

The **overload-timeout** option controls the length of time that IS-IS is in the overload state when the **prefix-limit** is reached. The system automatically attempts to restart IS-IS at the end of this duration. If the **overload-timeout forever** option is used, IS-IS is not restarted automatically and stays in overload until the condition is manually cleared by the administrator. This is also the default behavior when the **overload-timeout** option is not configured.

The **no** form of this command removes the **prefix-limit**.

Default

no prefix-limit

Parameters

log-only

Enables a warning message to be sent at the specified threshold percentage and also when the limit is exceeded. However, overload is not set when this parameter is configured.

limit

Specifies the number of prefixes that can be learned expressed as a decimal integer.

Values 1 to 4294967296

percent

Specifies the threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

seconds

Specifies the time in minutes before IS-IS is restarted.

Values 1 to 1800

forever

Specifies that IS-IS should be restarted only after the execution of the **clear router isis overload prefix-limit** command.

Platforms

7705 SAR Gen 2

23.10 prefix-list

prefix-list

Syntax

prefix-list *prefix-list-name* [{**all** | **none** | **any**}] [**router-instance** *router-instance-name*]

no prefix-list [*prefix-list-name*] [{**all** | **none** | **any**}] [**router-instance** *router-instance-name*]

Context

[Tree] (config>service>vprn>static-route-entry>indirect prefix-list)

[Tree] (config>service>vprn>static-route-entry>black-hole prefix-list)

[Tree] (config>service>vprn>static-route-entry>next-hop prefix-list)

Full Context

configure service vprn static-route-entry indirect prefix-list

configure service vprn static-route-entry black-hole prefix-list

configure service vprn static-route-entry next-hop prefix-list

Description

This command associates a constraint to the associated static route such that the static route is only active if **any**, **none**, or **all** of the routes in the prefix list are present and active in the route table.

If the conditional static route is configured in a VPRN and the *router-instance-name* is configured as “Base”, the activation of the static route is dependent on the existence of routes in the Base router; the prefix-list and flag are evaluated in this context.

No router instance is specified by default, and the conditional static route is dependent on the existence of routes in the same router instance as the static route itself, subject to the details of the prefix list and the flag setting.

Entries in a referenced prefix list that are not match type ‘exact’ are interpreted as though they are ‘exact’.

The **no** form of this command disables these constraints on the static route.

Default

no prefix-list

Parameters

prefix-list-name

Specifies the name of a currently configured prefix list.

all

Specifies that the static route condition is met if all prefixes in the prefix list are present in the active static route.

none

Specifies that the static route condition is met if none of the prefixes in the named prefix-list are present in the active static route.

any

Specifies that the static route condition is met if any prefixes in the prefix list are present in the active static route.

router-instance-name

Specifies the name of the router instance. Must be "Base".

Platforms

7705 SAR Gen 2

prefix-list

Syntax

prefix-list *prefix-list-name* [{**all** | **none**}]

no prefix-list [*prefix-list-name*] [{**all** | **none**}]

Context

[Tree] (config>router>static-route-entry>black-hole prefix-list)

[Tree] (config>router>static-route-entry>next-hop prefix-list)

[Tree] (config>router>static-route-entry>indirect prefix-list)

Full Context

configure router static-route-entry black-hole prefix-list

configure router static-route-entry next-hop prefix-list

configure router static-route-entry indirect prefix-list

Description

This command associates a new constraint to the associated static route such that the static route is only active if **none** or **all** of the routes in the prefix list are present and active in the route-table.

Default

no prefix-list

Parameters

prefix-list-name

Specifies the name of a currently configured prefix-list.

all

Specifies that the static route condition is met if all prefixes in the prefix-list must be present in the active route-table.

none

Specifies that the static route condition is met if none of the prefixes in the named prefix-list can be present in the active route-table.

Platforms

7705 SAR Gen 2

prefix-list

Syntax

[**no**] **prefix-list** *name*

Context

[Tree] (config>router>policy-options prefix-list)

Full Context

configure router policy-options prefix-list

Description

This command creates a context to configure a prefix list to use in route policy entries.

The **no** form of this command deletes the named prefix list.

Parameters

name

Specifies the prefix list name. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@," "start@variable@end"," @variable@end", or "start@variable@".

An empty prefix list can be configured for pre-provisioning. This empty prefix list will not find a match when referred to by a policy. When removing member prefixes from a prefix list, the prefix list will not be automatically removed when the last member is removed. If required, an empty prefix list must be explicitly removed using the **no** form of this command.

Platforms

7705 SAR Gen 2

prefix-list

Syntax

prefix-list *name* [*name*]

no prefix-list

Context

[Tree] (config>router>policy-options>policy-statement>entry>to prefix-list)

[Tree] (config>router>policy-options>policy-statement>entry>from prefix-list)

Full Context

configure router policy-options policy-statement entry to prefix-list

configure router policy-options policy-statement entry from prefix-list

Description

This command configures a prefix list as a match criterion for a route policy statement entry.

If no prefix list is specified, any network prefix is considered a match.

An empty prefix list will evaluate as if 'no match' was found.

The prefix lists specify the network prefix (this includes the prefix and length) a specific policy entry applies.

A maximum of 28 prefix names can be specified.

The **no** form of this command removes the prefix list match criterion.

Default

no prefix-list

Parameters

name

Specifies the prefix list name. Allowed values are any string up to 64 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@," "start@variable@end"," @variable@end", or "start@variable@".

Platforms

7705 SAR Gen 2

23.11 prefix-list-override

prefix-list-override

Syntax

no prefix-list-override

prefix-list-override *name* {**exact** | **longer**}

prefix-list-override *name* **prefix-length-range** *length1-length2*

prefix-list-override *name* **through** *length*

Context

[Tree] (config>router>policy-options>policy-statement>entry>from prefix-list-override)

Full Context

configure router policy-options policy-statement entry from prefix-list-override

Description

This command converts a prefix list to a specific match type. The routing policy uses the converted list as a match condition.

The prefix list to be converted can be specified by its name, as an expression containing the name of a global variable that holds the name of the prefix list, or as an expression containing the name of a subroutine variable that holds the name of the prefix list.

Parameters

name

Specifies the prefix list to be converted, up to 64 characters.

exact

Keyword to convert all entries in the specified prefix list to the exact match type.

longer

Keyword to convert all entries in the specified prefix list to the longer match type.

length1-length2

Specifies the start and end length of the prefix range.

Values 0 to 128

Note: *length2* (end length) must be equal to or greater than *length1* (start length).

length

Specifies the through length of the prefix.

Values 0 to 128

Note: Configure the through length of the prefix to a value higher than the prefix length configured in the **configure router policy-options prefix-list prefix** command.

Platforms

7705 SAR Gen 2

23.12 prefix-map

prefix-map

Syntax

prefix-map *ip-prefix/length* **subscriber-type** *nat-sub-type* **nat-policy** *nat-policy-name* [**create**]

prefix-map *ip-prefix/length* **subscriber-type** *nat-sub-type*

no prefix-map *ip-prefix/length* **subscriber-type** *nat-sub-type*

Context

[Tree] (config>service>vprn>nat>inside>deterministic prefix-map)

[Tree] (config>router>nat>inside>deterministic prefix-map)

Full Context

configure service vprn nat inside deterministic prefix-map

configure router nat inside deterministic prefix-map

Description

This command is applicable to deterministic NAT and static 1:1 NAT. It is used to configure source IP prefixes on the inside and their association with outside deterministic NAT pools via the NAT policy. Hosts within the source IP prefix are deterministically mapped to outside IP addresses and port ranges in the associated deterministic NAT pool.

Multiple source IP prefixes within an inside routing instance can be defined and they can reference different NAT policies (and therefore, outside deterministic NAT pools and routing instances). Source IP prefixes from multiple routing instances can share the same deterministic NAT pool.

With this command, multiple NAT policies based on a destination prefix or filter criteria can be used together with deterministic NAT.

Non-deterministic NAT can be used simultaneously with deterministic NAT within the same inside routing instance. However, they cannot share the same NAT pool.

Source IP prefixes can be added or removed as long as the associated deterministic NAT pool is in a **no shutdown** mode.

Removing a prefix or modifying the map statement under it requires that the source IP prefix be in a **shutdown** mode.

Parameters

ip-prefix/length

Specifies source IP prefix on the inside whose hosts is deterministically mapped to an outside IP address and port block in the corresponding deterministic NAT pool.

Values

<ip-prefix/ip-pref*> <ipv4-prefix>/<ipv4-prefix-length>

	<ipv6-prefix>/<ipv6-prefix-length>
<ipv4-prefix>	a.b.c.d (host bits must be 0)
<ipv4-prefix-length>	0 to 32
<ipv6-prefix>	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D
<ipv6-prefix-length>	0 to 128

nat-sub-type

Specifies the subscriber type.

Values classic-lsn-sub: LSN44 subscriber

dslit-lsn-sub: DT-lite subscriber

nat-policy-name

Specifies a NAT policy, up to 32 characters, that points to an outside pool and outside routing instance.

create

Keyword used to create the particular prefix instance.

Platforms

7705 SAR Gen 2

23.13 prefix-policy

prefix-policy

Syntax

prefix-policy *prefix-policy* [*prefix-policy*]

no prefix-policy

Context

[Tree] (config>service>vprn>isis>loopfree-alternates>exclude prefix-policy)

Full Context

configure service vprn isis loopfree-alternates exclude prefix-policy

Description

This command specifies the name of the policy for the prefixes to exclude from the LFA SPF calculation in this ISIS instance.

The **no** form of this command deletes the exclude prefix policy.

Default

no prefix-policy

Parameters

prefix-policy prefix-policy

Specifies the name of the prefix policy, up to 32 characters. Up to five prefix policies can be specified. The specified name must have been already defined.

Platforms

7705 SAR Gen 2

prefix-policy

Syntax

[**no**] **prefix-policy** *prefix-policy* [*prefix-policy*]

Context

[Tree] (config>service>vprn>ospf3>loopfree-alternates>exclude prefix-policy)

[Tree] (config>service>vprn>ospf>loopfree-alternates>exclude prefix-policy)

Full Context

configure service vprn ospf3 loopfree-alternates exclude prefix-policy

configure service vprn ospf loopfree-alternates exclude prefix-policy

Description

This command specifies the name of the policy for the prefixes to exclude from the LFA SPF calculation in this OSPF or OSPF3 instance.

The **no** form of this command deletes the exclude prefix policy.

Default

no prefix-policy

Parameters

prefix-policy prefix-policy

Specifies the name of the prefix policy, up to 32 characters. Up to five prefix policies can be specified. The specified name must have been already defined.

Platforms

7705 SAR Gen 2

prefix-policy

Syntax

prefix-policy *prefix-policy* [*prefix-policy*]

no prefix-policy

Context

[Tree] (config>router>isis>loopfree-alternates>exclude prefix-policy)

Full Context

configure router isis loopfree-alternates exclude prefix-policy

Description

This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.

The implementation already allows the user to exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority. The prefix tag will, however, be used in the main SPF.

This command specifies the name of the policy for the prefixes to exclude from the LFA SPF calculation in this IS-IS instance.

Note:
Prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action, when not explicitly specified by the user in the prefix policy, is a "reject". Thus, regardless if the user did or did not explicitly add the statement "default-action reject" to the prefix policy, a prefix that did not match any entry in the policy will be accepted into LFA SPF.

The **no** form of this command deletes the exclude prefix policy.

Default

no prefix-policy

Parameters

prefix-policy prefix-policy

Specifies the name of the prefix policy, up to 32 characters. Up to five prefix policies can be specified. The specified name must have been already defined.

Platforms

7705 SAR Gen 2

prefix-policy

Syntax

prefix-policy *prefix-policy* [*prefix-policy*]

no prefix-policy

Context

[Tree] (config>router>ospf>loopfree-alternates>exclude prefix-policy)

[Tree] (config>router>ospf3>loopfree-alternates>exclude prefix-policy)

Full Context

configure router ospf loopfree-alternates exclude prefix-policy

configure router ospf3 loopfree-alternates exclude prefix-policy

Description

This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.

The implementation already allows the user to exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority. The prefix tag will, however, be used in the main SPF.

This command specifies the name of the policy for the prefixes to exclude from the LFA SPF calculation in this OSPF or OSPF3 instance.

Note:
Prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action, when not explicitly specified by the user in the prefix policy, is a "reject". Thus, regardless if the user did or did not explicitly add the statement "default-action reject" to the prefix policy, a prefix that did not match any entry in the policy will be accepted into LFA SPF.

The **no** form of this command deletes the exclude prefix policy.

Default

no prefix-policy

Parameters

prefix-policy

Specifies the name of the prefix policy, up to 32 characters. Up to five prefix policies can be specified. The specified name must have been already defined.

Platforms

7705 SAR Gen 2

23.14 prefix-sid-range

prefix-sid-range

Syntax

prefix-sid-range global

prefix-sid-range start-label *start-label* **max-index** *max-index*

no prefix-sid-range

Context

[Tree] (config>router>bgp>segment-routing prefix-sid-range)

Full Context

configure router bgp segment-routing prefix-sid-range

Description

This command configures the label block that BGP segment routing is allowed to use.

The s**tart-label** and **max-index** parameters specify that BGP should be restricted to a subrange of the SRGB, with the subrange starting at **start-label** and ending at **max-index**.

It is not possible to enable segment routing (perform a **no shutdown**) unless the **prefix-sid-range** is configured using the **global** keyword or using the **start-label** and **max-index** parameters.

The **no** form of the command allocates no labels for BGP segment-routing.

Default

no prefix-sid-range

Parameters

global

Specifies that BGP is allowed to allocate labels from the entire space of the SRGB, as defined under **config**>**router**>**mpls-labels**>**sr-labels**.

start-label

Specifies the first label value that is available to BGP in a contiguous range of labels.

Values 0 to 524287

max-index

Specifies the last label value that is available to BGP in a contiguous range of labels.

Values 1 to 524287

Platforms

7705 SAR Gen 2

prefix-sid-range

Syntax

prefix-sid-range {**global** | **start-label** *label-value* **max-index** *index-value*}

no prefix-sid-range

Context

[Tree] (config>router>isis>segment-routing prefix-sid-range)

Full Context

configure router isis segment-routing prefix-sid-range

Description

This command configures the prefix SID index range and offset label value for a given IGP instance.

The key parameter is the configuration of the prefix SID index range and the offset label value which this IGP instance will use. Since each prefix SID represents a network global IP address, the SID index for a prefix must be network-wide unique. Thus, all routers in the network are expected to configure and advertise the same prefix SID index range for a given IGP instance. However, the label value used by each router to represent this prefix; that is, the label programmed in the ILM can be local to that router by the use of an offset label, referred to as a start label:

Local Label (Prefix SID) = start-label + {SID index}

The label operation in the network becomes thus very similar to LDP when operating in the independent label distribution mode (RFC 5036, *LDP Specification*) with the difference that the label value used to forward a packet to each downstream router is computed by the upstream router based on advertised prefix SID index using the above formula.

There are two mutually exclusive modes of operation for the prefix SID range on the router. In the **global** mode of operation, the user configures the global value and this IGP instance will assume the start label value is the lowest label value in the SRGB and the prefix SID index range size equal to the range size of the SRGB. Once one IGP instance selected the global option for the prefix SID range, all IGP instances on the system will be restricted to do the same. The user must shutdown the segment routing context and delete the **prefix-sid-range** command in all IGP instances in order to change the SRGB. Once the SRGB is changed, the user must re-enter the **prefix-sid-range** command again. The SRGB range change will be failed if an already allocated SID index/label goes out of range.

In the per-instance mode of operation, the user partitions the SRGB into non-overlapping sub-ranges among the IGP instances. The user thus configures a subset of the SRGB by specifying the start label value and the prefix SID index range size. All resulting net label values (start-label + index} must be within the SRGB or the configuration will be failed. Furthermore, the code checks for overlaps of the resulting net label value range across IGP instances and will strictly enforce that these ranges do not overlap. The user must shutdown the segment routing context of an IGP instance in order to change the SID index/label range of that IGP instance using the **prefix-sid-range** command. In addition, any range change will be failed if an already allocated SID index/label goes out of range. The user can however change the SRGB on the fly as long as it does not reduce the current per IGP instance SID index/label range defined with the

prefix-sid-range. Otherwise, the user must shutdown the segment routing context of the IGP instance and delete and re-configure the **prefix-sid-range** command.

Default

no prefix-sid-range

Parameters

label-value

Specifies the label offset for the SR label range of this IGP instance.

Values 0 to 524287

index-value

Specifies the maximum value of the prefix SID index range for this IGP instance.

Values 1 to 524287

Platforms

7705 SAR Gen 2

prefix-sid-range

Syntax

prefix-sid-range global

prefix-sid-range start-label *label-value* **max-index** *index-value*

no prefix-sid-range

Context

[Tree] (config>router>ospf>segm-rtng prefix-sid-range)

Full Context

configure router ospf segment-routing prefix-sid-range

Description

This command configures the prefix SID index range and offset label value for an IGP instance.

The key parameter is the configuration of the prefix SID index range and the offset label value that this IGP instance will use. Because each prefix SID represents a network global IP address, the SID index for a prefix must be unique network-wide. Therefore, all routers in the network are expected to configure and advertise the same prefix SID index range for an IGP instance. However, the label value used by each router to represent this prefix, that is, the label programmed in the ILM, can be local to that router by the use of an offset label, referred to as a start label:

Local Label (Prefix SID) = start-label + {SID index}

The label operation in the network is very similar to LDP when operating in independent label distribution mode (RFC 5036, *LDP Specification*), with the difference being that the label value used to forward a

packet to each downstream router is computed by the upstream router based on the advertised prefix SID index using the above formula.

There are two mutually exclusive modes of operation for the prefix SID range on the router. In the **global** mode of operation, the user configures the global value and this IGP instance will assume the start label value is the lowest label value in the SRGB and the prefix SID index range size equal to the range size of the SRGB. After one IGP instance selected the global option for the prefix SID range, all IGP instances on the system will be restricted to do the same. The user must shutdown the segment routing context and delete the **prefix-sid-range** command in all IGP instances in order to change the SRGB. After the SRGB is changed, the user must re-enter the **prefix-sid-range** command again. The SRGB range change will be failed if an already allocated SID index/label goes out of range.

In per-instance mode, the user partitions the SRGB into non-overlapping sub-ranges among the IGP instances. The user configures a subset of the SRGB by specifying the start label value and the prefix SID index range size. All resulting net label values (start-label + index) must be within the SRGB or the configuration will fail. The 7705 SAR Gen 2 checks for overlaps of the resulting net label value range across IGP instances and will strictly enforce no overlapping of these ranges. The user must shut down the segment routing context of an IGP instance in order to change the SID index/label range of that IGP instance using the **prefix-sid-range** command. A range change will fail if an already allocated SID index/label goes out of range. The user can change the SRGB without shutting down the segment routing context as long as it does not reduce the current per-IGP instance SID index/label range defined with the **prefix-sid-range** command. Otherwise, shut down the segment routing context of the IGP instance, and disable and re-enable the **prefix-sid-range** command.

Default

no prefix-sid-range

Parameters

label-value

Specifies the label offset for the SR label range of this IGP instance.

Values 0 to 524287

index-value

Specifies the maximum value of the prefix SID index range for this IGP.

Values 1 to 524287

Platforms

7705 SAR Gen 2

23.15 prefix-sids

prefix-sids

Syntax

prefix-sids *ip-int-name*

no prefix-sids *ip-int-name*

Context

[Tree] (config>router>segment-routing>sr-mpls prefix-sids)

Full Context

configure router segment-routing sr-mpls prefix-sids

Description

This command configures the prefix SIDs for an interface.

The **no** form of this command removes the prefix SIDs list instance.

Default

no prefix-sids

Parameters

ip-int-name

Specifies the loopback or system interface name that owns the prefix to be advertised, up to 32 characters.

Platforms

7705 SAR Gen 2

23.16 prefix-unreachable

prefix-unreachable

Syntax

prefix-unreachable

Context

[Tree] (config>router>isis prefix-unreachable)

Full Context

configure router isis prefix-unreachable

Description

Commands in this context configure the prefix-unreachable context.

Platforms

7705 SAR Gen 2

23.17 preserve-key

preserve-key

Syntax

[**no**] **preserve-key**

Context

[Tree] (config>system>security>ssh preserve-key)

Full Context

configure system security ssh preserve-key

Description

After enabling this command, private keys, public keys, and host key file are saved by the server. It is restored following a system reboot or the ssh server restart.

The **no** form of this command specifies that the keys are held in memory by an SSH server and is not restored following a system reboot.

Default

no preserve-key

Platforms

7705 SAR Gen 2

23.18 primary

primary

Syntax

primary *path-name*

no primary

Context

[Tree] (config>router>mpls>lsp primary)

Full Context

configure router mpls lsp primary

Description

This command specifies a preferred path for the LSP. This command is optional only if the **secondary** *path-name* is included in the LSP definition. Only one primary path can be defined for an LSP.

Some of the attributes of the LSP such as the bandwidth, and hop-limit can be optionally specified as the attributes of the primary path. The attributes specified in the **primary path** *path-name* command, override the LSP attributes.

The **no** form of this command deletes the association of this *path-name* from the LSP *lsp-name*. All configurations specific to this primary path, such as record, bandwidth, and hop limit, are deleted. The primary path must be shutdown first in order to delete it. The **no primary** command will not result in any action except a warning message on the console indicating that the primary path is administratively up.

Parameters

path-name

Specifies the case-sensitive alphanumeric name label for the LSP path up to 64 characters in length.

Platforms

7705 SAR Gen 2

primary

Syntax

primary *mda-id*

no primary

Context

[Tree] (config>isa>tunnel-grp primary)

Full Context

configure isa tunnel-group primary

Description

This command assigns an ISA IPsec module configured in the specified slot to this IPsec group. The backup ISA IPsec provides the IPsec group with warm redundancy when the primary ISA IPsec in the group is configured. Primary and backup ISA IPsec have equal operational status and when both MDAs are coming up, the one that becomes operational first becomes the active ISA IPsec.

All configuration information is pushed down to the backup MDA from the CPM once the CPM gets notice that the primary module has gone down. This allows multiple IPsec groups to use the same backup module. Any statistics not yet spooled will be lost. Auto-switching from the backup to primary, once the primary becomes available again, is supported.

The operator is notified through SNMP events when:

- When the ISA IPsec service goes down (all modules in the group are down) or comes back up (a module in the group becomes active).

- When ISA IPsec redundancy fails (one of the modules in the group is down) or recovers (the failed module comes back up).
- When an ISA IPsec activity switch took place.

The **no** form of this command removes the specified primary ID from the group's configuration.

Default

no primary

Parameters

mda-id

Specifies the card/slot identifying a provisioned IPsec ISA.

Platforms

7705 SAR Gen 2

primary

Syntax

primary *primary* **secondary** *secondary*

Context

[Tree] (config>router>if>ipsec>ipsec-tun>dyn>cert>status-verify primary)

[Tree] (config>ipsec>trans-mode-prof>dyn>cert>status-verify primary)

[Tree] (config>service>vprn>if>sap>ipsec-tun>dyn>cert>status-verify primary)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>dyn>cert>status-verify primary)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>dyn>cert>status-verify primary)

[Tree] (config>service>vprn>if>sap>ipsec-gw>cert>status-verify primary)

[Tree] (config>service>ies>if>sap>ipsec-gw>cert>status-verify primary)

Full Context

configure router interface ipsec ipsec-tunnel dynamic-keying cert status-verify primary

configure ipsec ipsec-transport-mode-profile dynamic-keying cert status-verify primary

configure service vprn interface sap ipsec-tunnel dynamic-keying cert status-verify primary

configure service vprn interface ipsec ipsec-tunnel dynamic-keying cert status-verify primary

configure service ies interface ipsec ipsec-tunnel dynamic-keying cert status-verify primary

configure service vprn interface sap ipsec-gw cert status-verify primary

configure service ies interface sap ipsec-gw cert status-verify primary

Description

This command specifies the primary and secondary CVS methods used to verify the revocation status of the peer's certificate.

OCSP or CRL uses the corresponding configuration in the CA profile of the issuer of the certificate in question.

Default

primary crl

Parameters

primary

Specifies the primary CSV method used to verify the revocation status of the peer's certificate.

Values **ocsp** — Specifies that the OCSP protocol should be used. The OCSP server is configured in the corresponding CA profile.

crl — Specifies that the local CRL file should be used. The CRL file is configured in the corresponding CA profile.

Default crl

secondary

Specifies the secondary CSV method used to verify the revocation status of the peer's certificate.

Values **ocsp** — Specifies that the OCSP protocol should be used. The OCSP server is configured in the corresponding CA profile.

crl — Specifies that the local CRL file should be used. The CRL file is configured in the corresponding CA profile.

none — Specifies that no secondary method of CSV is used.

Default none

Platforms

7705 SAR Gen 2

23.19 primary-config

primary-config

Syntax

primary-config *file-url*

no primary-config

Context

[Tree] (bof primary-config)

Full Context

bof primary-config

Description

This command specifies the name and location of the primary configuration file.

The system attempts to use the configuration specified in **primary-config**. If the specified file cannot be located, the system automatically attempts to obtain the configuration from the location specified in **secondary-config** and then the **tertiary-config**.

If an error in the configuration file is encountered, the boot process aborts.

The **no** form of this command removes the **primary-config** configuration.

Parameters

file-url

Specifies the primary configuration file location, expressed as a file URL.

Values

file-url	{*local-url* | *remote-url*} (up to 180 characters)
local-url	[*cflash-id*/][*file-path*]
remote-url	[{ftp://| tftp://} *login:pswd@remote-locn*/][*file-path*]
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

23.20 primary-dns

primary-dns

Syntax

primary-dns *ip-address*

no primary-dns

Context

[Tree] (config>service>vprn>dns primary-dns)

Full Context

configure service vprn dns primary-dns

Description

This command configures the primary DNS server used for DNS name resolution. DNS name resolution can be used when executing ping, traceroute, and service-ping, and also when defining file URLs. DNS name resolution is not supported when DNS names are embedded in configuration files.

The **no** form of this command removes the primary DNS server from the configuration.

Default

no primary-dns — No primary DNS server is configured.

Parameters

ip-address

The IP or IPv6 address of the primary DNS server.

Values

ipv4-address -a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x[-interface]

x:x:x:x:x:x:d.d.d.d[-interface]

x: [0..FFFF]H

d: [0..255]D

interface - 32 characters max, for link local addresses.

Platforms

7705 SAR Gen 2

primary-dns

Syntax

primary-dns *ip-address*

no primary-dns [*ip-address*]

Context

[Tree] (bof primary-dns)

Full Context

bof primary-dns

Description

This command configures the primary DNS server used for DNS name resolution. DNS name resolution can be used when executing ping, traceroute, and service-ping, and also when defining file URLs. DNS name resolution is not supported when DNS names are embedded in configuration files.

The **no** form of this command removes the primary DNS server from the configuration.

Default

no primary-dns

Parameters

ip-address

Specifies the IP or IPv6 address of the primary DNS server.

Values

ipv4-address	*a.b.c.d*
ipv6-address	x:x:x:x:x:x:x:x[-interface]
	x:x:x:x:x:x:d.d.d.d[-interface]
	x: [0 to FFFF]H
	d: [0 to 255]D
interface	32 chars max, for link local addresses

Platforms

7705 SAR Gen 2

23.21 primary-image

primary-image

Syntax

primary-image *file-url*

no primary image

Context

[Tree] (bof primary-image)

Full Context

bof primary-image

Description

This command specifies the primary directory location for runtime image file loading.

The system attempts to load all runtime image files configured in the **primary-image** first. If this fails, the system attempts to load the runtime images from the location configured in the **secondary-image**. If the secondary image load fails, the tertiary image specified in **tertiary-image** is used.

All runtime image files (*.tim files) must be located in the same directory.

The **no** form of this command removes the **primary-image** configuration.

Parameters

file-url

Specifies the *file-url* can be either local (this CPM) or a remote FTP server.

Values

file-url	{*local-url* | *remote-url*} (up to 180 characters)
local-url	[*cflash-id*/][*file-path*]
remote-url	[{ftp://| tftp://} *login:pswd@remote-locn*/][*file-path*]
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

23.22 primary-ip-address

primary-ip-address

Syntax

primary-ip-address *ipv4-address*

no primary-ip-address

Context

[Tree] (config>router>bgp>orr>location primary-ip-address)

Full Context

configure router bgp optimal-route-reflection location primary-ip-address

Description

This command specifies the primary IP address of a reference location used for BGP optimal route reflection. Up to three IPv4 addresses and three IPv6 addresses can be specified per location.

If the TE DB is unable find a node in its topology database that matches a primary address of the location, then it tries to find a node matching a secondary address. If this attempt also fails, the TE DB tries to find a node matching a tertiary address.

The IP addresses specified for a location should be topologically "close” to a set of clients that should all receive the same optimal path for that location.

The **no** form of this command removes the primary IP address information.

Default

no primary-ip-address

Parameters

ipv4-address

Specifies the primary IPv4 address of a location expressed in dotted decimal notation.

Values a.b.c.d

Platforms

7705 SAR Gen 2

23.23 primary-ipv6-address

primary-ipv6-address

Syntax

primary-ipv6-address *ipv6-address*

no primary-ipv6-address

Context

[Tree] (config>router>bgp>orr>location primary-ipv6-address)

Full Context

configure router bgp optimal-route-reflection location primary-ipv6-address

Description

This command specifies the primary IPv6 address of a reference location used for BGP optimal route reflection. Up to three IPv4 addresses and three IPv6 addresses can be specified per location.

If the TE DB is unable find a node in its topology database that matches a primary address of the location, then it tries to find a node matching a secondary address. If this attempt also fails, the TE DB tries to find a node matching a tertiary address.

The IP addresses specified for a location should be topologically "close” to a set of clients that should all receive the same optimal path for that location.

The **no** form of this command removes the primary IPv6 address information.

Default

no primary-ipv6-address

Parameters

ipv6-address

Specifies the primary IPv6 address of a location expressed in dotted decimal notation.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

23.24 primary-next-hop

primary-next-hop

Syntax

[**no**] **primary-next-hop**

Context

[Tree] (config>router>mpls>fwd-policies>fwd-policy>nh-grp primary-next-hop)

Full Context

configure router mpls forwarding-policies forwarding-policy next-hop-group primary-next-hop

Description

Commands in this context configure the primary next hop of an NHG entry in a forwarding policy.

The **no** form of this command removes the primary next-hop context from an NHG entry in a forwarding policy.

Platforms

7705 SAR Gen 2

23.25 primary-ports

primary-ports

Syntax

primary-ports

Context

[Tree] (config>service>template>vpls-template>mac-move primary-ports)

[Tree] (config>service>vpls>mac-move primary-ports)

Full Context

configure service template vpls-template mac-move primary-ports

configure service vpls mac-move primary-ports

Description

Commands in this context define primary VPLS ports. VPLS ports that were declared as secondary prior to the execution of this command will be moved from secondary port-level to primary port-level. Changing a port to the tertiary level can only be done by first removing it from the secondary port-level.

Platforms

7705 SAR Gen 2

23.26 priority

priority

Syntax

[**no**] **priority** *level*

Context

[Tree] (config>card>fp>ingress>access>queue-group>policer-control-override>priority-mbs-thresholds priority)

[Tree] (config>card>fp>ingress>network>queue-group>policer-control-override>priority-mbs-thresholds priority)

Full Context

configure card fp ingress access queue-group policer-control-override priority-mbs-thresholds priority

configure card fp ingress network queue-group policer-control-override priority-mbs-thresholds priority

Description

The **priority** level command contains the **mbs-contribution** configuration command for a given strict priority level. Eight levels are supported numbered 1 through 8 with 8 being the highest strict priority.

Each of the eight priority CLI nodes always exists and do not need to be created. While parameters exist for each priority level, the parameters are only applied when the priority level within a parent policer instance is currently supporting child policers.

Parameters

level

Specifies the priority level.

Values 1 to 8

Platforms

7705 SAR Gen 2

priority

Syntax

[**no**] **priority** *level*

Context

[Tree] (config>service>epipe>sap>egress>policy-ctrl-over>mbs-thrshlds priority)

[Tree] (config>service>epipe>sap>ingress>policy-ctrl-over>mbs-thrshlds priority)

Full Context

configure service epipe sap egress policer-control-override priority-mbs-thresholds priority

configure service epipe sap ingress policer-control-override priority-mbs-thresholds priority

Description

The priority-level level override CLI node contains the specified priority level's mbs-contribution override value.

This node does not need to be created and will not be output in show or save configurations unless an mbs-contribution override exist for level.

Parameters

level

The level parameter is required when specifying priority-level and identifies which of the parent policer instances priority level's the mbs-contribution is overriding.

Values 1 to 8

Platforms

7705 SAR Gen 2

priority

Syntax

priority *stp-priority*

no priority [*stp-priority*]

Context

[Tree] (config>service>template>vpls-template>stp priority)

[Tree] (config>service>vpls>stp priority)

[Tree] (config>service>template>vpls-sap-template>stp priority)

Full Context

configure service template vpls-template stp priority

configure service vpls stp priority

configure service template vpls-sap-template stp priority

Description

The bridge-priority command is used to populate the priority portion of the bridge ID field within outbound BPDUs (the most significant 4 bits of the bridge ID). It is also used as part of the decision process when determining the best BPDU between messages received and sent. All values are truncated to multiples of 4096, conforming with IEEE 802.1t and 802.1D-2004.

The **no** form of this command returns the bridge priority to the default value.

Default

priority 4096

Parameters

bridge-priority

Specifies the bridge priority for the STP instance

Values Allowed values are integers in the range of 4096 to 65535 with 4096 being the highest priority. The actual bridge priority value stored/used is the number entered with the lowest 12 bits masked off which means the actual range of values is 4096 to 61440 in increments of 4096.

Platforms

7705 SAR Gen 2

priority

Syntax

priority *stp-priority*

no priority

Context

[Tree] (config>service>vpls>spoke-sdp>stp priority)

[Tree] (config>service>vpls>sap>stp priority)

Full Context

configure service vpls spoke-sdp stp priority

configure service vpls sap stp priority

Description

This command configures the Nokia Spanning Tree Protocol (STP) priority for the SAP or spoke SDP.

STP priority is a configurable parameter associated with a SAP or spoke SDP. When configuration BPDUs are received, the priority is used in some circumstances as a tie breaking mechanism to determine whether the SAP or spoke SDP be designated or blocked.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255 also) to create a 16 bit value. In the latest STP standard (802.1D-2004) only the upper 4 bits of the port priority field are used to encode the SAP or spoke SDP priority. The remaining 4 bits are used to extend the port ID field into a 12 bit virtual port number field. The virtual port number uniquely references a SAP or spoke SDP within the STP instance.

STP computes the actual priority by taking the input value and masking out the lower four bits. The result is the value that is stored in the SDP priority parameter. For instance, if a value of 0 is entered, masking out the lower 4 bits results in a parameter value of 0. If a value of 255 is entered, the result is 240.

The **no** form of this command returns the STP priority to the default value.

Default

priority 128

Parameters

stp-priority

Specifies the STP priority value for the SAP or spoke SDP. 0 is the highest priority. The actual value used for STP priority (and stored in the configuration) is the result of masking out the lower 4 bits, therefore the actual value range is 0 to 240 in increments of 16.

Values 0 to 255

Platforms

7705 SAR Gen 2

priority

Syntax

[**no**] **priority** *level*

Context

[Tree] (config>service>vpls>sap>egress>policy-ctrl-over>mbs-thrshlds priority)

[Tree] (config>service>vpls>sap>ingress>policy-ctrl-over>mbs-thrshlds priority)

Full Context

configure service vpls sap egress policer-control-override priority-mbs-thresholds priority

configure service vpls sap ingress policer-control-override priority-mbs-thresholds priority

Description

The priority-level level override CLI node contains the specified priority level's mbs-contribution override value.

This node does not need to be created and will not be output in show or save configurations unless an mbs-contribution override exist for level.

The **no** form of this command sets the MBS contribution for the associated priority to its default value.

Parameters

level

Specifies that the level parameter is required when specifying priority-level and identifies which of the parent policer instances priority level's the mbs-contribution is overriding

Values 1 to 8

Platforms

7705 SAR Gen 2

priority

Syntax

priority *base-priority*

no priority

Context

[Tree] (config>service>ies>if>ipv6>vrrp priority)

Full Context

configure service ies interface ipv6 vrrp priority

Description

This command provides the ability to configure a specific priority value to the virtual router instance. In conjunction with an optional policy command, the base-priority is used to derive the in-use priority of the virtual router instance.

This command is only available in the non-owner vrrp virtual-router-id nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed. For non-owner virtual router instances, if the priority command is not executed, the base-priority will be set to 100.

The **no** form of this command restores the default value of 100 to base-priority.

Default

priority 100

Parameters

base-priority

The base-priority parameter configures the base priority used by the virtual router instance. If a VRRP Priority Control policy is not also defined, the base-priority will be the in-use priority for the virtual router instance.

Values 1 to 254

Default 100

Platforms

7705 SAR Gen 2

priority

Syntax

[**no**] **priority** *level*

Context

[Tree] (config>service>ies>if>sap>egress>policer-ctrl-over>mbs-thrshlds priority)

[Tree] (config>service>ies>if>sap>ingress>policer-ctrl-over>mbs-thrshlds priority)

Full Context

configure service ies interface sap egress policer-control-override priority-mbs-thresholds priority

configure service ies interface sap ingress policer-control-override priority-mbs-thresholds priority

Description

The priority-level level override CLI node contains the specified priority level's mbs-contribution override value.

This node does not need to be created and will not be output in show or save configurations unless an mbs-contribution override exist for level.

The **no** form of this command sets the MBS contribution for the associated priority to its default value.

Parameters

level

Specifies that the level parameter is required when specifying priority-level and identifies which of the parent policer instances priority level's the mbs-contribution is overriding.

Values 1 to 8

Platforms

7705 SAR Gen 2

priority

Syntax

priority *base-priority*

no priority

Context

[Tree] (config>service>ies>if>vrrp priority)

Full Context

configure service ies interface vrrp priority

Description

The priority command provides the ability to configure a specific priority value to the virtual router instance. In conjunction with an optional policy command, the base-priority is used to derive the in-use priority of the virtual router instance.

The priority command is only available in the non-owner vrrp virtual-router-id nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed. For non-owner virtual router instances, if the priority command is not executed, the base-priority will be set to 100.

The **no** form of this command restores the default value of 100 to base-priority.

Parameters

base-priority

The base-priority parameter configures the base priority used by the virtual router instance. If a VRRP Priority Control policy is not also defined, the base-priority will be the in-use priority for the virtual router instance.

Values 1 to 254

Default 100

Platforms

7705 SAR Gen 2

priority

Syntax

[**no**] **priority** *level*

Context

[Tree] (config>service>vprn>if>sap>ingress>policer-ctrl-over>mbs-thrshlds priority)

[Tree] (config>service>vprn>if>sap>egress>policer-ctrl-over>mbs-thrshlds priority)

Full Context

configure service vprn interface sap ingress policer-control-override priority-mbs-thresholds priority

configure service vprn interface sap egress policer-control-override priority-mbs-thresholds priority

Description

The priority-level level override CLI node contains the specified priority level's mbs-contribution override value.

This node does not need to be created and will not be output in show or save configurations unless an mbs-contribution override exist for level.

The **no** form of this command sets the MBS contribution for the associated priority to its default value.

Parameters

level

Specifies that the level parameter is required when specifying priority-level and identifies which of the parent policer instances priority level's the mbs-contribution is overriding.

Values 1 to 8

Platforms

7705 SAR Gen 2

priority

Syntax

priority *priority*

no priority

Context

[Tree] (config>service>vprn>if>ipv6>vrrp priority)

[Tree] (config>service>vprn>if>vrrp priority)

Full Context

configure service vprn interface ipv6 vrrp priority

configure service vprn interface vrrp priority

Description

The priority command provides the ability to configure a specific priority value to the virtual router instance. In conjunction with an optional policy command, the base-priority is used to derive the in-use priority of the virtual router instance.

The priority command is only available in the non-owner **vrrp** *virtual-router-id* nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed. For non-owner virtual router instances, if the priority command is not executed, the base-priority will be set to 100.

The **no** form of this command restores the default value of 100 to base-priority.

Parameters

base-priority

The base-priority parameter configures the base priority used by the virtual router instance. If a VRRP priority control policy is not also defined, the base-priority will be the in-use priority for the virtual router instance.

Values 1 to 254

Default 100

Platforms

7705 SAR Gen 2

priority

Syntax

priority *number*

no priority

Context

[Tree] (config>service>vprn>isis>if>level priority)

Full Context

configure service vprn isis interface level priority

Description

This command configures the priority of the IS-IS router interface for designated router election on a multi-access network.

This priority is included in hello PDUs transmitted by the interface on a multi-access network. The router with the highest priority is the preferred designated router. The designated router is responsible for sending LSPs with regard to this network and the routers that are attached to it.

The **no** form of this command reverts to the default value.

Default

priority 64

Parameters

number

Specifies the priority for this interface at this level.

Values 0 to 127

Platforms

7705 SAR Gen 2

priority

Syntax

priority *number*

no priority

Context

[Tree] (config>service>vprn>ospf>area>if priority)

[Tree] (config>service>vprn>ospf3>area>if priority)

Full Context

configure service vprn ospf area interface priority

configure service vprn ospf3 area interface priority

Description

This command configures the priority of the OSPF interface that is used to elect the designated router (DR) on the subnet.

This parameter is only used if the interface is of type **broadcast**. The router with the highest priority interface becomes the DR. A router with priority 0 is not eligible to be the designated router or backup designated router.

The **no** form of this command resets the interface priority to the default value.

Default

priority 1

Parameters

number

The interface priority expressed as a decimal integer. A value of 0 indicates the router is not eligible to be the Designated Router of Backup Designated Router on the interface subnet.

Values 0 to 255

Platforms

7705 SAR Gen 2

priority

Syntax

priority *dr-priority*

no priority

Context

[Tree] (config>service>vprn>pim>if priority)

Full Context

configure service vprn pim interface priority

Description

This command sets the priority value to become the rendezvous point (RP) that is included in bootstrap messages sent by the router. The RP is sometimes called the bootstrap router. The **priority** command indicates whether the router is eligible to be a bootstrap router.

The **no** form of this command disqualifies the router to participate in the bootstrap election.

Default

priority 1 (The router is the least likely to become the designated router.)

Parameters

dr-priority

Specifies the priority to become the designated router. The higher the value, the higher the priority.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

priority

Syntax

priority *bootstrap-priority*

Context

[Tree] (config>service>vprn>pim>rp>bsr-candidate priority)

[Tree] (config>service>vprn>pim>rp>ipv6>bsr-candidate priority)

Full Context

configure service vprn pim rp bsr-candidate priority

configure service vprn pim rp ipv6 bsr-candidate priority

Description

This command defines the priority used to become the rendezvous point (RP). The higher the priority value the more likely that this router becomes the RP. If there is a tie, the router with the highest IP address is elected.

Parameters

bootstrap-priority

The priority to become the bootstrap router.

Values 0 to 255

Default 0 (the router is not eligible to be the bootstrap router)

Platforms

7705 SAR Gen 2

priority

Syntax

priority *priority*

no priority

Context

[Tree] (config>service>vprn>pim>rp>rp-candidate priority)

Full Context

configure service vprn pim rp rp-candidate priority

Description

This command defines the priority used to become the rendezvous point (RP). The higher the priority value, the more likely that this router will become the RP.

Use the **no** form of this command to revert to the default value.

Default

priority 192

Parameters

priority

Specifies the priority to become the designated router. The higher the value the more likely the router will become the RP.

Values 0 to 255

Platforms

7705 SAR Gen 2

priority

Syntax

priority *setup-priority hold-priority*

no priority

Context

[Tree] (config>router>mpls>lsp-template priority)

[Tree] (config>router>mpls>lsp>secondary priority)

[Tree] (config>router>mpls>lsp>primary priority)

Full Context

configure router mpls lsp-template priority

configure router mpls lsp secondary priority

configure router mpls lsp primary priority

Description

This command enables the soft preemption procedures for this LSP path. The operator enables the soft preemption mechanism on a specific LSP name by explicitly configuring the setup and holding priorities for the primary path at the head-end node. The operator can similarly configure priority values for a secondary path for this LSP name. Different values could be used for the primary and for any of the secondary paths. In the absence of explicit user configuration, the setup priority is internally set to the default value of 7 and the holding priority is set to the default value of 0.

Note:
Valid user-entered values for these two parameters require that the holding priority be numerically lower than or equal to the setup priority, otherwise preemption loops can occur.

Preemption is effected when a router preempting node processes a new RSVP session reservation and there is not enough available bandwidth on the RSVP interface, or the Class Type (CT) when Diff-Serv is enabled, to satisfy the bandwidth in the FlowSpec object while there exist other session reservations for LSP paths with a strictly lower holding priority (numerically higher holding priority value) than the setup priority of the new LSP reservation. If enough available bandwidth is freed on the link or CT to accommodate the new reservation by preempting one or more lower priority LSP paths, the preempting node allows temporary overbooking of the RSVP interface and honors the new reservation.

The preempting node will immediately set the 'Preemption pending' flag (0x10) in the IPv4 Sub-Object in the RRO object in the Resv refresh for each of the preempted LSP paths. The IPv4 Sub-Object corresponds to the outgoing interface being used by the preempting and preempted LSP paths; however, the bandwidth value in the FlowSpec object is not changed. The Resv flag must also be set if the preempting node is a merge point for the primary LSP path and the backup bypass LSP or detour LSP and the backup LSP is activated.

When evaluating if enough available bandwidth will be freed, the preempting node considers the reservations in order from the lowest holding priority (numerically higher holding priority value) to the holding priority just below the setup priority of the new reservation. A new reservation cannot preempt a reservation which has a value of the holding priority equal to the new reservation setup priority.

When Diff-Serv is enabled on the preempting node and the MAM bandwidth allocation model is used, a new reservation can only preempt a reservation in the same Class Type (CT).

LSP paths which were not flagged at the head-end for soft preemption will be hard preempted. LSP paths with the default holding priority of 0 cannot be preempted. LSP paths with zero bandwidth do not preempt other LSP paths regardless of the values of the path setup priority and the path holding priority. They can also not be preempted.

When evaluating if enough available bandwidth will be freed, the preempting node considers the reservations in order from the lowest holding priority (numerically higher holding priority) to the holding priority just below the setup priority of the new reservation. There is no specific order in which the reservations in the same holding priority are considered.

The preempting node starts a preemption timer for each of the preempted LSP paths. While this timer is on, the node should continue to refresh the Path and Resv for the preempted LSP paths. When the preemption timer expires, the node tears down the reservation if the head-end node has not already done so.

A head-end node upon receipt of the Resv refresh message with the 'Preemption pending' flag must immediately perform a make-before-break on the affected adaptive CSPF LSP. Both IGP metric and TE metric based CSPF LSPs are included. If an alternative path that excludes the flagged interface is not found, then the LSP is put on a retry in a similar way to the Global Revertive procedure at a head-end node. However, the number of retries and the retry timer are governed by the values of the **retry-limit** and **retry-timer** parameters: **config>router>mpls>lsp>retry-limit**; **config>router>mpls>lsp>retry-timer**.

MPLS will keep the address list of flagged interfaces for a maximum of 60 s (not user-configurable) from the time the first Resv message with the 'Preemption pending' flag is received. This actually means that MPLS will request CSPF to find a path that excludes the flagged interfaces in the first few retries until success or until 60 s have elapsed. Subsequent retries after the 60 s will not exclude the flagged interfaces as it is assumed IGP has converged by then and the Unreserved Bandwidth sub-TLV for that priority, or TE Class, in the TE database will show the updated value taking into account the preempting LSP path reservation or a value of zero if overbooked.

If the LSP has a configured secondary standby which is operationally UP, the router will switch the path of the LSP to it and then start the MBB. If no standby path is available and a secondary non-standby is configured, the router will start the MBB and signal the path of the secondary. The LSP path will be switched to either the secondary or the new primary, whichever comes up first.

The **no** form of this command reverts the LSP path priority to the default values and results in setting the setup priority to 7, in setting the hold priority to 0, and in clearing the 'soft preemption desired' flag in the RRO in the Resv refresh message.

Default

no priority

Parameters

setup-priority

Specifies the priority of the reservation for this session at setup time.

Values 0 to 7 (0 is the highest priority and 7 is the lowest priority.)

Default 7 — This session does not preempt any other session.

holding-priority

Specifies the priority of the reservation for this session at preemption action.

Values 0 to 7 (0 is the highest priority and 7 is the lowest priority.)

Default 0 — This session does not get preempted by any other session.

Platforms

7705 SAR Gen 2

priority

Syntax

priority *priority*

no priority

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ipsec>tunnel-group priority)

Full Context

configure redundancy multi-chassis peer mc-ipsec tunnel-group priority

Description

This command specifies the local priority of the tunnel-group, this is used to elect master, higher number win. If priority are same, then the peer has more active ISA win; and priority and the number of active ISA are same, then the peer with higher IP address win.

The **no** form of this command removes the priority value from the configuration.

Default

priority 100

Parameters

priority

Specifies the priority of this tunnel-group.

Values 0 to 255

Platforms

7705 SAR Gen 2

priority

Syntax

priority *dr-priority*

no priority

Context

[Tree] (config>router>pim>interface priority)

Full Context

configure router pim interface priority

Description

This command sets the priority value to elect the designated router (DR). The DR election priority is a 32-bit unsigned number and the numerically larger priority is always preferred.

The **no** form of this command reverts to the default value.

Default

priority 1

Parameters

priority

Specifies the priority to become the designated router. The higher the value, the higher the priority.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

priority

Syntax

priority *priority*

no priority

Context

[Tree] (config>router>pim>rp>ipv6>rp-candidate priority)

[Tree] (config>router>pim>rp>rp-candidate priority)

Full Context

configure router pim rp ipv6 rp-candidate priority

configure router pim rp rp-candidate priority

Description

This command configures the Candidate-RP priority for becoming a rendezvous point (RP). This value is used to elect RP for a group range.

The **no** form of this command reverts to the default value.

Default

priority 192

Parameters

priority

Specifies the priority to become a rendezvous point (RP). A value of 0 is considered as the highest priority.

Values 0 to 255

Platforms

7705 SAR Gen 2

priority

Syntax

priority *level*

Context

[Tree] (config>qos>plcr-ctrl-plcy>root>priority-mbs-thresholds priority)

Full Context

configure qos policer-control-policy root priority-mbs-thresholds priority

Description

The **priority** level command contains the **mbs-contribution** configuration command for a given strict priority level. Eight levels are supported numbered 1 through 8 with 8 being the highest strict priority.

Each of the eight priority CLI nodes always exists and do not need to be created. While parameters exist for each priority level, the parameters are only applied when the priority level within a parent policer instance is currently supporting child policers.

Platforms

7705 SAR Gen 2

priority

Syntax

priority [*priority*]

no priority

Context

[Tree] (config>filter>redirect-policy>dest priority)

Full Context

configure filter redirect-policy destination priority

Description

Redirect policies can contain multiple destinations. Each destination is assigned an initial or base **priority** which describes its relative importance within the policy.

Default

priority 100

Parameters

priority

Specifies the priority, expressed as a decimal integer, used to weigh the destination's relative importance within the policy.

Values 1 to 255

Platforms

7705 SAR Gen 2

priority

Syntax

priority *priority*

no priority

Context

[Tree] (config>router>fad>flex-algo priority)

Full Context

configure router flexible-algorithm-definitions flex-algo priority

Description

This command configures the priority of the FAD. This priority is used as a tie-breaker when the router has received multiple FADs for the same flexible algorithm.

Every router that is configured to participate in a particular flexible algorithm uses the same tie-breaker logic to select the winning FAD. This allows for consistent FAD definition selection in cases where routers advertise different definitions for a specific flexible algorithm. The following rules apply to the breaker mechanism.

- From the advertisements of the FAD in the area (including both locally generated advertisements and received advertisements), select the one with the highest priority value.
- If there are multiple advertisements of the FAD with the same highest priority, select the one that is originated from the router with either the highest system ID or router ID.

The **no** form of this command sets the priority to the default value.

Default

priority 100

Parameters

priority

Configures the priority of this FAD.

Values 0 to 255

Default 100

Platforms

7705 SAR Gen 2

priority

Syntax

priority *priority*

no priority

Context

[Tree] (config>router>if>ipv6>vrrp priority)

[Tree] (config>router>if>vrrp priority)

Full Context

configure router interface ipv6 vrrp priority

configure router interface vrrp priority

Description

This command configures the base router priority for the virtual router instance used in the master election process.

The priority is the most important parameter set on a non-owner virtual router instance. The priority defines a virtual router's selection order in the master election process. Together, the priority value and the **preempt** mode allow the virtual router with the best priority to become the master virtual router.

The *base-priority* is used to derive the in-use priority of the virtual router instance as modified by any optional VRRP priority control policy. VRRP priority control policies can be used to either override or adjust the base priority value depending on events or conditions within the chassis.

The **priority** command is only available in the non-owner **vrrp** nodal context. The priority of **owner** virtual router instances is permanently set to 255 and cannot be changed.

For non-owner virtual router instances, the default base priority value is 100.

The **no** form of the command reverts to the default value.

Default

priority 100

Parameters

priority

The base priority used by the virtual router instance expressed as a decimal integer. If no VRRP priority control policy is defined, the *base-priority* is the in-use priority for the virtual router instance.

Values 1 to 254

Platforms

7705 SAR Gen 2

priority

Syntax

priority *priority-level* [{**delta** | **explicit**}]

no priority

Context

[Tree] (config>vrrp>policy>priority-event>route-unknown priority)

[Tree] (config>vrrp>policy>priority-event>port-down priority)

[Tree] (config>vrrp>policy>priority-event>lag-port-down>number-down priority)

[Tree] (config>vrrp>policy>priority-event>lag-port-down>weight-down priority)

[Tree] (config>vrrp>policy>priority-event>host-unreachable priority)

Full Context

configure vrrp policy priority-event route-unknown priority

configure vrrp policy priority-event port-down priority

configure vrrp policy priority-event lag-port-down number-down priority

configure vrrp policy priority-event lag-port-down weight-down priority

configure vrrp policy priority-event host-unreachable priority

Description

This command controls the effect the set event has on the virtual router instance in-use priority.

When the event is set, the *priority-level* is either subtracted from the base priority of each virtual router instance or it defines the explicit in-use priority value of the virtual router instance depending on whether the **delta** or **explicit** keywords are specified.

Multiple set events in the same policy have interaction constraints:

- If any set events have an explicit **priority** value, all the delta **priority** values are ignored.
- The set event with the lowest explicit **priority** value defines the in-use priority that are used by all virtual router instances associated with the policy.
- If no set events have an explicit **priority** value, all the set events delta **priority** values are added and subtracted from the base priority value defined on each virtual router instance associated with the policy.
- If the delta priorities sum exceeds the **delta-in-use-limit** parameter, then the **delta-in-use-limit** parameter is used as the value subtracted from the base priority value defined on each virtual router instance associated with the policy.

If the **priority** command is not configured on the priority event, the *priority-value* defaults to 0 and the qualifier keyword defaults to **delta**, therefore, there is no impact on the in-use priority.

The **no** form of the command configures the set event to subtract 0 from the base priority (no effect).

Default

no priority

Parameters

priority-level

The priority level adjustment value expressed as a decimal integer.

Values 0 to 254

delta

Configures what effect the *priority-level* will have on the base priority value. The default base priority value is **delta**.

When **delta** is specified, the *priority-level* value is subtracted from the associated virtual router instance's base priority when the event is set and no explicit events are set. The sum of the priority event *priority-level* values on all set delta priority events are subtracted from the virtual router base priority to derive the virtual router instance in-use priority value. If the **delta** priority event is cleared, the *priority-level* is no longer used in the in-use priority calculation.

explicit

Configures what effect the *priority-level* will have on the base priority value.

When **explicit** is specified, the *priority-level* value is used to override the base priority of the virtual router instance if the priority event is set and no other **explicit** priority event is set with a lower *priority-level*. The set **explicit** priority value with the lowest *priority-level* determines the actual in-use protocol value for all virtual router instances associated with the policy.

Platforms

7705 SAR Gen 2

priority

Syntax

priority *priority-level* **explicit**

no priority

Context

[Tree] (config>vrrp>policy>priority-event>mc-ipsec-non-forwarding priority)

Full Context

configure vrrp policy priority-event mc-ipsec-non-forwarding priority

Description

This command controls the effect the set event has on the virtual router instance in-use priority.

When the event is set, the *priority-level* is either subtracted from the base priority of each virtual router instance or it defines the explicit in-use priority value of the virtual router instance depending on whether the **delta** or **explicit** keywords are specified.

Multiple set events in the same policy have interaction constraints:

- If any set events have an explicit **priority** value, all the delta **priority** values are ignored.
- The set event with the lowest explicit **priority** value defines the in-use priority that are used by all virtual router instances associated with the policy.
- If no set events have an explicit **priority** value, all the set events delta **priority** values are added and subtracted from the base priority value defined on each virtual router instance associated with the policy.
- If the delta priorities sum exceeds the **delta-in-use-limit** parameter, then the **delta-in-use-limit** parameter is used as the value subtracted from the base priority value defined on each virtual router instance associated with the policy.

If the **priority** command is not configured on the priority event, the *priority-value* defaults to 0 and the qualifier keyword defaults to **delta**, therefore, there is no impact on the in-use priority.

The **no** form of the command configures the set event to subtract 0 from the base priority (no effect).

Default

no priority

Parameters

priority-level

The priority level adjustment value expressed as a decimal integer.

Values 0 to 254

explicit

When **explicit** is specified, the *priority-level* value is used to override the base priority of the virtual router instance if the priority event is set and no other **explicit** priority event is set with a lower *priority-level*. The set **explicit** priority value with the lowest *priority-level* determines the actual in-use protocol value for all virtual router instances associated with the policy.

Platforms

7705 SAR Gen 2

priority

Syntax

priority *bridge-priority*

no priority

Context

[Tree] (config>service>pw-template>stp priority)

Full Context

configure service pw-template stp priority

Description

The **bridge-priority** command is used to populate the priority portion of the bridge ID field within outbound BPDUs (the most significant 4 bits of the bridge ID). It is also used as part of the decision process when determining the best BPDU between messages received and sent. All values will be truncated to multiples of 4096, conforming with IEEE 802.1t and 802.1D-2004.

The **no** form of this command returns the bridge priority to the default value.

Default

priority 4096

Parameters

bridge-priority

Specifies the bridge priority for the STP instance.

Values Allowed values are integers in the range of 4096 to 65535 with 4096 being the highest priority. The actual bridge priority value stored/used is the number entered with the lowest 12 bits masked off which means the actual range of values is 4096 to 61440 in increments of 4096.

Platforms

7705 SAR Gen 2

priority

Syntax

priority *number*

no priority

Context

[Tree] (config>router>isis>if>level priority)

Full Context

configure router isis interface level priority

Description

This command configures the priority of the IS-IS router interface for designated router election on a multi-access network.

This priority is included in hello PDUs transmitted by the interface on a multi-access network. The router with the highest priority is the preferred designated router. The designated router is responsible for sending LSPs with regard to this network and the routers that are attached to it.

The **no** form of this command reverts to the default value.

Default

priority 64

Parameters

number

Specifies the priority for this interface at this level.

Values 0 to 127

Platforms

7705 SAR Gen 2

priority

Syntax

priority *number*

no priority

Context

[Tree] (config>router>ospf>area>interface priority)

[Tree] (config>router>ospf3>area>interface priority)

Full Context

configure router ospf area interface priority

configure router ospf3 area interface priority

Description

This command configures the priority of the OSPF interface that is used in an election of the designated router on the subnet.

This parameter is only used if the interface is of type broadcast. The router with the highest priority interface becomes the designated router. A router with priority 0 is not eligible to be Designated Router or Backup Designated Router.

The **no** form of this command reverts the interface priority to the default value.

Default

priority 1

Parameters

number

Specifies the interface priority expressed as a decimal integer. A value of 0 indicates the router is not eligible to be the Designated Router or Backup Designated Router on the interface subnet.

Values 0 to 255

Platforms

7705 SAR Gen 2

priority

Syntax

priority [*value*]

no priority

Context

[Tree] (config>redundancy>multi-chassis>ipsec-domain priority)

Full Context

configure redundancy multi-chassis ipsec-domain priority

Description

This command configures the priority for the tunnel group in the IPsec domain. The node with the higher priority is more likely to be elected as active within the domain.

The **no** form of this command reverts to the default value.

Default

priority 100

Parameters

value

Specifies the IPsec domain tunnel group priority.

Platforms

7705 SAR Gen 2

23.27 priority-event

priority-event

Syntax

[**no**] **priority-event**

Context

[Tree] (config>vrrp>policy priority-event)

Full Context

configure vrrp policy priority-event

Description

This command creates the context to configure VRRP priority control events used to define criteria to modify the VRRP in-use priority.

A priority control event specifies an object to monitor and the effect on the in-use priority level for an associated virtual router instance.

Up to 32 priority control events can be configured within the **priority-event** node.

The **no** form of the command clears any configured priority events.

Platforms

7705 SAR Gen 2

23.28 priority-mbs-thresholds

priority-mbs-thresholds

Syntax

priority-mbs-thresholds

Context

[Tree] (config>card>fp>ingress>access>queue-group>policer-control-override priority-mbs-thresholds)

[Tree] (config>card>fp>ingress>network>queue-group>policer-control-override priority-mbs-thresholds)

Full Context

configure card fp ingress access queue-group policer-control-override priority-mbs-thresholds

configure card fp ingress network queue-group policer-control-override priority-mbs-thresholds

Description

This command contains the root arbiter parent policer's **min-thresh-separation** command and each priority level's **mbs-contribution** command that is used to internally derive each priority level's shared-portion and fair-portion values. The system uses each priority level's shared-portion and fair-portion value to calculate each priority level's discard-unfair and discard-all MBS thresholds that enforce priority sensitive rate-based discards within the root arbiter's parent policer.

The **priority-mbs-thresholds** CLI node always exists and does not need to be created.

Platforms

7705 SAR Gen 2

priority-mbs-thresholds

Syntax

priority-mbs-thresholds

Context

[Tree] (config>service>epipe>sap>egress>policer-control-override priority-mbs-thresholds)

[Tree] (config>service>epipe>sap>ingress>policer-control-override priority-mbs-thresholds)

Full Context

configure service epipe sap egress policer-control-override priority-mbs-thresholds

configure service epipe sap ingress policer-control-override priority-mbs-thresholds

Description

This command overrides the CLI node contains the configured min-thresh-separation and the various priority level mbs-contribution override commands.

Platforms

7705 SAR Gen 2

priority-mbs-thresholds

Syntax

priority-mbs-thresholds

Context

[Tree] (config>service>vpls>sap>ingress>policer-ctrl-over priority-mbs-thresholds)

[Tree] (config>service>vpls>sap>egress>policer-ctrl-over priority-mbs-thresholds)

Full Context

configure service vpls sap ingress policer-control-override priority-mbs-thresholds

configure service vpls sap egress policer-control-override priority-mbs-thresholds

Description

This command overrides the CLI node contains the configured min-thresh-separation and the various priority level mbs-contribution override commands.

Platforms

7705 SAR Gen 2

priority-mbs-thresholds

Syntax

priority-mbs-thresholds

Context

[Tree] (config>service>ies>if>sap>egress>policer-ctrl-over priority-mbs-thresholds)

[Tree] (config>service>ies>if>sap>ingress>policer-ctrl-over priority-mbs-thresholds)

Full Context

configure service ies interface sap egress policer-control-override priority-mbs-thresholds

configure service ies interface sap ingress policer-control-override priority-mbs-thresholds

Description

This command overrides the CLI node contains the configured min-thresh-separation and the various priority level mbs-contribution override commands.

Platforms

7705 SAR Gen 2

priority-mbs-thresholds

Syntax

priority-mbs-thresholds

Context

[Tree] (config>service>vprn>if>sap>egress>policer-ctrl-over priority-mbs-thresholds)

[Tree] (config>service>vprn>if>sap>ingress>policer-ctrl-over priority-mbs-thresholds)

Full Context

configure service vprn interface sap egress policer-control-override priority-mbs-thresholds

configure service vprn interface sap ingress policer-control-override priority-mbs-thresholds

Description

This command overrides the CLI node contains the configured min-thresh-separation and the various priority level mbs-contribution override commands.

Platforms

7705 SAR Gen 2

priority-mbs-thresholds

Syntax

priority-mbs-thresholds

Context

[Tree] (config>qos>plcr-ctrl-plcy>root priority-mbs-thresholds)

Full Context

configure qos policer-control-policy root priority-mbs-thresholds

Description

The **priority-mbs-thresholds** command contains the root arbiter parent policer's **min-thresh-separation** command and each priority level's **mbs-contribution** command that is used to internally derive each priority level's shared-portion and fair-portion values. The system uses each priority level's shared-portion and fair-portion value to calculate each priority level's discard-unfair and discard-all MBS thresholds that enforce priority-sensitive rate-based discards within the root arbiter's parent policer.

The **priority-mbs-thresholds** CLI node always exists and does not need to be created.

Platforms

7705 SAR Gen 2

23.29 priv-lvl

priv-lvl

Syntax

priv-lvl *priv-lvl user-profile-name*

no priv-lvl *priv-lvl*

Context

[Tree] (config>service>vprn>aaa>rmt-srv>tacplus>priv-lvl-map priv-lvl)

[Tree] (config>system>security>tacplus>priv-lvl-map priv-lvl)

Full Context

configure service vprn aaa remote-servers tacplus priv-lvl-map priv-lvl

configure system security tacplus priv-lvl-map priv-lvl

Description

This command maps a specific TACACS+ priv-lvl to a locally configured profile for authorization. This mapping is used when the **use-priv-lvl** option is specified for TACPLUS authorization.

Parameters

priv-lvl

Specifies the privilege level used when sending a TACACS+ ENABLE request.

Values 0 to 15

user-profile-name

Specifies the user profile for this mapping.

Platforms

7705 SAR Gen 2

23.30 priv-lvl-map

priv-lvl-map

Syntax

[**no**] **priv-lvl-map**

Context

[Tree] (config>service>vprn>aaa>rmt-srv>tacplus priv-lvl-map)

[Tree] (config>system>security>tacplus priv-lvl-map)

Full Context

configure service vprn aaa remote-servers tacplus priv-lvl-map

configure system security tacplus priv-lvl-map

Description

Commands in this context specify a series of mappings between TACACS+ priv-lvl and locally configured profiles for authorization. These mappings are used when the use-priv-lvl option is specified for tacplus authorization.

The **no** form of this command reverts to the default.

Default

priv-lvl-map

Platforms

7705 SAR Gen 2

23.31 private-interface

private-interface

Syntax

private-interface *ip-int-name*

no private-interface

Context

[Tree] (config>ipsec>client-db>client private-interface)

Full Context

configure ipsec client-db client private-interface

Description

This command specifies the private interface name that is used for tunnel setup.

The **no** form of this command reverts to the default.

Default

no private-interface

Parameters

ip-int-name

Specifies the name of the private interface.

Platforms

7705 SAR Gen 2

23.32 private-service

private-service

Syntax

private-service *service-id*

private-service name *service-name*

no private-service

Context

[Tree] (config>ipsec>client-db>client private-service)

Full Context

configure ipsec client-db client private-service

Description

This command specifies the private service ID that is used for tunnel setup.

The **no** form of this command reverts to the default.

Default

no private-service

Parameters

service-id

Specifies the service ID of the tunnel delivery service.

This variant of this command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **private-service name** *service-name* variant can be used in all configuration modes.

Values {*id* | *svc-name*}

id: 1 to 2147483647

svc-name: up to 64 characters (*svc-name* is an alias for input only. The *svc-name* gets replaced with an id automatically by SR OS in the configuration).

name service-name

Identifies the service, up to 64 characters.

Platforms

7705 SAR Gen 2

23.33 private-tcp-mss-adjust

private-tcp-mss-adjust

Syntax

private-tcp-mss-adjust *bytes*

private-tcp-mss-adjust *octets*

no private-tcp-mss-adjust

Context

[Tree] (config>router>if>ipsec>ipsec-tunnel private-tcp-mss-adjust)

[Tree] (config>ipsec>tnl-temp private-tcp-mss-adjust)

[Tree] (config>service>ies>if>sap>ip-tunnel private-tcp-mss-adjust)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel private-tcp-mss-adjust)

[Tree] (config>service>vprn>if>sap>ip-tunnel private-tcp-mss-adjust)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel private-tcp-mss-adjust)

[Tree] (config>service>vprn>if>sap>ipsec-tun private-tcp-mss-adjust)

Full Context

configure router interface ipsec ipsec-tunnel private-tcp-mss-adjust

configure ipsec tunnel-template private-tcp-mss-adjust

configure service ies interface sap ip-tunnel private-tcp-mss-adjust

configure service ies interface ipsec ipsec-tunnel private-tcp-mss-adjust

configure service vprn interface sap ip-tunnel private-tcp-mss-adjust

configure service vprn interface ipsec ipsec-tunnel private-tcp-mss-adjust

configure service vprn interface sap ipsec-tunnel private-tcp-mss-adjust

Description

This command enables TCP MSS to adjust for L2TPv3 tunnels, IPsec, or IP tunnels on the private side. When the command is configured, the system updates the TCP MSS option to the value of the received TCP SYN packet on the private side.

The **no** form of this command disables TCP MSS adjust on the private side.

Default

no private-tcp-mcc-adjust

Parameters

bytes

Specifies the new TCP MSS value in bytes.

Values 512 to 9000

octets

Specifies the new TCP MSS value in octets.

Values 512 to 9000

Platforms

7705 SAR Gen 2

23.34 probe-fail-enable

probe-fail-enable

Syntax

[**no**] **probe-fail-enable**

Context

[Tree] (config>saa>test>trap-gen probe-fail-enable)

Full Context

configure saa test trap-gen probe-fail-enable

Description

This command enables the generation of an SNMP trap when the consecutive probe failure threshold (configured using the **probe-fail-threshold** command) is reached during the execution of the SAA ping test. This command is not applicable to SAA trace route tests.

The **no** form of this command disables the generation of an SNMP trap.

Platforms

7705 SAR Gen 2

23.35 probe-fail-threshold

probe-fail-threshold

Syntax

probe-fail-threshold *threshold*

no probe-fail-threshold

Context

[Tree] (config>saa>test>trap-gen probe-fail-threshold)

Full Context

configure saa test trap-gen probe-fail-threshold

Description

This command configures the threshold for trap generation after ping probe failure.

This command has no effect when **probe-fail-enable** is disabled. This command is not applicable to SAA trace route tests.

The **no** form of this command returns the threshold value to the default.

Default

probe-fail-threshold 1

Parameters

threshold

Specifies the number of consecutive ping probe failures required to generate a trap.

Values 0 to 15

Platforms

7705 SAR Gen 2

23.36 probe-history

probe-history

Syntax

probe-history {**keep** | **drop** | **auto**}

Context

[Tree] (config>saa>test probe-history)

Full Context

configure saa test probe-history

Description

Specifies history probe behavior. Defaults are associated with various configured parameters within the SAA test. Auto (keep) is used for test with probe counts of 100 or less, and intervals of 1 second and above. Auto (drop) only maintains summary information for tests marked as continuous with file functions, probe counts more than 100 and intervals of less than 1 second. SAA tests that are not continuous with a write to file defaults to Auto (keep). The operator is free to change the default behaviors for each type. Each test that maintains per probe history consumes more system memory. When per probe entries are required, the probe history is available at the completion of the test.

Default

probe-history auto

Parameters

auto

An auto selector that determines the storage of the history information.

drop

Stores summarized min/max/avg data not per probe information for test runs. This may be configured for all tests to conserve memory.

keep

Stores per probe information for tests. This consumes significantly more memory than summary information and should only be used if necessary.

Platforms

7705 SAR Gen 2

23.37 process-arp-probes

process-arp-probes

Syntax

[**no**] **process-arp-probes**

Context

[Tree] (config>service>vpls>proxy-arp process-arp-probes)

Full Context

configure service vpls proxy-arp process-arp-probes

Description

This command enables router proxy ARP function replies to Duplicate Address Detection (DAD) ARP probes upon a successful proxy ARP table lookup.

The **no** form of this command disables the router from replying to DAD ARP probes.

Default

process-arp-probes

Platforms

7705 SAR Gen 2

23.38 process-cpm-traffic-on-sap-down

process-cpm-traffic-on-sap-down

Syntax

[**no**] **process-cpm-traffic-on-sap-down**

Context

[Tree] (config>service>vpls>sap process-cpm-traffic-on-sap-down)

Full Context

configure service vpls sap process-cpm-traffic-on-sap-down

Description

This command is applicable to simple SAPs configured on LAGs that are not part of any "endpoint" configurations or complicated resiliency schemes like MC-LAG with inter-chassis-backup (ICB) configurations. When configured, a simple LAG SAP is not removed from the forwarding plane and flooded traffic (unknown unicast, broadcast and multicast) is dropped on egress. This allows applicable control traffic that is extracted at the egress interface to be processed by the CPM. This command will not prevent a VPLS service from entering an operationally down state if it is the last active connection to enter a nonoperational state. By default, without this command, when a SAP on a LAG enters a nonoperational state, it is removed from the forwarding plane and no forwarding occurs to the egress.

The **no** form of this command removes a SAP over a LAG that is not operational from the forwarding process.

Default

no process-cpm-traffic-on-sap-down

Platforms

7705 SAR Gen 2

23.39 process-dad-neighbor-solicitations

process-dad-neighbor-solicitations

Syntax

[**no**] **process-dad-neighbor-solicitations**

Context

[Tree] (config>service>vpls>proxy-nd process-dad-neighbor-solicitations)

Full Context

configure service vpls proxy-nd process-dad-neighbor-solicitations

Description

This command enables the router proxy ND replies to Duplicate Address Detection (DAD) neighbor solicitations upon a successful proxy ND table lookup.

The **no** form of this command disables the router from replying to DAD neighbor solicitations.

Default

process-dad-neighbor-solicitations

Platforms

7705 SAR Gen 2

23.40 process-received-upa

process-received-upa

Syntax

[**no**] **process-received-upa**

Context

[Tree] (config>router>isis>upa process-received-upa)

Full Context

configure router isis prefix-unreachable process-received-upa

Description

This command enables processing of UPAs received from other routers. When configured, received UPAs are inserted into the unicast routing table as unreachable prefixes. When configured on an Area Boundary Router (ABR), received UPAs are inserted into the unreachable prefix table and redistributed into the other areas.

The **no** form of this command disables the processing of UPAs received from other routers. When disabled, received UPAs are ignored by the router.

Default

no process-received-upa

Platforms

7705 SAR Gen 2

23.41 profile

profile

Syntax

[**no**] **profile** *user-profile-name*

Context

[Tree] (config>system>security profile)

Full Context

configure system security profile

Description

This command creates a context to create user profiles for command authorization and other functions associated with a user.

Profiles can be used to deny or permit user access to entire command branches or to specific commands.

Once the profiles are created, the **user** command assigns users to one or more profiles. You can define up to 16 user profiles but a maximum of 8 profiles can be assigned to a user.

The **no** form of this command deletes a user profile.

Parameters

user-profile-name

Specifies the user profile name entered as a character string. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

Platforms

7705 SAR Gen 2

profile

Syntax

profile {**in** | **out**}

no profile

Context

[Tree] (config>saa>test>type-multi-line>lsp-ping>sr-policy profile)

[Tree] (config>saa>test>type-multi-line>lsp-ping profile)

Full Context

configure saa test type-multi-line lsp-ping sr-policy profile

configure saa test type-multi-line lsp-ping profile

Description

This command configures the profile state of the MPLS echo request packet.

The **no** form of this command reverts to the default value.

Default

profile out

Parameters

in

Specifies "in” as the profile state of the MPLS echo request packet.

out

Specifies "out” as the profile state of the MPLS echo request packet.

Platforms

7705 SAR Gen 2

profile

Syntax

profile {**in** | **out**}

no profile

Context

[Tree] (config>oam-pm>session>ip profile)

Full Context

configure oam-pm session ip profile

Description

This command defines whether the TWAMP Light PDU packet should be treated as in-profile or out-of-profile. The default has been selected because the forwarding class defaults to best effort.

The **no** form of this command restores the default value.

Default

profile out

Parameters

in

Specifies that the TWAMP Light PDU packet is sent as in-profile.

out

Specifies that the TWAMP Light PDU packet is sent as out-of-profile.

Platforms

7705 SAR Gen 2

profile

Syntax

profile {**in** | **out**}

no profile

Context

[Tree] (config>qos>sap-ingress>fc profile)

Full Context

configure qos sap-ingress fc profile

Description

This command places a forwarding class or subclass into a color aware profile mode. Normally, packets associated with a class are considered in-profile or out-of-profile solely based on the dynamic rate of the ingress queue relative to its CIR. Explicitly defining a class as in-profile or out-of-profile overrides this function by handling each packet with the defined profile state.

The profile command may only be executed when the forwarding class or the parent forwarding class (for a subclass) is mapped to a queue that has been enabled to support color aware profile packets. The queue may only be configured for profile-mode at the time the queue is created in the SAP ingress QoS policy.

A queue operating in profile-mode may support in-profile, out-of-profile, and non-profiled packets simultaneously. However, the high- and low-priority classification actions are ignored when the queue is in profile-mode.

The **no** form of this command removes an explicit in-profile or out-of-profile configuration on a forwarding class or subclass.

Default

no profile — The default profile state of a forwarding class or subclass is not to treat ingress packets as color aware. An explicit definition for in-profile or out-of-profile must be specified on the forwarding class or subclass.

Parameters

in

The **in** keyword is mutually exclusive to the **out** keyword. When the profile in command is executed, all packets associated with the class will be handled as in-profile. Packets explicitly handled as in-profile or out-of-profile still flow through the ingress service queue associated with the class to preserve order within flows. In-profile packets will count against the CIR of the queue, diminishing the amount of CIR available to other classes using the queue that are not configured with an explicit profile.

out

The **out** keyword is mutually exclusive to the **in** keyword. When the profile out command is executed, all packets associated with the class will be handled as out-of-profile. Packets explicitly handled as in-profile or out-of-profile still flow through the ingress service queue associated with the class to preserve order within flows. Out-of-profile packets will not count against the CIR of the queue, allowing other classes using the queue that are not configured with an explicit profile to be measured against the full CIR.

Platforms

7705 SAR Gen 2

profile

Syntax

profile *name* [**create**]

no profile *name*

Context

[Tree] (config>system>network-element-discovery profile)

Full Context

configure system network-element-discovery profile

Description

This command configures a profile to be used by IGP to advertise the network element information to its neighbors.

The **no** form of this command deletes the specified profile.

Parameters

name

Specifies the name of the profile, up to 32 characters.

Platforms

7705 SAR Gen 2

profile

Syntax

profile *user-profile-name*

no profile

Context

[Tree] (config>system>security>user-template profile)

Full Context

configure system security user-template profile

Description

This command configures the command authorization profile to associate with a user template. See the **user-template** command for more details.

Parameters

user-profile-name

The user profile name entered as a character string. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

Platforms

7705 SAR Gen 2

profile

Syntax

profile *cert-update-profile*

Context

[Tree] (config>system>security>pki>cert-auto-upd>cert profile)

Full Context

configure system security pki certificate-auto-update cert profile

Description

This command configures a **certificate-update-profile** to reference the update behavior.

Parameters

cert-update-profile

Specifies the certificate profile name, up to 32 characters.

Platforms

7705 SAR Gen 2

23.42 profile-capped

profile-capped

Syntax

[**no**] **profile-capped**

Context

[Tree] (config>qos>sap-ingress>policer profile-capped)

[Tree] (config>qos>sap-egress>policer profile-capped)

Full Context

configure qos sap-ingress policer profile-capped

configure qos sap-egress policer profile-capped

Description

Profile-capped mode enforces an overall in-profile burst limit to the CIR bucket for ingress undefined, ingress explicit in-profile, egress soft-in-profile, and egress explicit in-profile packets. The default behavior when profile-capped mode is not enabled is to ignore the CIR output state when an explicit in-profile packet is handled by an ingress or egress policer.

The profile-capped mode makes two changes:

- At egress, soft-in-profile packets (packets received from ingress as in-profile) are treated the same as explicit in-profile (unless explicitly reclassified as out-of-profile) and have an initial policer state of in-profile.

- At both ingress and egress, any packet output from the policer with a non-conforming CIR state are treated as out-of-profile (out-of-profile state is ignored for initial in-profile packets when profile-capped mode is not enabled).

Default

no profile-capped

Platforms

7705 SAR Gen 2

profile-capped

Syntax

[**no**] **profile-capped**

Context

[Tree] (config>qos>qgrps>egr>qgrp>policer profile-capped)

[Tree] (config>qos>qgrps>ing>qgrp>policer profile-capped)

Full Context

configure qos queue-group-templates egress queue-group policer profile-capped

configure qos queue-group-templates ingress queue-group policer profile-capped

Description

This command enables a limit on the profile.

Default

no profile-capped

Platforms

7705 SAR Gen 2

23.43 profile-out-preserve

profile-out-preserve

Syntax

[**no**] **profile-out-preserve**

Context

[Tree] (config>qos>sap-egress>policer profile-out-preserve)

Full Context

configure qos sap-egress policer profile-out-preserve

Description

This command specifies whether to preserve the color of offered out-of-profile traffic at sap-egress policer (profile of the packet can change based on egress CIR state).

When enabled, traffic determined as out-of-profile at ingress policer will be treated as out-of-profile at sap-egress policer.

Platforms

7705 SAR Gen 2

23.44 profile-preferred

profile-preferred

Syntax

profile-preferred

no profile-preferred

Context

[Tree] (config>qos>plcr-ctrl-plcy>root profile-preferred)

Full Context

configure qos policer-control-policy root profile-preferred

Description

The **profile-preferred** command ensures that the root policer provides a preference to consume its PIR bucket tokens at a given priority level to packets that have their profile state set to in-profile by the output of the child policer CIR bucket.

Default

no profile-preferred

Platforms

7705 SAR Gen 2

23.45 progress-indicator

progress-indicator

Syntax

progress-indicator

Context

[Tree] (config>system>management-interface>cli>md-cli>environment progress-indicator)

Full Context

configure system management-interface cli md-cli environment progress-indicator

Description

Commands in this context configure progress indicator parameters.

Platforms

7705 SAR Gen 2

23.46 prompt

prompt

Syntax

prompt

Context

[Tree] (config>system>management-interface>cli>md-cli>environment prompt)

Full Context

configure system management-interface cli md-cli environment prompt

Description

Commands in this context configure prompt parameters.

Platforms

7705 SAR Gen 2

23.47 propagate-admin-group

propagate-admin-group

Syntax

[**no**] **propagate-admin-group**

Context

[Tree] (config>router>mpls>lsp-template>fast-reroute propagate-admin-group)

[Tree] (config>router>mpls>lsp>fast-reroute propagate-admin-group)

Full Context

configure router mpls lsp-template fast-reroute propagate-admin-group

configure router mpls lsp fast-reroute propagate-admin-group

Description

The command enables the signaling of the primary LSP path admin-group constraints in the FRR object at the ingress.

When this command is executed, the admin-group constraints configured in the context of the P2P LSP primary path, or the ones configured in the context of the LSP and inherited by the primary path, are copied into the FAST_REROUTE object. The admin-group constraints are copied into the 'include-any' or 'exclude-any' fields.

The ingress LER thus propagates these constraints to the downstream nodes during the signaling of the LSP to allow them to include the admin-group constraints in the selection of the FRR backup LSP for protecting the LSP primary path.

The ingress LER inserts the FAST_REROUTE object by default in a primary LSP path message. If the user disables the object using the following command, the admin-group constraints will not be propagated: **config>router>mpls>no frr-object**.

Note that the same admin-group constraints can be copied into the Session Attribute object. They are intended for the use of an LSR, typically an ABR, to expand the ERO of an inter-area LSP path. They are also used by any LSR node in the path of a CSPF or non-CSPF LSP to check the admin-group constraints against the ERO regardless if the hop is strict or loose. These are governed strictly by the command:

config>router>mpls>lsp>propagate-admin-group

In other words, the user may decide to copy the primary path admin-group constraints into the FAST_REROUTE object only, or into the Session Attribute object only, or into both. Note, however, that the PLR rules for processing the admin-group constraints can make use of either of the two object admin-group constraints.

This feature is supported with the following LSP types and in both intra-area and inter-area TE where applicable:

- Primary path of a RSVP P2P LSP.
- S2L path of an RSVP P2MP LSP instance

- LSP template for an S2L path of an RSVP P2MP LSP instance.

The **no** form of this command disables the signaling of administrative group constraints in the FRR object.

Default

no propagate-admin-group

Platforms

7705 SAR Gen 2

propagate-admin-group

Syntax

[**no**] **propagate-admin-group**

Context

[Tree] (config>router>mpls>lsp propagate-admin-group)

[Tree] (config>router>mpls>lsp-template propagate-admin-group)

Full Context

configure router mpls lsp propagate-admin-group

configure router mpls lsp-template propagate-admin-group

Description

This command enables propagation of session attribute object with resource affinity (C-type 1) in PATH message. If an LSR receives a session attribute with resource affinity, then it will check the compatibility of admin-groups received in PATH message against configured admin-groups on the egress interface of LSP.

To support admin-group for inter-area LSP, the ingress node must configure propagating admin-groups within the session attribute object. If a PATH message is received by an LSR node that has the **cspf-on-loose-hop** option enabled and the message includes admin-groups, then the ERO expansion by CSPF to calculate the path to the next loose hop includes the admin-group constraints received from ingress node.

If this option is disabled, then the session attribute object without resource affinity (C-Type 7) is propagated in PATH message and CSPF at the LSR node does not include admin-group constraints.

This admin group propagation is supported with a P2P LSP, a P2MP LSP instance, and an LSP template.

The user can change the value of the **propagate-admin-group** option on the fly. A RSVP P2P LSP performs a Make-Before-Break (MBB) on changing the configuration. A S2L path of an RSVP P2MP LSP performs a Break-Before-Make on changing the configuration.

The **no** form of this command reverts to the default value.

Default

no propagate-admin-group

Platforms
7705 SAR Gen 2

23.48 propagate-mac-flush

propagate-mac-flush

Syntax
[**no**] **propagate-mac-flush**

Context
[Tree] (config>service>vpls propagate-mac-flush)

Full Context
configure service vpls propagate-mac-flush

Description
This command enabled propagation of mac-flush messages received from the specified T-LDP on all spoke and mesh-SDPs within the context of the VPLS service. The propagation will follow split-horizon principles and any data-path blocking in order to avoid looping of these messages.

Default
no propagate-mac-flush

Platforms
7705 SAR Gen 2

23.49 propagate-metric

propagate-metric

Syntax
[**no**] **propagate-metric**

Context
[Tree] (config>service>vprn>rip propagate-metric)

Full Context
configure service vprn rip propagate-metric

Description

This command enables the BGP MED to be used to configure the RIP metric at the BGP to RIP transition on egress routers. BGP always configures the BGP MED to the RIP metric at the ingress router. When **propagate-metric** is configured, the RIP metric at egress routers is configured as the BGP MED attribute added to the optional value configured with the **metric-out** command.

The **no** version of this command sets the RIP metric to the optional value configured with the **metric-out** command plus 1.

Default

no propagate-metric

Platforms

7705 SAR Gen 2

23.50 propagate-pmtu-v4

propagate-pmtu-v4

Syntax

[**no**] **propagate-pmtu-v4**

Context

[Tree] (config>router>if>ipsec>ipsec-tunnel propagate-pmtu-v4)

[Tree] (config>service>ies>if>sap>ip-tunnel propagate-pmtu-v4)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel propagate-pmtu-v4)

[Tree] (config>service>vprn>if>sap>ip-tunnel propagate-pmtu-v4)

[Tree] (config>ipsec>tnl-temp propagate-pmtu-v4)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel propagate-pmtu-v4)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel propagate-pmtu-v4)

Full Context

configure router interface ipsec ipsec-tunnel propagate-pmtu-v4

configure service ies interface sap ip-tunnel propagate-pmtu-v4

configure service ies interface ipsec ipsec-tunnel propagate-pmtu-v4

configure service vprn interface sap ip-tunnel propagate-pmtu-v4

configure ipsec tunnel-template propagate-pmtu-v4

configure service vprn interface ipsec ipsec-tunnel propagate-pmtu-v4

configure service vprn interface sap ipsec-tunnel propagate-pmtu-v4

Description

This command enables the system to propagate the path MTU learned from public side to private side (IPv4 hosts).

The **no** form of this command prevents the learned path MTU propagation.

Default

propagate-pmtu-v4

Platforms

7705 SAR Gen 2

23.51 propagate-pmtu-v6

propagate-pmtu-v6

Syntax

[**no**] **propagate-pmtu-v6**

Context

[Tree] (config>ipsec>tnl-temp propagate-pmtu-v6)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel propagate-pmtu-v6)

[Tree] (config>router>if>ipsec>ipsec-tunnel propagate-pmtu-v6)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel propagate-pmtu-v6)

[Tree] (config>service>ies>if>sap>ip-tunnel propagate-pmtu-v6)

[Tree] (config>service>vprn>if>sap>ip-tunnel propagate-pmtu-v6)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel propagate-pmtu-v6)

Full Context

configure ipsec tunnel-template propagate-pmtu-v6

configure service vprn interface ipsec ipsec-tunnel propagate-pmtu-v6

configure router interface ipsec ipsec-tunnel propagate-pmtu-v6

configure service ies interface ipsec ipsec-tunnel propagate-pmtu-v6

configure service ies interface sap ip-tunnel propagate-pmtu-v6

configure service vprn interface sap ip-tunnel propagate-pmtu-v6

configure service vprn interface sap ipsec-tunnel propagate-pmtu-v6

Description

This command enables the system to propagate the path MTU learned from public side to private side (IPv6 hosts).

The **no** form of this command prevents the learned path MTU propagation.

Default

propagate-pmtu-v6

Platforms

7705 SAR Gen 2

23.52 protection-type

protection-type

Syntax

protection-type {**link** | **node**}

no protection-type

Context

[Tree] (config>router>route-next-hop-policy>template protection-type)

Full Context

configure router route-next-hop-policy template protection-type

Description

This command configures the protection type constraint into the route next-hop policy template.

The user can select if link protection or node protection is preferred in the selection of an LFA next-hop for all IP prefixes and LDP FEC prefixes to which a route next-hop policy template is applied. The default in SR OS implementation is node protection. The implementation will fall back to the other type if no LFA next-hop of the preferred type is found.

When the route next-hop policy template is applied to an IP interface, all prefixes using this interface as a primary next-hop will follow the protection type preference specified in the template.

The **no** form deletes the protection type constraint from the route next-hop policy template.

Default

protection-type node

Parameters

{link | node}

Specifies the two possible values for the protection type.

Default node

Platforms

7705 SAR Gen 2

23.53 proto-version

proto-version

Syntax

proto-version {v070 | latest}

Context

[Tree] (config>system>grpc>gnmi proto-version)

Full Context

configure system grpc gnmi proto-version

Description

This command sets the gnmi.proto version that the GRPC server should use for all gNMI RPCs.

Default

proto-version latest

Parameters

v070

Specifies to use v0.7.0 for gNMI RPCs. Only use this option for backward compatibility with legacy collectors.

latest

Specifies to use the latest gnmi.proto version for gNMI RPCs. The latest version is v0.8.0.

Platforms

7705 SAR Gen 2

23.54 protocol

protocol

Syntax

protocol *protocol* **profile-name** *profile-name*

Context

[Tree] (config>system>security>pki>cert-upd-prof protocol)

Full Context

configure system security pki certificate-update-profile protocol

Description

This command configures the protocol to update the certificate.

Default

protocol cmpv2

Parameters

protocol

Specifies the protocol type.

Values cmpv2, est

profile-name

Specifies the name of the CA or EST profile to be used for the certificate update.

Platforms

7705 SAR Gen 2

protocol

Syntax

protocol *ipsec-protocol*

no protocol

Context

[Tree] (config>ipsec>static-sa protocol)

Full Context

configure ipsec static-sa protocol

Description

This command configures the security protocol to use for an IPsec manual SA. The **no** statement resets to the default value.

Default

protocol esp

Parameters

ipsec-protocol

Identifies the IPsec protocol used with this static SA.

Values **ah** — Specifies the Authentication Header protocol.**esp** — Specifies the Encapsulation Security Payload protocol.

Platforms

7705 SAR Gen 2

protocol

Syntax

protocol any

protocol *protocol-id* **port opaque**

protocol *protocol-id* **port any**

protocol *protocol-id* **port from** *begin-port-id* **to** *end-port-id*

no protocol

Context

[Tree] (config>ipsec>ts-list>remote>entry protocol)

[Tree] (config>ipsec>ts-list>local>entry protocol)

Full Context

configure ipsec ts-list remote entry protocol

configure ipsec ts-list local entry protocol

Description

This command specifies the protocol and port range in the IKEv2 traffic selector.

The SR OS supports OPAQUE ports and port ranges for the following protocols:

- TCP
- UDP
- SCTP
- ICMP
- ICMPv6
- MIPv6

For ICMP and ICMPv6, the *port* value takes the form *icmp-type*/*icmp-code*. For MIPv6, the *port* value is the mobility header type. For other protocols, only the **port any** configuration can be used.

Default

no protocol

Parameters

protocol-id

Specifies the protocol ID. The value can be a number, a protocol name, or **any**.

begin-port-id

Specifies the beginning of the port range.

Values For TCP, UDP, and SCTP, the value is the port number.

For ICMP and ICMPv6, the value takes the form *icmp-type*/*icmp-code*; for example, 0/0.

For MIPv6, the value is the mobility header type.

end-port-id

Specifies the end of the port range

Values For TCP, UDP, and SCTP, the value is the port number.

For ICMP and ICMPv6, the value takes the form *icmp-type*/*icmp-code*; for example, 0/0.

For MIPv6, the value is the mobility header type.

opaque

Specifies OPAQUE ports.

any

Specifies any port.

Platforms

7705 SAR Gen 2

protocol

Syntax

[**no**] **protocol** *protocol-id*

Context

[Tree] (config>filter>match-list>protocol-list protocol)

Full Context

configure filter match-list protocol-list protocol

Description

This command adds a protocol to the match protocol list.

The **no** form of this command removes the protocol from the **protocol-list**.

Parameters

protocol-id

protocol-number, *protocol-name*

protocol-number

Specifies the protocol number value to be added or removed from the protocol list. The value can be expressed as a decimal integer, or in hexadecimal or binary format.

Values [0 to 255]D

[0x0 to 0xFF]H

[0b0 to 0b11111111]B

protocol-name

Specifies the protocol name to be added or removed from the protocol list.

Values icmp, igmp, ip, tcp, egp, igp, udp, rdp, ipv6, ipv6-route, ipv6-frag, idrp, rsvp, gre, ipv6-icmp, ipv6-no-nxt, ipv6-opts, iso-ip, eigrp, ospf-igp, ether-ip, encap, pnni, pim, vrrp, l2tp, stp, ptp, isis, crtp, crudp, sctp.

Platforms

7705 SAR Gen 2

protocol

Syntax

protocol *protocol*

no protocol [*protocol*]

Context

[Tree] (config>vrrp>policy>priority-event>route-unknown protocol)

Full Context

configure vrrp policy priority-event route-unknown protocol

Description

This command adds one or more route sources to match the route unknown IP route prefix for a route unknown priority control event.

If the route source does not match one of the defined protocols, the match is considered unsuccessful and the **route-unknown** event transitions to the set state.

The **protocol** command is optional. If the **protocol** command is not executed, the comparison between the RTM prefix return and the **route-unknown** IP route prefix will not include the source of the prefix. The **protocol** command cannot be executed without at least one associated route source parameter.

All parameters are reset each time the **protocol** command is executed and only the explicitly defined protocols are allowed to match.

The **no** form of the command removes protocol route source as a match criteria for returned RTM route prefixes.

To remove specific existing route source match criteria, execute the **protocol** command and include only the specific route source criteria. Any unspecified route source criteria is removed.

Default

no protocol — No route source for the route unknown priority event is defined.

Parameters

protocol

Explicitly defined protocols

Values **bgp** - This parameter defines BGP as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **bgp** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **bgp** parameter, a returned route prefix with a source of BGP will not be considered a match and will cause the event to enter the set state.

bgp-vpn - This parameter defines **bgp-vpn** as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **bgp-vpn** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **bgp-vpn** parameter, a returned route prefix with a source of **bgp-vpn** will not be considered a match and will cause the event to enter the set state.

ospf - This parameter defines OSPF as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **ospf** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **ospf** parameter, a returned route prefix with a source of OSPF will not be considered a match and will cause the event to enter the set state.

is-is - This parameter defines IS-IS as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **is-is** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **is-is** parameter, a returned route prefix with a source of IS-IS will not be considered a match and will cause the event to enter the set state.

rip - This parameter defines RIP as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **rip** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **rip** parameter, a returned route prefix with a source of RIP will not be considered a match and will cause the event to enter the set state.

static - This parameter defines a static route as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **static** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **static** parameter, a returned route prefix with a source of static route will not be considered a match and will cause the event to enter the set state.

Platforms

7705 SAR Gen 2

protocol

Syntax

protocol *protocol-id*

no protocol

Context

[Tree] (config>system>security>mgmt-access-filter>ip-filter>entry protocol)

Full Context

configure system security management-access-filter ip-filter entry protocol

Description

This command configures an IP protocol type to be used as a management access filter match criterion.

The protocol type, such as TCP, UDP, and OSPF, is identified by its respective protocol number. Well-known protocol numbers include ICMP (1), TCP (6), and UDP (17).

The **no** form the command removes the protocol from the match criteria.

Parameters

protocol

Specifies the protocol number for the match criterion.

Values 1 to 255 (decimal)

Platforms

7705 SAR Gen 2

protocol

Syntax

[**no**] **protocol** *name* [**create**]

Context

[Tree] (config>sys>security>dist-cpu-protection>policy protocol)

Full Context

configure system security dist-cpu-protection policy protocol

Description

This command creates the control protocol for the policy.

The **no** form of this command means packets of the specified protocol are not monitored or enforced (although they count in the FP protocol queue) on the objects to which this DCP policy is assigned. The packets are treated as part of the **all-unspecified** protocol if the protocol is created in the policy.

Parameters

names

Signifies the protocol name.

The following explanatory notes for specific protocols apply:

- bfd-cpm - includes all BFD handled on the CPM, including the cpm-np type, single hop and multi hop, and MPLS-TP CC and CV BFD
- dhcp - includes DHCP for IPv4 and IPv6
- eth-cfm - 802.1ag and includes Y.1731. ETH-CFM packets on port and LAG-based facility MEPs are not included (but packets on tunnel MEPs are)
- icmp - includes IPv4 and IPv6 ICMP (including RS/RA/Redirect), except NS/NA Neighbor Discovery packets that are classified as a separate NDIS protocol
- icmp-ping-check - includes packets associated with ping-template functions
- isis - includes IS-IS used for SPBM
- ldp - includes LDP and T-LDP
- mpls-ttl - includes MPLS packets that are extracted because of an expired MPLS-TTL field
- ndis - includes IPv6 NS/NA Neighbor Discovery (not including RS/RA/Redirect which are classified as part of the ICMP protocol)
- ospf - includes all OSPFv2 and OSPFv3 packets
- pppoe-pppoa - includes PADx, LCP, PAP/CHAP, and NCPs
- ssh - includes TCP port 22 or a user-configured SSH server port
- vrrp - includes VRRP and SRRP packets
- multi-chassis - includes SR OS Multi-Chassis UDP port 1025 packets
- multi-chassis-sync - includes SR OS Multi-Chassis Sync TCP port 45067 packets
- all-unspecified - a special aggregate entry for protocols that are not explicitly specified. When configured, this option treats all extracted control packets that are not explicitly configured in the DCP policy as a single aggregate flow (or virtual protocol). It lumps together remaining control traffic to allow it to be rate limited as one flow. It includes all control traffic of all protocols that are extracted and sent to the CPM (including protocols that cannot be explicitly configured with the distributed CPU protection feature). Control

packets that are both forwarded and copied for extraction are not included. If a user later explicitly configures a protocol, that protocol is suddenly no longer part of the **all-unspecified** flow. This protocol must be explicitly configured to operate.

Values all-unspecified, arp, bfd-cpm, bgp, dhcp, eth-cfm, http-redirect, icmp, icmp-ping-check, igmp, isis, ldp, mld, mpls-ttl, multi-chassis, multi-chassis-sync, ndis, ospf, pim, pppoe-pppoa, radius, rsvp, ssh, vrrp

create

Mandatory keyword to create the protocol.

Platforms

7705 SAR Gen 2

protocol

Syntax

protocol *protocol* [**all** | { **instance** *instance*}]

protocol *protocol2* [*protocol2* (up to 5 max)]

no protocol

Context

[Tree] (config>router>policy-options>policy-statement>entry>from protocol)

Full Context

configure router policy-options policy-statement entry from protocol

Description

This command configures a routing protocol as a match criterion for a route policy statement entry. This command is used for both import and export policies depending how it is used.

The **protocol direct-interface** route type matches the specific direct interface host IPv4 /32 and IPv6 /128 routes. The **protocol direct** route type matches direct routes and does not match the specific /32 or /128 interface route itself.

Note:
The **instance** command cannot be used if multiple protocol names are specified for the *protocol2* parameter.

The **no** form of this command removes the protocol match criterion.

Default

no protocol

Parameters

protocol

Specifies the protocol name for the match criterion.

Values direct, static, bgp, isis, ospf, rip, aggregate, bgp-vpn, igmp, pim, ospf3, ldp, sub-mgmt, mld, managed, vpn-leak, nat, periodic, ipsec, dhcpv6-pd, dhcpv6-na, dhcpv6-ta, dhcpv6-pd-excl, ripng, bgp-label, direct-interface, arp-nd, rib-api, evpn-ifl, evpn-iff, srv6

instance

Specifies the OSPF, OSPFv3, or IS-IS protocol instance.

Values isis-inst — 0 to 127

ospf-inst — 0 to 31

ospf3-inst — 0 to 31, 64 to 95

protocol2

Specifies up to five protocol names to match on.

Values direct, static, isis, aggregate, bgp, bgp-label, direct-interface

all

Keyword that specifies to match on any OSPF, OSPFv3, or IS-IS protocol instance.

Platforms

7705 SAR Gen 2

protocol

Syntax

protocol *protocol* [**all** | **instance** *instance*]

protocol bgp bgp-label

no protocol

Context

[Tree] (config>router>policy-options>policy-statement>entry>to protocol)

Full Context

configure router policy-options policy-statement entry to protocol

Description

This command configures a routing protocol as a match criterion for a route policy statement entry. This command is used for both import and export policies depending how it is used.

The **no** form of this command removes the protocol match criterion.

Default

no protocol

Parameters

protocol

Specifies the protocol name to match on.

Values bgp, isis, ospf, rip, bgp-vpn, ospf3, vpn-leak, ldp, ripng, bgp-label

instance

Specifies the OSPF, OSPFv3, or IS-IS instance.

Values isis-inst — 0 to 127
ospf-inst — 0 to 31
ospf3-inst — 0 to 31, 64 to 95

all

Keyword that specifies to match on any OSPF, OSPFv3, or IS-IS protocol instance.

Platforms

7705 SAR Gen 2

23.55 protocol-list

protocol-list

Syntax

protocol-list *protocol-list-name* [**create**]

no protocol-list *protocol-list-name*

Context

[Tree] (config>filter>match-list protocol-list)

Full Context

configure filter match-list protocol-list

Description

This command creates a list of IP protocols that can be used in line card IP and IPv6 filters.

The **no** form of this command removes the IP protocol list.

Default

no protocol-list

Parameters

protocol-list-name

Specifies the name of the protocol list.

create

This keyword is required to create the protocol list. After it is created, the protocol list can be enabled with or without the **create** keyword.

Platforms

7705 SAR Gen 2

23.56 protocol-version

protocol-version

Syntax

protocol-version *TLS version*

no protocol-version

Context

[Tree] (config>system>security>tls>client-tls-profile protocol-version)

Full Context

configure system security tls client-tls-profile protocol-version

Description

This command configures the TLS version to be negotiated between the client and server.

When configured, the client adds the specified version as a supported version in its Hello message to the server. If **tls-version-all** is specified, the client adds both TLS 1.2 and TLS 1.3 as supported versions in its Hello message.

The **no** form of this command reverts to the default TLS version.

Default

protocol-version tls-version12

Parameters

TLS version

Specifies the TLS version to include in the client Hello message.

Values tls-version12, tls-version13, tls-version-all

Platforms

7705 SAR Gen 2

protocol-version

Syntax

protocol-version *TLS version*

no protocol-version

Context

[Tree] (config>system>security>tls>server-tls-profile protocol-version)

Full Context

configure system security tls server-tls-profile protocol-version

Description

This command configures the TLS version to be negotiated between the server and client.

When configured, the server adds the specified version as a supported version in its Hello message to the client. If **tls-version-all** is specified, the server adds both TLS 1.2 and TLS 1.3 as supported versions in its Hello message.

The **no** form of this command reverts to the default TLS version.

Default

protocol-version tls-version12

Parameters

TLS version

Specifies the TLS version to include in the server Hello message.

Values tls-version12, tls-version13, tls-version-all

Platforms

7705 SAR Gen 2

23.57 proxy-arp

proxy-arp

Syntax

[**no**] **proxy-arp**

Context

[Tree] (config>service>vpls proxy-arp)

Full Context

configure service vpls proxy-arp

Description

Commands in this context configure the proxy-ARP parameters in a VPLS service.

Default

no proxy-arp

Platforms

7705 SAR Gen 2

proxy-arp

Syntax

[**no**] **proxy-arp** [**mac** [*ieee-address*]] [**ip** [*ipaddr*] **all**]]

Context

[Tree] (debug>service>id proxy-arp)

Full Context

debug service id proxy-arp

Description

This command enables the debug of the proxy-arp function for a specified service. Alternatively, the debug can be enabled only for certain entries given by their IP or MAC addresses.

Platforms

7705 SAR Gen 2

23.58 proxy-arp-nd

proxy-arp-nd

Syntax

proxy-arp-nd

Context

[Tree] (config>service proxy-arp-nd)

Full Context

configure service proxy-arp-nd

Description

Commands in this context configure the service-level **proxy-arp-nd** commands.

Platforms

7705 SAR Gen 2

23.59 proxy-arp-policy

proxy-arp-policy

Syntax

[**no**] **proxy-arp-policy** *policy-name* [*policy-name*]

Context

[Tree] (config>service>vprn>if proxy-arp-policy)

[Tree] (config>service>ies>if proxy-arp-policy)

Full Context

configure service vprn interface proxy-arp-policy

configure service ies interface proxy-arp-policy

Description

This command specifies an existing policy-statement to analyze match and action criteria that controls the flow of routing information to and from a given protocol, set of protocols, or a neighbor.

The **no** form of this command disables the proxy ARP capability.

Parameters

policy-name

Specifies the export route policy name. Allowed values are any string, up to 32 characters, composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

The specified name must already be defined.

Platforms

7705 SAR Gen 2

proxy-arp-policy

Syntax

proxy-arp-policy *policy-name* [*policy-name*]

no proxy-arp-policy

Context

[Tree] (config>router>if proxy-arp-policy)

Full Context

configure router interface proxy-arp-policy

Description

This command enables and configure proxy ARP on the interface and specifies an existing policy-statement to analyze match and action criteria that controls the flow of routing information to and from a given protocol, set of protocols, or a specific neighbor. The policy-name is configured in the **config>router>policy-options** context.

Use proxy ARP so the router responds to ARP requests on behalf of another device. Static ARP is used when a router needs to know about a device on an interface that cannot or does not respond to ARP requests. Therefore, the router configuration can state that if it has a packet that has a certain IP address to send it to the corresponding ARP address.

Default

no proxy-arp-policy

Parameters

policy-name

Specifies the export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. A maximum of five policy names can be specified in a single statement. The specified policy names must already be defined.

Platforms

7705 SAR Gen 2

23.60 proxy-nd

proxy-nd

Syntax

[**no**] **proxy-nd**

Context

[Tree] (config>service>vpls proxy-nd)

Full Context

configure service vpls proxy-nd

Description

Commands in this context configure the proxy-ND parameters in a VPLS service.

Default

no proxy-nd

Platforms

7705 SAR Gen 2

proxy-nd

Syntax

[**no**] **proxy-nd** [**mac** [*ieee-address*]] [**ip** [*ipaddr*] **all**]]

Context

[Tree] (debug>service>id proxy-nd)

Full Context

debug service id proxy-nd

Description

This command enables the debug of the proxy-nd function for a specified service. Alternatively, the debug can be enabled only for certain entries given by their IPv6 or MAC addresses.

Platforms

7705 SAR Gen 2

23.61 proxy-nd-policy

proxy-nd-policy

Syntax

proxy-nd-policy *policy-name* [*policy-name*]

no proxy-nd-policy

Context

[Tree] (config>service>ies>if>ipv6 proxy-nd-policy)

Full Context

configure service ies interface ipv6 proxy-nd-policy

Description

This command configures a proxy neighbor discovery policy for the interface. This policy determines networks and sources for which proxy ND is attempted, when local proxy neighbor discovery is enabled.

The **no** form of this command reverts to the default value.

Parameters

policy-name

Specifies up to five the export route policy names. Allowed values are any string, up to 32 characters, composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

Up to 5 policy-names can be specified in a single statement.

Platforms

7705 SAR Gen 2

proxy-nd-policy

Syntax

proxy-nd-policy *policy-name* [*policy-name*]

no proxy-nd-policy

Context

[Tree] (config>service>vprn>if>ipv6 proxy-nd-policy)

Full Context

configure service vprn interface ipv6 proxy-nd-policy

Description

This command configures a proxy neighbor discovery policy for the interface.

Parameters

policy-name

Specifies up to five existing policy names.

Platforms

7705 SAR Gen 2

proxy-nd-policy

Syntax

proxy-nd-policy *policy-name* [*policy-name*]

no proxy-nd-policy

Context

[Tree] (config>router>if>ipv6 proxy-nd-policy)

Full Context

configure router interface ipv6 proxy-nd-policy

Description

This command configure a proxy neighbor discovery policy for the interface.

Parameters

policy-name

The neighbor discovery policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. A maximum of five policy names can be specified in a single statement. The specified policy names must already be defined.

Platforms

7705 SAR Gen 2

23.62 proxy-server

proxy-server

Syntax

proxy-server

Context

[Tree] (config>service>vpls>sap>dhcp proxy-server)

[Tree] (config>service>ies>if>dhcp proxy-server)

Full Context

configure service vpls sap dhcp proxy-server

configure service ies interface dhcp proxy-server

Description

Commands in this context configure DHCP proxy server parameters.

Platforms

7705 SAR Gen 2

23.63 psnp-authentication

psnp-authentication

Syntax

[**no**] **psnp-authentication**

Context

[Tree] (config>service>vprn>isis>level psnp-authentication)

[Tree] (config>service>vprn>isis psnp-authentication)

Full Context

configure service vprn isis level psnp-authentication

configure service vprn isis psnp-authentication

Description

This command enables authentication of individual ISIS packets of partial sequence number PDU (PSNP) type.

The **no** form of this command suppresses authentication of PSNP packets.

Platforms

7705 SAR Gen 2

psnp-authentication

Syntax

[**no**] **psnp-authentication**

Context

[Tree] (config>router>isis>level psnp-authentication)

[Tree] (config>router>isis psnp-authentication)

Full Context

configure router isis level psnp-authentication

configure router isis psnp-authentication

Description

This command enables authentication of individual IS-IS packets of partial sequence number PDU (PSNP) type.

The **no** form of this command suppresses authentication of PSNP packets.

Default

psnp-authentication

Platforms

7705 SAR Gen 2

23.64 public-key-authentication

public-key-authentication

Syntax

[**no**] **public-key-authentication**

Context

[Tree] (config>system>security>ldap public-key-authentication)

Full Context

configure system security ldap public-key-authentication

Description

This command enables public key retrieval from the LDAP server. If disabled (**no public-key-authentication**), password authentication is attempted via LDAP.

Default

no public-key-authentication

Platforms

7705 SAR Gen 2

23.65 public-key-min-bits

public-key-min-bits

Syntax

public-key-min-bits *bits*

no public-key-min-bits

Context

[Tree] (config>service>ies>if>ipv6>secure-nd public-key-min-bits)

Full Context

configure service ies interface ipv6 secure-nd public-key-min-bits

Description

This command configures the minimum acceptable key length for public keys used in the generation of a Cryptographically Generated Address (CGA).

Parameters

bits

Specifies the number of bits.

Values 512 to 1024

Platforms

7705 SAR Gen 2

public-key-min-bits

Syntax

public-key-min-bits *bits*

[**no**] **public-key-min-bits**

Context

[Tree] (config>service>vprn>if>send public-key-min-bits)

Full Context

configure service vprn interface ipv6 secure-nd public-key-min-bits

Description

This command configures the minimum acceptable key length for public keys used in the generation of a Cryptographically Generated Address (CGA).

Parameters

bits

Specifies the number of bits.

Values 512 to 1024

Platforms

7705 SAR Gen 2

public-key-min-bits

Syntax

public-key-min-bits *bits*

no public-key-min-bits

Context

[Tree] (config>router>if>ipv6>secure-nd public-key-min-bits)

Full Context

configure router interface ipv6 secure-nd public-key-min-bits

Description

This command configures the minimum acceptable key length for public keys used in the generation of a Cryptographically Generated Address (CGA).

Parameters

bits

Specifies the number of bits.

Values 512 to 1024

Platforms

7705 SAR Gen 2

23.66 public-key-only

public-key-only

Syntax

[**no**] **public-key-only**

Context

[Tree] (config>system>security>ssh>auth-method>server public-key-only)

Full Context

configure system security ssh authentication-method server public-key-only

Description

This command configures the SSH server to accept only the public-key authentication method.

The **no** form of this command configures the SSH server to accept public-key or password client authentication. If **interactive-authentication** is enabled in the **configure system security aaa remote-servers radius** or **configure system security aaa remote-servers tacplus** contexts, the SSH server also accepts interactive keyboard authentication.

Default

no public-key-only

Platforms

7705 SAR Gen 2

public-key-only

Syntax

public-key-only {**false**|**true**|**system**}

Context

[Tree] (config>system>security>user>ssh-auth-method>server public-key-only)

Full Context

configure system security user ssh-authentication-method server public-key-only

Description

This command configures the accepted SSH authentication method for the user connection.

Default

system

Parameters

false

Specifies the use of public-key only, or public-key and password for client authentication. If **interactive-authentication** is enabled in the **configure system security aaa remote-servers radius** or **configure system security aaa remote-servers tacplus** contexts, the SSH server also accepts interactive keyboard authentication.

true

Specifies the use of public-key authentication only.

system

Specifies the use of the SSH authentication method configured at the system level.

Platforms

7705 SAR Gen 2

23.67 public-keys

public-keys

Syntax

public-keys

Context

[Tree] (config>system>security>user public-keys)

Full Context

configure system security user public-keys

Description

This command allows the user to enter the context to configure public keys for SSH.

Platforms

7705 SAR Gen 2

23.68 public-tcp-mss-adjust

public-tcp-mss-adjust

Syntax

public-tcp-mss-adjust *bytes*

public-tcp-mss-adjust *octets*

public-tcp-mss-adjust auto

no public-tcp-mss-adjust

Context

[Tree] (config>service>vprn>if>sap>ipsec-tun public-tcp-mss-adjust)

[Tree] (config>ipsec>tnl-temp public-tcp-mss-adjust)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel public-tcp-mss-adjust)

[Tree] (config>router>if>ipsec>ipsec-tunnel public-tcp-mss-adjust)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel public-tcp-mss-adjust)

[Tree] (config>service>ies>if>sap>ip-tunnel public-tcp-mss-adjust)

Full Context

configure service vprn interface sap ipsec-tunnel public-tcp-mss-adjust

configure ipsec tunnel-template public-tcp-mss-adjust

configure service vprn interface ipsec ipsec-tunnel public-tcp-mss-adjust

configure router interface ipsec ipsec-tunnel public-tcp-mss-adjust

configure service ies interface ipsec ipsec-tunnel public-tcp-mss-adjust

configure service ies interface sap ip-tunnel public-tcp-mss-adjust

Description

This command enables the Maximum Segment Size (MSS) for the TCP traffic in an IPsec tunnel which is sent from the public network to the private network. The system may use this value to adjust or insert the MSS option in TCP SYN packet.

If the **auto** parameter is specified, the system derives the new MSS value based on the public MTU and IPsec overhead.

The **no** form of this command disables TCP MSS adjust on the public side.

Default

no public-tcp-mss-adjust

Parameters

auto

Derive the new MSS value based on the public MTU and IPsec overhead.

bytes

Specifies the new TCP MSS value in bytes.

Values 512 to 9000

octets

Specifies the new TCP MSS value in octets

Values 512 to 9000

Platforms

7705 SAR Gen 2

23.69 purge-timer

purge-timer

Syntax

purge-timer *minutes*

no purge-timer

Context

[Tree] (config>router>bgp purge-timer)

Full Context

configure router bgp purge-timer

Description

When the system sends a VPN-IP Route-Refresh to a peer it sets all the VPN-IP routes received from that peer (in the RIB-IN) to stale and starts the purge-timer. If the routes are not updated (refreshed) before the purge-timer has expired then the routes are removed.

The BGP purge timer configures the time before stale routes are purged.

The **no** form of this command reverts to the default.

Default

purge-timer 10

Parameters

minutes

Specifies the maximum time before stale routes are purged.

Values 1 to 60

Platforms

7705 SAR Gen 2

23.70 push

push

Syntax

push {*label* | **implicit-null-label**} **nexthop** *ip-address*

no push {*out-label* | **implicit-null-label**}

Context

[Tree] (config>router>mpls>static-lsp push)

Full Context

configure router mpls static-lsp push

Description

This command specifies the label to be pushed on the label stack and the next hop IP address for the static LSP.

The **no** form of this command removes the association of the label to push for the static LSP.

Parameters

implicit-null-label

Specifies the use of the implicit label value for the push operation.

label

The label to push on the label stack. Label values 16 through 1,048,575 are defined as follows:

- label values 16 through 31 are reserved
- label values 32 through 1,023 are available for static assignment

- label values 1,024 through 2,047 are reserved for future use
- label values 2,048 through 18,431 are statically assigned for services
- label values 28,672 through 131,071 are dynamically assigned for both MPLS and services
- label values 131,072 through 1,048,575 are reserved for future use

Values 16 to 1048575

nexthop *ip-address*

Specifies the IP address of the next hop towards the LSP egress router. If an ARP entry for the next hop exists, then the static LSP is marked operational. If ARP entry does not exist, software sets the operational status of the static LSP to down and continues to ARP for the configured nexthop. Software continuously tries to ARP for the configured nexthop at a fixed interval.

Platforms

7705 SAR Gen 2

23.71 pw-routing

pw-routing

Syntax

pw-routing

Context

[Tree] (config>service pw-routing)

Full Context

configure service pw-routing

Description

Commands in this context configure dynamic multi-segment pseudowire (MS-PW) routing. Pseudowire routing must be configured on each node that will be a T-PE or an S-PE.

Platforms

7705 SAR Gen 2

23.72 pw-status-signaling

pw-status-signaling

Syntax

[**no**] **pw-status-signaling**

Context

[Tree] (config>service>epipe>spoke-sdp pw-status-signaling)

Full Context

configure service epipe spoke-sdp pw-status-signaling

Description

This command enables pseudowire status signaling for this spoke SDP binding.

The **no** form of this command disables the status signaling.

Default

pw-status-signaling

Platforms

7705 SAR Gen 2

pw-status-signaling

Syntax

[**no**] **pw-status-signaling**

Context

[Tree] (config>service>vpls>spoke-sdp pw-status-signaling)

Full Context

configure service vpls spoke-sdp pw-status-signaling

Description

This command specifies the type of signaling used by this multi-segment pseudowire provider-edge for this service.

When no pw-status-signaling is enabled, the 7705 SAR Gen 2 will not include the pseudowire status TLV in the initial label mapping message of the pseudowire used for a spoke-SDP. This will force both 7705 SAR Gen 2 PEs to use the pseudowire label withdrawal method for signaling pseudowire status.

If pw-status-signaling is configured, the node will include the use of the pseudowire status TLV in the initial label mapping message for the pseudowire.

Platforms

7705 SAR Gen 2

23.73 pw-template

pw-template

Syntax

pw-template *policy-id* [**use-provisioned-sdp** | [**prefer-provisioned-sdp**] [**auto-gre-sdp**]][**create**] [**name** *name*]

no **pw-template** *policy-id*

Context

[Tree] (config>service pw-template)

Full Context

configure service pw-template

Description

This command configures an SDP template.

Parameters

policy-id

Specifies a number that uniquely identifies a template for the creation of an SDP.

Values *policy-id*: 1 to 2147483647

use-provisioned-sdp

Specifies whether to use an already provisioned SDP. When specified, the tunnel manager is consulted for an existing active SDP (with a matching far-end address), and the SDP with the lowest metric is chosen. If there are multiple SDPs with the same metric, then the highest SDP identifier that is oper-up is chosen. The choice of SDP can be configured by applying **sdp-include/exclude** in the PW template together with an sdp-group in the provisioned SDPs. This option, and the **auto-gre-sdp** option, are mutually exclusive.

prefer-provisioned-sdp

Specifies that if an existing matching SDP that conforms to any restrictions defined in the **pw-template** is found (for example, **sdp-include**/**exclude** *group*), then it will be used, following the same logic as for the **use-provisioned-sdp** parameter. Otherwise, the command will automatically create an SDP in the same manner as if the user did not specify any option. This option and the **use-provisioned-sdp** option are mutually exclusive.

auto-gre-sdp

Specifies that an SDP should automatically be created using a GRE tunnel. This option and the **use-provisioned-sdp** option are mutually exclusive. The PW template parameters **hash-label**, **entropy-label** and **sdp-include/exclude** are ignored when an GRE SDP is auto-created.

auto-mpls-sdp

Specifies that an SDP should automatically be created using an MPLS tunnel. This is the default.

create

This keyword is required when first creating the configuration context. Once the context is created, it is possible to navigate into the context without the **create** keyword.

name name

A name of the operator's choice, up to 64 characters. The name is saved as part of the configuration.

If a name is not specified at creation time, then SR OS assigns a string version of the policy-id as the name.

Values *name*: 64 characters maximum

Platforms

7705 SAR Gen 2

23.74 pw-template-bind

pw-template-bind

Syntax

pw-template-bind *policy-id*

no pw-template-bind

Context

[Tree] (config>service>epipe>spoke-sdp-fec pw-template-bind)

Full Context

configure service epipe spoke-sdp-fec pw-template-bind

Description

This command binds includes the parameters included in a specific PW template to a spoke SDP.

The **no** form of this command removes the values from the configuration.

Parameters

policy-id

Specifies the existing policy ID.

Values 1 to 2147483647

Platforms

7705 SAR Gen 2

23.75 pw-template-binding

pw-template-binding

Syntax

pw-template-binding *policy-id* [**import-rt** { *ext-community* [*ext-community*]}] [**endpoint** *endpoint-name*]

no pw-template-binding *policy-id*

Context

[Tree] (config>service>epipe>bgp pw-template-binding)

Full Context

configure service epipe bgp pw-template-binding

Description

This command binds the advertisements received with the route targets (RT) that match the configured list (either the generic or the specified import) to a specific pw-template. If the RT list is not present, or if multiple matches are found, the numerically lowest pw-template is used.

The pw-template-binding applies to BGP-VPWS when enabled in the Epipe.

For BGP VPWS, the following additional rules govern the use of pseudowire-template:

- On transmission, the settings for the L2-Info extended community in the BGP updates are derived from the pseudowire template attributes. If multiple pseudowire template bindings (with or without import-rt) are specified for the same VPWS instance the first pw-template entry will be used for the information in the BGP update sent.
- On reception, the values of the parameters in the L2-Info extended community of the BGP updates are compared with the settings from the corresponding pseudowire template bindings. The following steps are used to determine the local pw-template:
  - The RT values are matched to determine the pw-template. The route targets configured for each pw-template-binding are compared to the route targets within the BGP update. The PW template corresponding to **pw-template-binding** with the first matching route target is used to for the SDP. The matching is performed from the lowest PW template binding identifier to the highest.
  - If no pw-template-binding matches are found from the previous step, the first (numerically lowest) configured pw-template entry without any route-target configured will be used.

If the value used for Layer 2 MTU (unless the value zero is received), or control word does not match, the pseudowire is created but with the operationally down state.

If the value used for the S (sequenced delivery) flags is not zero the pseudowire is not created.

The **tools perform** commands can be used to control the application of changes in pw-template for BGP-VPWS.

The **no** form of this command removes the values from the configuration.

Parameters

policy-id

Specifies an existing policy ID.

Values 1 to 2147483647

import-rt ext-comm

Specifies the communities, up to five, allowed to be accepted from remote PE neighbors. An extended BGP community in the type:x:y format. The value x can be an integer or IP address. The type can be the target or origin.

Values target:{*ip-addr*:*comm-val* | *2byte-asnumber*:*ext-comm-val*| *4byte-snumber*:*comm-val*}

ip-addr	a.b.c.d
comm-val	0 to 65535
2byte-asnumber	0 to 65535
ext-comm-val	0 to 4294967295
4byte-asnumber	0 to 4294967295

endpoint-name

Specifies the name of the endpoint the BGP PW template is associated with, up to 32 characters. When the configured endpoint is associated to the **pw-template-binding** of a BGP VPWS service, EVPN MPLS can also be configured and associated to the same endpoint in the same Epipe service. Modifying this element causes the parent element to be recreated automatically in order for the new value to take effect.

Platforms

7705 SAR Gen 2

pw-template-binding

Syntax

pw-template-binding *policy-id* [**split-horizon-group** *group-name*] [**import-rt** {*ext-community*}]

no pw-template-bind *policy-id*

Context

[Tree] (config>service>vpls>bgp pw-template-binding)

Full Context

configure service vpls bgp pw-template-binding

Description

This command binds the advertisements received with the route target (RT) that matches the configured list (either the generic or the specified import) to a specific PW template. If the RT list is not present the pw-template is used for all of them.

The **pw-template-binding** applies to both BGP-AD and BGP-VPLS if these features are enabled in the VPLS.

For BGP VPLS the following additional rules govern the use of pseudowire-template.

- On transmission, the settings for the L2-Info extended community in the BGP update are derived from the pseudowire template attributes. If multiple pseudowire template bindings (with or without import-rt) are specified, the first pw-template entry will be used for the information in the BGP update sent.
- On reception, the values of the parameters in the L2-Info extended community of the BGP update are compared with the settings from the corresponding pw-template. The following steps are used to determine the local pw-template.
  - The RT values are matched to determine the pw-template. The route targets configured for each pw-template-binding are compared to the route targets within the BGP update. The PW template corresponding to pw-template-binding with the first matching route target is used to for the SDP. The matching is performed from the lowest PW template binding identifier to the highest
  - If no pw-templates matches are found from the previous step, the first (numerically lowest) configured pw-template entry without any route-target configured will be used.

If the values used for Layer 2 MTU (unless the value zero is received) or control word flag do not match, the pseudowire is created but with the operationally down state.

If the value used for the S (sequenced delivery) flags is not zero, the pseudowire is not created.

The tools perform commands can be used to control the application of changes in pw-template for both BGP-AD and BGP-VPLS.

The **no** form of this command removes the values from the configuration.

Parameters

policy-id

Specifies an existing policy ID

Values 1 to 2147483647

group-name

The specified group-name overrides the split horizon group template settings

import-rt ext-comm

Specifies communities allowed to be accepted from remote PE neighbors. An extended BGP community in the **type**:**x**:**y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. **x** and **y** are 16-bit integers. A maximum of five *import-rt ext-com* can be specified.

Values target:{*ip-addr:comm-val*| *2byte-asnumber:ext-comm-val*| *4byte-asnumber:comm-val*}

ip-addr: a.b.c.d

comm-val: [0 to 65535]

2byte-as-number: [0 to 4294967295]

ext-comm-val: [0 to 4294967295]

4byte-asnumber: [0 to 4294967295]

Platforms

7705 SAR Gen 2

23.76 pw-template-id-range

pw-template-id-range

Syntax

pw-template-id-range start *pw-template-id* **end** *pw-template-id*

no pw-template-id-range

Context

[Tree] (config>service>md-auto-id pw-template-id-range)

Full Context

configure service md-auto-id pw-template-id-range

Description

This command specifies the range of IDs used by SR OS to automatically assign an ID to PW templates that are created in model-driven interfaces without an ID explicitly specified by the user or client.

A PW template created with an explicitly-specified ID cannot use an ID in this range. In the classic CLI and SNMP, the ID range cannot be changed while objects exist inside the previous or new range. In MD-CLI interfaces, the range can be changed, which causes any previously existing objects in the previous ID range to be deleted and re-created using a new ID in the new range.

The **no** form of this command removes the range values.

See the **config>service md-auto-id** command for further details.

Default

no pw-template-id-range

Parameters

start *pw-template-id*

Specifies the lower value of the ID range. The value must be less than or equal to the **end** value.

Values 1 to 2147483647

end *pw-template-id*

Specifies the upper value of the ID range. The value must be greater than or equal to the **start** value.

Values 1 to 2147483647

Platforms

7705 SAR Gen 2

23.77 pwc

pwc

Syntax

pwc [**previous**]

Context

[Tree] (pwc)

Full Context

pwc

Description

This command displays the present or previous working context of the CLI session. The **pwc** command provides a user who is in the process of dynamically configuring a chassis a way to display the current or previous working context of the CLI session. The **pwc** command displays a list of the CLI nodes that hierarchically define the current context of the CLI instance of the user.

The following shows an output example:

```
A:ALA-1>config>router>bgp>group# pwc
------------------------------------------------
Present Working Context :
------------------------------------------------
<root>
  configure
  router Base
  bgp
  group test
  ospf
```

```
  area 1
-----------------------------------------------
A:ALA-1>config>router>bgp>group#
```

When the **previous** keyword is specified, the previous context displays. This is the context entered by the CLI parser upon execution of the **exit** command. The current context of the CLI is not affected by the **pwc** command.

Parameters

previous

Displays the previous present working context.

Platforms

7705 SAR Gen 2

23.78 pxc

```
pxc
```

Syntax

pxc *pxc-id* [**create**]

no pxc *pxc-id*

Context

[Tree] (config>port-xc pxc)

Full Context

configure port-xc pxc

Description

This command creates a port cross-connect (PXC) object. Referencing an Ethernet port within the PXC object will automatically configure this Ethernet port as a loopback port. The node will automatically create two PXC sub-ports under this Ethernet port. The configuration of PXC sub-ports can be accessed through the CLI.

Parameters

pxc-id

Specifies the port cross-connect identifier.

Values 1 to 64

Platforms

7705 SAR Gen 2

24 q Commands

24.1 qinq-etype

qinq-etype

Syntax

qinq-etype *qinq-etype-value*

no qinq-etype

Context

[Tree] (config>port>ethernet qinq-etype)

Full Context

configure port ethernet qinq-etype

Description

This command configures the Ethertype used for Q-in-Q encapsulation.

The **no** form of this command reverts the qinq-etype value to the default.

Default

no qinq-etype

Parameters

qinq-etype-value

Specifies the qinq-etype to expect in the form of 0x600 to 0xfff.

Values 1536 to 65535 in decimal or hex formats

Platforms

7705 SAR Gen 2

24.2 qinq-mark-top-only

qinq-mark-top-only

Syntax

[**no**] **qinq-mark-top-only**

Context

[Tree] (config>service>vprn>if>sap>egress qinq-mark-top-only)

Full Context

configure service vprn interface sap egress qinq-mark-top-only

Description

When the encapsulation type is qinq for the access port for the specified SAP, enabling this command specifies which P-bits or DEI bit to mark during packet egress. Only the P-bits or DEI bit in the top Q tag are marked. When this command is disabled, both sets of P-bits and the DEI bit are marked.

Default

no qinq-mark-top-only

Platforms

7705 SAR Gen 2

qinq-mark-top-only

Syntax

[**no**] **qinq-mark-top-only**

Context

[Tree] (config>service>epipe>sap>egress qinq-mark-top-only)

Full Context

configure service epipe sap egress qinq-mark-top-only

Description

When enabled (the encapsulation type of the access port where this SAP is defined as qinq), the **qinq-mark-top-only** command specifies which P-bits/DEI bit to mark during packet egress. When disabled, both set of P-bits/DEI bit are marked. When the enabled, only the P-bits/DEI bit in the top Q-tag are marked.

Default

no qinq-mark-top-only

Platforms

7705 SAR Gen 2

qinq-mark-top-only

Syntax

[**no**] **qinq-mark-top-only**

Context

[Tree] (config>service>vpls>sap>egress qinq-mark-top-only)

Full Context

configure service vpls sap egress qinq-mark-top-only

Description

When enabled (the encapsulation type of the access port where this SAP is defined as qinq), the **qinq-mark-top-only** command specifies which P-bits/DEI bit to mark during packet egress. When disabled, both set of P-bits/DEI bit are marked. When enabled, only the P-bits/DEI bit in the top Q-tag are marked.

The **no** form of this command disables the command.

Default

no qinq-mark-top-only

Platforms

7705 SAR Gen 2

qinq-mark-top-only

Syntax

[**no**] **qinq-mark-top-only**

Context

[Tree] (config>service>ies>if>sap>egress qinq-mark-top-only)

Full Context

configure service ies interface sap egress qinq-mark-top-only

Description

When enabled (the encapsulation type of the access port where this SAP is defined as qinq), the **qinq-mark-top-only** command specifies which P-bits/DEI bit to mark during packet egress. When disabled, both set of P-bits/DEI bit are marked. When the enabled, only the P-bits/DEI bit in the top Q-tag are marked.

Default

no qinq-mark-top-only

Platforms

7705 SAR Gen 2

24.3 qos

qos

Syntax

qos *policy-id*

qos *policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*

no qos [**policy-id**]

Context

[Tree] (config>service>vprn>if>sap>egress qos)

[Tree] (config>service>ies>if>sap>egress qos)

[Tree] (config>service>vpls>sap>egress qos)

Full Context

configure service vprn interface sap egress qos

configure service ies interface sap egress qos

configure service vpls sap egress qos

Description

This command associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP) or IP interface.

QoS egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the policy ID does not exist, an error is returned.

The **qos** command associates both ingress and egress QoS policies. The **qos** command only allows ingress policies to be associated on SAP or IP interface ingress and egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type returns an error.

By default, no specific QoS policy is associated with the SAP or IP interface for ingress or egress, so the default QoS policy is used.

The **no** form of this command removes the QoS policy association from the SAP or IP interface, and the QoS policy reverts to the default.

Parameters

policy-id

The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.

Values 1 to 65535

port-redirect-group

This keyword associates a SAP egress with an instance of a named queue group template on the egress port of a given IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command.

queue-group-name

Specifies the name of the egress port queue group of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid egress queue group, created under config>port>ethernet>access>egress.

instance *instance-id*

Specifies the instance of the named egress port queue group on the IOM/IMM/XMA.

Values 1 to 40960

Default 1

Platforms

7705 SAR Gen 2

qos

Syntax

qos policy-id [**port-redirect-group** *queue-group-name* **instance** *instance-id*]

no qos

Context

[Tree] (config>service>vpls>sap>egress qos)

Full Context

configure service vpls sap egress qos

Description

This command associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP).

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the policy-id does not exist, an error will be returned.

The qos command is used to associate both ingress and egress QoS policies. The qos command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a specified type will return an error.

When an egress QoS policy is associated with an IES IP interface that has been bound to a VPLS, the policy becomes associated with every SAP on the VPLS and augments the egress QoS policy that is defined on each SAP. Packets that are bridged will be processed using the policy defined on the VPLS SAP; packets that are routed will be processed using the policy defined in the IES IP interface- binding context.

By default, if no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.

The **no** form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

Parameters

port-redirect-group

Associates a SAP egress with an instance of a named queue group template on the egress port of a specified IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command.

queue-group-name

Specifies the name of the egress port queue group of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid egress queue group, created under *config>port>ethernet>access>egress*.

instance *instance-id*

Specifies the instance of the named egress port queue group on the IOM/IMM/XMA

Values 1 to 40960

Default 1

Platforms

7705 SAR Gen 2

qos

Syntax

qos *policy-id* [**shared-queuing** | **multipoint-shared**] [**fp-redirect-group** *queue-group-name* **instance** *instance-id*]

qos *policy-id* [**shared-queuing** | **multipoint-shared**]

no qos [*policy-id*]

Context

[Tree] (config>service>ies>if>sap>ingress qos)

[Tree] (config>service>vprn>if>sap>ingress qos)

Full Context

configure service ies interface sap ingress qos

configure service vprn interface sap ingress qos

Description

This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP).

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the policy ID does not exist, an error is returned.

The **qos** command is used to associate both ingress and egress QoS policies. The qos command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type returns an error.

By default, no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.

The no form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

The **no** form of this command removes the QoS policy association from the SAP or IP interface, and the QoS policy reverts to the default.

Parameters

policy-id

The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.

1 to 65535

shared-queuing

Specifies the ingress shared queue policy used by this SAP. When the value of this object is null it means that the SAP uses individual ingress QoS queues instead of the shared ones.

multipoint-shared

Specifies that this queue-id is for multipoint forwarded traffic only. This queue-id can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. Attempting to map forwarding class unicast traffic to a multipoint queue generates an error; no changes are made to the current unicast traffic queue mapping.

A queue must be created as multipoint. The multipoint designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the multipoint keyword, an error is generated and the command is not executed.

The multipoint keyword can be entered in the command line on a preexisting multipoint queue to edit queue ID parameters.

Default Present (the queue is created as non-multipoint).

Values **Multipoint** or not present.

fp-redirect-group

Creates an instance of a named queue group template on the ingress forwarding plane of a given IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command. The named queue group template can contain only policers. If it contains queues, then the command fails.

queue-group-name

Specifies the name of the queue group template to be instantiated on the forwarding plane of the IOM/IMM/XMA, up to 32 characters. The queue-group-name must correspond to a valid ingress queue group template name, configured in the **config>qos>queue-group-templates** context.

instance-id

Specifies the instance of the named queue group to be created on the IOM/IMM/XMA ingress forwarding plane.

Platforms

7705 SAR Gen 2

qos

Syntax

qos *policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*

qos *policy-id*

no qos [*policy-id*]

Context

[Tree] (config>service>epipe>sap>egress qos)

Full Context

configure service epipe sap egress qos

Description

This command associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP).

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the *policy-id* does not exist, an error will be returned.

The **qos** command, when used under the egress context, is used to associate egress QoS policies.

The **qos** command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a specified type will return an error.

By default, if no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.

The **no** form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

Parameters

policy-id

The egress policy ID to associate with SAP on egress. The policy ID must already exist.

Values 1 to 65535

queue-group-name

Specifies the name of the egress port queue group of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid egress queue group, created under *config>port>ethernet>access>egress*.

instance-id

Specifies the instance of the named egress port queue group on the IOM/IMM/XMA.

Values 1 to 40960

Default 1

Platforms

7705 SAR Gen 2

qos

Syntax

qos *policy-id* [**shared-queuing**] [**fp-redirect-group** *queue-group-name* **instance** *instance-id*]

no qos

Context

[Tree] (config>service>epipe>sap>ingress qos)

Full Context

configure service epipe sap ingress qos

Description

This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP).

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the policy-id does not exist, an error will be returned.

The **qos** command, when used under the ingress context, is used to associate ingress QoS policies. The **qos** command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a specified type will return an error.

By default, if no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.

The **no** form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

Parameters

policy-id

The ingress policy ID to associate with SAP or IP interface on ingress. The policy ID must already exist.

Values 1 to 65535

shared-queuing

This keyword can only be specified on SAP ingress. The shared-queuing keyword specifies the shared queue policy will be used by this SAP. When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

fp-redirect-group

This keyword can only be used on SAP ingress and associates a SAP ingress with an instance of a named queue group template on the ingress forwarding plane of a specified IOM/IMM/XMA. The queue-group-name and **instance** *instance-id* are mandatory parameters when executing the command.

queue-group-name

Specifies the name of the queue group to be instance on the forwarding plane of the IOM/IMM/XMA, up to 32 characters in length. The *queue-group-name* must correspond to a valid ingress forwarding plane queue group, created under **config**>**card**>**fp**>**ingress**>**access**.

instance-id

Specifies the instance of the named queue group on the IOM/IMM/XMA ingress forwarding plane.

Platforms

7705 SAR Gen 2

qos

Syntax

qos *network-policy-id* **port-redirect-group** *queue-group-name* [**instance** *instance-id*]

no qos

Context

[Tree] (config>service>vpls>spoke-sdp>egress qos)

[Tree] (config>service>vpls>mesh-sdp>egress qos)

[Tree] (config>service>epipe>spoke-sdp>egress qos)

Full Context

configure service vpls spoke-sdp egress qos

configure service vpls mesh-sdp egress qos

configure service epipe spoke-sdp egress qos

Description

This command is used to redirect pseudowire (PW) packets to an egress port queue-group for the purpose of shaping.

The egress PW shaping provisioning model allows the mapping of one or more PWs to the same instance of queues, or policers and queues, that are defined in the queue-group template.

Operationally, the provisioning model consists of the following steps:

1. Create an egress queue-group template and configure queues only, or policers and queues for each FC that needs to be redirected.
2. Apply the queue-group template to the network egress context of all ports where there exists a network IP interface that the PW packets can be forwarded on. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.
3. Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue-group in the egress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different PWs to different queue-group templates.
4. Apply this network QoS policy to the egress context of a spoke-sdp inside a service, or to the egress context of a PW template and specify the redirect queue-group name.

One or more spoke-sdps can have their FCs redirected to use queues only, or queues and policers in the same queue-group instance.

The following are the constraints and rules of this provisioning model:

1. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name does not exist, the association is failed at the time the user associates the egress context of a spoke-sdp to the named queue-group. In such a case, the PW packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the PW packet is forwarded on. This queue can be a queue-group queue or the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. This is the existing implementation and default behavior for a PW packet.
2. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name exists but the policer-id and/or the queue-id is not defined in the queue-group template, the association is failed at the time the user associates the egress context of a spoke-sdp to the named queue-group. In such a case, the PW packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the PW packet is forwarded on.
3. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name exists and the policer-id or policer-id plus queue-id exist, it is not required to check that

an instance of that queue-group exists in all egress network ports that have network IP interfaces. The handling of this is dealt with in the data path as follows:

- When a PW packet for that FC is forwarded and an instance of the referenced queue-group name exists on that egress port, the packet is processed by the queue-group policer and will then be fed to the queue-group queue.
- When a PW packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the PW packet will be fed directly to the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

4. If a network QoS policy is applied to the egress context of a PW, any PW FC that is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the corresponding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

When the queue-group name the PW is redirected to exists and the redirection succeeds, the marking of the packet's DEI/dot1p/DSCP and the tunnel's DEI/dot1p/DSCP/EXP is performed according to the relevant mappings of the {FC, profile} in the egress context of the network QoS policy applied to the PW. This is true regardless of whether an instance of the queue-group exists or not on the egress port the PW packet is forwarded to. If the packet's profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet's DEI/dot1p and the tunnel's DEI/dot1p/EXP, and the DSCP/prec will be remarked if **enable-dscp-prec-marking** is enabled under the policer.

When the queue-group name the PW is redirected does not exist, the redirection command is failed. In this case, the marking of the packet's DEI/dot1p/DSCP and the tunnel's DEI/dot1p/DSCP/EXP fields is performed according to the relevant commands in the egress context of the network QoS policy applied to the network IP interface to which the PW packet is forwarded.

The **no** version of this command removes the redirection of the PW to the queue-group.

Parameters

network-policy-id

Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 to 65535

queue-group-name

Specifies the name of the queue group template up to 32 characters in length.

instance-id

Specifies the optional identification of a specific instance of the queue-group.

Values 1 to 65535

Platforms

7705 SAR Gen 2

qos

Syntax

qos *network-policy-id* **port-redirect-group** *queue-group-name* [**instance** *instance-id*]

no qos [*network-policy-id*]

Context

[Tree] (config>service>ies>if>spoke-sdp>egress qos)

[Tree] (config>service>vprn>if>spoke-sdp>egress qos)

Full Context

configure service ies interface spoke-sdp egress qos

configure service vprn interface spoke-sdp egress qos

Description

This command is used to redirect pseudowire packets to an egress port queue-group for the purpose of shaping.

The egress pseudowire shaping provisioning model allows the mapping of one or more pseudowires to the same instance of queues, or policers and queues, which are defined in the queue-group template.

Operationally, the provisioning model consists of the following steps:

1. Create an egress queue-group template and configure queues only or policers and queues for each FC that needs to be redirected.
2. Apply the queue-group template to the network egress context of all ports where there exists a network IP interface on which the pseudowire packets can be forwarded. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.
3. Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue-group in the egress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.
4. Apply this network QoS policy to the egress context of a spoke-SDP inside a service or to the egress context of a pseudowire template and specify the redirect queue-group name.

One or more spoke-SDPs can have their FCs redirected to use queues only or queues and policers in the same queue-group instance.

The following are the constraints and rules of this provisioning model:

1. When a pseudowire FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name does not exist, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface on which the pseudowire packet is forwarded. This queue can be a queue-group queue, or the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. This is the existing implementation and default behavior for a pseudowire packet.
2. When a pseudowire FC is redirected to use a queue or a policer, and a queue in a queue-group and the queue-group name exists, but the policer-id and/or the queue-id is not defined in the queue-group

template, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the pseudowire packet is forwarded on.

3. When a pseudowire FC is redirected to use a queue, or a policer and a queue in a queue-group, and the queue-group name exists and the policer-id or policer-id plus queue-id exist, it is not required to check that an instance of that queue-group exists in all egress network ports which have network IP interfaces. The handling of this is dealt with in the data path as follows:
   a. When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name exists on that egress port, the packet is processed by the queue-group policer and will then be fed to the queue-group queue.
   b. When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the pseudowire packet will be fed directly to the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.
4. If a network QoS policy is applied to the egress context of a pseudowire, any pseudowire FC, which is not explicitly redirected in the network QoS policy, will have the corresponding packets feed directly the corresponding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

When the queue-group name the pseudowire is redirected to exists and the redirection succeeds, the marking of the packet DEI/dot1p/DSCP and the tunnel DEI/dot1p/DSCP/EXP is performed; according to the relevant mappings of the (FC, profile) in the egress context of the network QoS policy applied to the pseudowire. This is true regardless, whether an instance of the queue-group exists or not on the egress port to which the pseudowire packet is forwarded. If the packet profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet DEI/dot1p and the tunnel DEI/dot1p/EXP, and the DSCP/prec will be remarked if e**nable-dscp-prec-marking** is enabled under the policer.

When the queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the marking of the packet DEI/dot1p/DSCP and the tunnel DEI/dot1p/DSCP/EXP fields is performed according to the relevant commands in the egress context of the network QoS policy applied to the network IP interface to which the pseudowire packet is forwarded.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

Parameters

network-policy-id

Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 to 65535

port-redirect-group *queue-group-name*

This optional parameter specifies that the *queue-group-name* will be used for all egress forwarding class redirections within the network QoS policy ID. The specified *queue-group-name* must exist as a port egress queue group on the port associated with the IP interface.

egress-instance *instance-id*

Specifies the identification of a specific instance of the queue-group.

Values 1 to 16384

Platforms

7705 SAR Gen 2

qos

Syntax

qos *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*

no qos [*network-policy-id*]

Context

[Tree] (config>service>vprn>if>spoke-sdp>ingress qos)

[Tree] (config>service>ies>if>spoke-sdp>ingress qos)

Full Context

configure service vprn interface spoke-sdp ingress qos

configure service ies interface spoke-sdp ingress qos

Description

This command is used to redirect pseudowire packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.

The ingress pseudowire rate-limiting feature uses a policer in queue-group provisioning model. This model allows the mapping of one or more pseudowires to the same instance of policers, which are defined in a queue-group template.

Operationally, the provisioning model in the case of the ingress pseudowire shaping feature consists of the following steps:

1. Create an ingress queue-group template and configure policers for each FC that needs to be redirected and optionally, for each traffic type (unicast, broadcast, unknown, or multicast).
2. Apply the queue-group template to the network ingress forwarding plane where there exists a network IP interface to which the pseudowire packets can be received. This creates one instance of the template on the ingress of the FP. One or more instances of the same template can be created.
3. Configure FC-to-policer mappings together with the policer redirect to a queue-group in the ingress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.
4. Apply this network QoS policy to the ingress context of a spoke-SDP inside a service, or to the ingress context of a pseudowire template, and specify the redirect queue-group name.
5. One or more spoke-SDPs can have their FCs redirected to use policers in the same policer queue-group instance.

The following are the constraints and rules of this provisioning model when used in the ingress pseudowire rate-limiting feature:

1. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name does not exist, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will feed directly the

ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the FP.

2. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists but the policer-id is not defined in the queue-group template, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the FP.
3. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists and the policer-id is defined in the queue-group template, it is not required to check that an instance of that queue-group exists in all ingress FPs which have network IP interfaces. The handling of this is dealt with in the data path as follows:

When a pseudowire packet for that FC is received and an instance of the referenced queue-group name exists on that FP, the packet is processed by the policer and will then feed the per-FP ingress shared queues referred to as *policer-output-queues*.

When a pseudowire packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the pseudowire packets will be fed directly into the corresponding ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the FP.

- If a network QoS policy is applied to the ingress context of a pseudowire, any pseudowire FC which is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the FP.
- If no network QoS policy is applied to the ingress context of the pseudowire, then all packets of the pseudowire will feed:

the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the FP. This is the default behavior.

a queue-group policer followed by the per-FP ingress shared queues referred to as *policer-output-queues* if the ingress context of the network IP interface from which the packet is received is redirected to a queue-group (csc-policing). The only exceptions to this behavior are for packets received from a IES/VPRN spoke interface and from an R-VPLS spoke-SDP, which is forwarded to the R-VPLS IP interface. In these two cases, the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the FP is used.

When a pseudowire is redirected to use a policer queue-group, the classification of the packet for the purpose of FC and profile determination is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the pseudowire. This is true regardless of whether an instance of the named policer queue-group exists on the ingress FP on which the pseudowire packet is received. The user can apply a QoS filter matching the dot1.p in the VLAN tag corresponding to the Ethernet port encapsulation, the EXP in the outer label when the tunnel is an LSP, the DSCP in the IP header if the tunnel encapsulation is GRE, and the DSCP in the payload IP header if the user enabled the **ler-use-dscp** option and the pseudowire terminates in IES or VPRN service (spoke-interface).

When the policer queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface on which the pseudowire packet is received.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

Parameters

network-policy-id

Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 to 65535

fp- redirect-group *queue-group-name*

Specifies the name of the queue group template up to 32 characters in length.

ingress-instance *instance-id*

Specifies the identification of a specific instance of the queue-group.

Values 1 to 16384

Platforms

7705 SAR Gen 2

qos

Syntax

qos *network-policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*

qos name *network-policy-name* **port-redirect-group** *queue-group-name* **instance** *instance-id*

no qos [*network-policy-id*]

Context

[Tree] (config>service>pw-template>egress qos)

Full Context

configure service pw-template egress qos

Description

This command is used to redirect PW packets to an egress port queue-group for the purpose of shaping.

The egress PW shaping provisioning model allows the mapping of one or more PWs to the same instance of queues, or policers and queues, that are defined in the queue-group template.

Operationally, the provisioning model consists of the following steps:

1. Create an egress queue-group template and configure queues only, or policers and queues for each FC that needs to be redirected.
2. Apply the queue-group template to the network egress context of all ports where there exists a network IP interface that the PW packets can be forwarded on. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.
3. Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue- group in the egress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different PWs to different queue-group templates.

4. Apply this network QoS policy to the egress context of a spoke-SDP inside a service, or to the egress context of a PW template and specify the redirect queue-group name.

One or more spoke-SDPs can have their FCs redirected to use queues only, or queues and policers in the same queue-group instance.

The following are the constraints and rules of this provisioning model.

1. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name does not exist, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the PW packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the PW packet is forwarded on. This queue can be a queue-group queue or the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. This is the existing implementation and default behavior for a PW packet.
2. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name exists but the policer-id and/or the queue-id is not defined in the queue-group template, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the PW packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the PW packet is forwarded on.
3. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name exists and the policer-id or policer-id plus queue-id exist, it is not required to check that an instance of that queue-group exists in all egress network ports that have network IP interfaces. The handling of this is dealt with in the data path as follows:
   - When a PW packet for that FC is forwarded and an instance of the referenced queue-group name exists on that egress port, the packet is processed by the queue-group policer and will then be fed to the queue-group queue.
   - When a PW packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the PW packet will be fed directly to the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.
4. If a network QoS policy is applied to the egress context of a PW, any PW FC that is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the corresponding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

When the queue-group name the PW is redirected to exists and the redirection succeeds, the marking of the packet's DEI/dot1p/DSCP and the tunnel's DEI/dot1p/DSCP/EXP is performed according to the relevant mappings of the {FC, profile} in the egress context of the network QoS policy applied to the PW. This is true regardless of whether an instance of the queue-group exists or not on the egress port the PW packet is forwarded to. If the packet's profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet's DEI/dot1p and the tunnel's DEI/dot1p/EXP, and the DSCP/prec will be remarked if **enable-dscp-prec-marking** is enabled under the policer.

When the queue-group name the PW is redirected does not exist, the redirection command is failed. In this case, the marking of the packet's DEI/dot1p/DSCP and the tunnel's DEI/ dot1p/DSCP/EXP fields is performed according to the relevant commands in the egress context of the network QoS policy applied to the network IP interface the PW packet is forwarded to.

The **no** version of this command removes the redirection of the PW to the queue-group.

Parameters

network-policy-id

Specifies the network policy identification. The value uniquely identifies the policy on the system.

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **qos name** *network-policy-name* variant can be used in all configuration modes.

Values 1 to 65535

queue-group-name

Specifies the name of the queue group template up to 32 characters in length.

instance-id

Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

name network-policy-name

Specifies the network policy name. The value uniquely identifies the policy on the system, up to 64 characters.

Platforms

7705 SAR Gen 2

qos

Syntax

qos *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*]

no qos

Context

[Tree] (config>service>epipe>spoke-sdp>ingress qos)

[Tree] (config>service>vpls>spoke-sdp>ingress qos)

[Tree] (config>service>vpls>mesh-sdp>ingress qos)

Full Context

configure service epipe spoke-sdp ingress qos

configure service vpls spoke-sdp ingress qos

configure service vpls mesh-sdp ingress qos

Description

This command is used to redirect pseudowire packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.

The ingress pseudowire rate-limiting feature uses a policer in queue-group provisioning model. This model allows the mapping of one or more pseudowires to the same instance of policers, which are defined in a queue-group template.

Operationally, the provisioning model in the case of the ingress pseudowire shaping feature consists of the following steps:

1. Create an ingress queue-group template and configure policers for each FC that needs to be redirected and optionally, for each traffic type (unicast or multicast).
2. Apply the queue-group template to the network ingress forwarding plane where there exists a network IP interface to which the pseudowire packets can be received. This creates one instance of the template on the ingress of the FP. One or more instances of the same template can be created.
3. Configure FC-to-policer mappings together with the policer redirect to a queue-group in the ingress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.
4. Apply this network QoS policy to the ingress context of a spoke-SDP inside a service, or to the ingress context of a pseudowire template, and specify the redirect queue-group name.
5. One or more spoke-SDPs can have their FCs redirected to use policers in the same policer queue-group instance.

The following are the constraints and rules of this provisioning model when used in the ingress pseudowire rate-limiting feature:

1. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name does not exist, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the FP.
2. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists but the policer-id is not defined in the queue-group template, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the FP.
3. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists and the policer-id is defined in the queue-group template, it is not required to check that an instance of that queue-group exists in all ingress FPs which have network IP interfaces. The handling of this is dealt with in the data path as follows:

When a pseudowire packet for that FC is received and an instance of the referenced queue-group name exists on that FP, the packet is processed by the policer and will then feed the per-FP ingress shared queues referred to as *policer-output-queues*.

When a pseudowire packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the pseudowire packets will be fed directly into the corresponding ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the FP.

- If a network QoS policy is applied to the ingress context of a pseudowire, any pseudowire FC which is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the FP.
- If no network QoS policy is applied to the ingress context of the pseudowire, then all packets of the pseudowire will feed:

- the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the FP. This is the default behavior.
- a queue-group policer followed by the per-FP ingress shared queues referred to as *policer-output-queues* if the ingress context of the network IP interface from which the packet is received is redirected to a queue-group. The only exceptions to this behavior are for packets received from a IES/VPRN spoke interface and from an R-VPLS spoke-SDP, which is forwarded to the R-VPLS IP interface. In these two cases, the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the FP is used.

When a pseudowire is redirected to use a policer queue-group, the classification of the packet for the purpose of FC and profile determination is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the pseudowire. This is true regardless of whether an instance of the named policer queue-group exists on the ingress FP on which the pseudowire packet is received. The user can apply a QoS filter matching the dot1-p in the VLAN tag corresponding to the Ethernet port encapsulation, the EXP in the outer label when the tunnel is an LSP, the DSCP in the IP header if the tunnel encapsulation is GRE, and the DSCP in the payload IP header if the user enabled the **ler-use-dscp** option and the pseudowire terminates in IES or VPRN service (spoke-interface).

When the policer queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface on which the pseudowire packet is received.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

Parameters

network-policy-id

Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 to 65535

queue-group-name

Specifies the name of the queue group template up to 32 characters in length

instance-id

Specifies the identification of a specific instance of the queue-group

Values 1 to 16384

Platforms

7705 SAR Gen 2

qos

Syntax

qos *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*

qos name *network-policy-name* **fp-redirect-group** *queue-group-name* **instance** *instance-id*

no qos [*network-policy-id*]

Context

[Tree] (config>service>pw-template>ingress qos)

Full Context

configure service pw-template ingress qos

Description

This command is used to redirect pseudowire packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.

The ingress pseudowire rate-limiting feature uses a policer in queue-group provisioning model. This model allows the mapping of one or more pseudowires to the same instance of policers which are defined in a queue-group template.

Operationally, the provisioning model in the case of the ingress pseudowire shaping feature consists of the following steps:

1. Create an ingress queue-group template and configure policers for each FC which needs to be redirected and optionally for each traffic type (unicast or multicast).
2. Apply the queue-group template to the network ingress forwarding plane where there exists a network IP interface which the pseudowire packets can be received on. This creates one instance of the template on the ingress of the FP. One or more instances of the same template can be created.
3. Configure FC-to-policer mappings together with the policer redirect to a queue-group in the ingress context of a network QoS policy. No queue-group name is specified in this step which means the same network QoS policy can redirect different pseudowires to different queue-group templates.
4. Apply this network QoS policy to the ingress context of a spoke-sdp inside a service or to the ingress context of a pseudowire template and specify the redirect queue-group name.

One or more spoke-sdps can have their FCs redirected to use policers in the same policer queue-group instance.

The following are the constraints and rules of this provisioning model when used in the ingress pseudowire rate-limiting feature:

1. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name does not exist, the association is failed at the time the user associates the ingress context of a spoke-sdp to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the FP.
2. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists but the policer-id is not defined in the queue-group template, the association is failed at the time the user associates the ingress context of a spoke-sdp to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the FP.
3. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists and the policer-id is defined in the queue-group template, it is not required to check that an instance of that queue-group exists in all ingress FPs which have network IP interfaces. The handling of this is dealt with in the data path as follows:

- When a pseudowire packet for that FC is received and an instance of the referenced queue-group name exists on that FP, the packet is processed by the policer and will then feed the per-FP ingress shared queues referred to as "policer-output-queues”.
- When a pseudowire packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the pseudowire packets will be fed directly into the corresponding ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the FP.

4. If a network QoS policy is applied to the ingress context of a pseudowire, any pseudowire FC which is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the FP.
5. If no network QoS policy is applied to the ingress context of the pseudowire, then all packets of the pseudowire will feed:
   - the ingress network shared queue for the packet’s FC defined in the network-queue policy applied to the ingress of the FP. This is the default behavior.
   - a queue-group policer followed by the per-FP ingress shared queues referred to as "policer-output-queues” Good received is redirected to a queue-group. The only exceptions to this behavior are for packets received from a IES/VPRN spoke interface and from a R-VPLS spoke-sdp which is forwarded to the R-VPLS IP interface. In these two cases, the ingress network shared queue for the packet’s FC defined in the network-queue policy applied to the ingress of the FP is used. When a pseudowire is redirected to use a policer queue-group, the classification of the packet for the purpose of FC and profile determination is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the pseudowire. This is true regardless if an instance of the named policer queue-group exists on the ingress FP the pseudowire packet is received on. The user can apply a QoS filter matching the dot1p in the VLAN tag corresponding to the Ethernet port encapsulation, the EXP in the outer label when the tunnel is an LSP, the DSCP in the IP header if the tunnel encapsulation is GRE, and the DSCP in the payload’s IP header if the user enabled the ler-use-dscp option and the pseudowire terminates in IES or VPRN service (spoke-interface).

When the policer queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface the pseudowire packet is received on.

The no version of this command removes the redirection of the pseudowire to the queue-group.

Parameters

network-policy-id

Specifies the network policy identification. The value uniquely identifies the policy on the system.

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **qos name** *network-policy-name* variant can be used in all configuration modes.

Values 1 to 65535

queue-group-name

Specifies the name of the queue group template up to 32 characters in length.

instance-id

Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

name network-policy-name

Specifies the network policy name. The value uniquely identifies the policy on the system, up to 64 characters.

Platforms

7705 SAR Gen 2

qos

Syntax

qos *policy-id* [**shared-queuing** | **multipoint-shared**]

qos name *sap-ingress-policy-name* [**shared-queuing** | **multipoint-shared**]

no qos [*policy-id*]

Context

[Tree] (config>service>template>vpls-sap-template>ingress qos)

Full Context

configure service template vpls-sap-template ingress qos

Description

This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP) for the Epipe SAP template.

Parameters

policy-id

The ingress policy ID to associate with SAP or IP interface on ingress. The policy ID must already exist.

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **qos name** *sap-ingress-policy-name* variant can be used in all configuration modes.

Values 1 to 65535

shared-queuing

This keyword can only be specified on SAP ingress. Specify the ingress shared queue policy used by this SAP. When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

multipoint-shared

This keyword can only be specified on SAP ingress. Multipoint shared queuing is a superset of shared queuing. When multipoint shared queuing keyword is set, as well as the unicast packets, multipoint packets also used shared queues.

Ingress unicast service queues are mapped one-for-one with hardware queues and unicast packets traverse the ingress forwarding plane twice, similar to the shared-queuing option. In addition, the multipoint queues defined in the ingress SAP QoS policy are not created. Instead, multipoint packets (broadcast, multicast and unknown unicast destined) are treated to the same dual pass ingress forwarding plane processing as unicast packets.

When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

Values Multipoint or not present

Default Present (the queue is created as non-multipoint)

sap-ingress-policy-name

The SAP ingress QoS policy name to associate with the SAP on ingress, up to 64 characters.

Platforms

7705 SAR Gen 2

qos

Syntax

qos *sap-egress-policy-id*

qos name *sap-egress-policy-name*

no qos

Context

[Tree] (config>service>template>vpls-sap-template>egress qos)

Full Context

configure service template vpls-sap-template egress qos

Description

This command associates an existing QoS policy with the template.

Parameters

sap-egress-policy-id

The egress policy ID to associate with SAP or IP interface on egress. The policy ID must already exist.

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **qos name** *sap-egress-policy-name* variant can be used in all configuration modes.

Values 1 to 65535

sap-egress-policy-name

The SAP egress QoS policy name to associate with the SAP on egress, up to 64 characters.

Platforms

7705 SAR Gen 2

qos

Syntax

qos *policy-id* [**shared-queuing** | **multipoint-shared**] [**fp-redirect-group** *queue-group-name* **instance** *instance-id*]

no qos

Context

[Tree] (config>service>vpls>sap>ingress qos)

Full Context

configure service vpls sap ingress qos

Description

This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP).

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the policy-id does not exist, an error will be returned.

The qos command is used to associate both ingress and egress QoS policies. The qos command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a specified type will return an error.

When an ingress QoS policy is defined on IES ingress IP interface that is bound to a VPLS, the policy becomes associated with every SAP on the VPLS and augments the QoS policy that is defined on each SAP. Packets that are bridged will be processed using the policy defined on the VPLS SAP; packets that are routed will be processed using the policy defined in the IES IP interface-binding context.

By default, if no specific QoS policy is associated with the SAP for ingress or egress, the default QoS policy is used.

The no form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

Parameters

policy-id

The ingress policy ID to associate with SAP or IP interface on ingress. The policy ID must already exist.

Values 1 to 65535

shared-queuing

This keyword can only be specified on SAP ingress. Specify the ingress shared queue policy used by this SAP. When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

multipoint-shared

This keyword can only be specified on SAP ingress. Multipoint shared queuing is a superset of shared queuing. When multipoint shared queuing keyword is set, as well as the unicast packets, multipoint packets also used shared queues.

Ingress unicast service queues are mapped one-for-one with hardware queues and unicast packets traverse the ingress forwarding plane twice, similar to the shared-queuing option. In addition, the multipoint queues defined in the ingress SAP QoS policy are not created. Instead, multipoint packets (broadcast, multicast and unknown unicast destined) are treated to the same dual pass ingress forwarding plane processing as unicast packets.

When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

Values Multipoint or not present

Default Present (the queue is created as non-multipoint)

fp-redirect-group

Creates an instance of a named queue group template on the ingress forwarding plane of a specified IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command. The named queue group template can contain only policers. If it contains queues, then the command will fail.

queue-group-name

Specifies the name of the queue group template to be instantiated on the forwarding plane of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid ingress queue group template name, configured under the **configure qos queue-group-templates** context.

instance-id

Specifies the instance of the named queue group to be created on the IOM/IMM/XMA ingress forwarding plane.

Platforms

7705 SAR Gen 2

qos

Syntax

qos *network-policy-id*

qos *network-policy-id* **egress-port-redirect-group** *queue-group-name* **egress-instance** *instance-id* **ingress-fp-redirect-group** *queue-group-name* **ingress-instance** *instance-id*

qos *network-policy-id* **egress-port-redirect-group** *queue-group-name* **egress-instance** *instance-id*

qos *network-policy-id* **ingress-fp-redirect-group** *queue-group-name* **ingress-instance** *instance-id*

no qos

Context

[Tree] (config>service>vprn>nw-if qos)

Full Context

configure service vprn network-interface qos

Description

This command associates a network Quality of Service (QoS) policy with a network IP interface. Only one network QoS policy can be associated with an IP interface at one time. Attempts to associate a second QoS policy return an error.

Associating a network QoS policy with a network interface is useful for the following purposes:

- To apply classification rules for determining the forwarding-class and profile of ingress packets on the interface.
- To associate ingress packets on the interface with a queue-group instance applied to the ingress context of the interface's forwarding plane (FP). (This is only applicable to interfaces on IOM4 and later cards.) The referenced ingress queue-group instance may have policers defined in order to rate limit ingress traffic on a per-forwarding class (and forwarding type: unicast vs. multicast) basis.
- To perform 802.1p, DSCP, IP precedence and/or MPLS EXP re-marking of egress packets on the interface.
- To associate egress packets on the interface with a queue-group instance applied to the egress context of the interface's port. The referenced egress queue-group instance may have policers and/or queues defined in order to rate limit egress traffic on a per-forwarding class basis.

The **no** form of this command removes the network QoS policy association from the network IP interface, and the QoS policy reverts to the default.

Default

no qos

Parameters

network-policy-id

An existing network policy ID to associate with the IP interface.

Values 1 to 65535

port-redirect-group *queue-group-name*

This optional parameter specifies the egress queue-group used for all egress forwarding-class redirections specified within the network QoS policy ID. The specified *queue-group-name* must exist as an egress queue group applied to the egress context of the port associated with the IP interface.

egress-instance *instance-id*

Since multiple instances of the same egress queue-group can be applied to the same port this optional parameter is used to specify which particular instance to associate with this particular network IP interface.

Values 1 to 16384

fp- redirect-group *queue-group-name*

This optional parameter specifies the ingress queue-group used for all ingress forwarding-class redirections specified within the network QoS policy ID. The specified queue-group-name must exist as an ingress queue group applied to the ingress context of the forwarding plane associated with the IP interface.

ingress-instance *instance-id*

Since multiple instances of the same ingress queue-group can be applied to the same forwarding plane this parameter is required to specify which particular instance to associate with this particular network IP interface.

Values 1 to 16384

Platforms

7705 SAR Gen 2

qos

Syntax

qos *network-policy-id* **fp-redirect-group** *queue-group-name instance instance-id*

no qos

Context

[Tree] (config>service>vprn>network>ingress qos)

Full Context

configure service vprn network ingress qos

Description

This command is used to redirect unicast packets arriving on an automatically (using the **auto-bind-tunnel** command) or manually configured (using a **spoke-sdp** command, but not the **spoke-sdp** command under the VPRN IP interface) binding in a VPRN to a policer in an ingress forwarding plane queue-group for the purpose of rate-limiting.

For the policer to be used, the following must be true:

1. The configured queue group template name must be applied to the forwarding plane on which the ingress traffic arrives using the instance id specified.
2. The policer referenced in the FC-to-policer mappings in the ingress context of a network QoS policy must be present in the specified queue group template.

The command fails if the queue group template name does not exist or if the policer specified in the network QoS policy does not exist in the queue group template. If the queue group template name with the specified instance is not applied to the forwarding plane on which the VPRN binding unicast traffic arrives then this traffic uses the ingress network queues related to the network interface, however, the ingress classification is still based on the applied network QoS policy.

The unicast traffic can be redirected to a policer under the forwarding class **fp-redirect-group** command in the ingress section of a network QoS policy; any **fp-redirect-group multicast-policer, broadcast-policer** or **unknown-policer** commands are ignored for this traffic. Multicast traffic would use the ingress network queues or queue group related to the network interface.

Ingress classification is based on the configuration of the ingress section of the specified network QoS policy, noting that the dot1p and exp classification is based on the outer Ethernet header and MPLS label whereas the DSCP applies to the outer IP header if the tunnel encapsulation is GRE, or the DSCP in the first IP header in the payload if **ler-use-dscp** is enabled in the ingress section of the referenced network QoS policy.

When this command is applied, it overrides the QoS applied to the related network interfaces for unicast traffic arriving on bindings in that VPRN.

The **no** version of this command removes the redirection of VPRN binding traffic to the queue-group policers.

Parameters

network-policy-id

Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 to 65535

fp-redirect-group *queue-group-name*

Specifies the name of the queue group template up to 32 characters in length.

instance *instance-id*

Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

Platforms

7705 SAR Gen 2

qos

Syntax

qos *policy-id*

qos *policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*

no qos

Context

[Tree] (config>mirror>mirror-dest>sap>egress qos)

Full Context

configure mirror mirror-dest sap egress qos

Description

This command associates a QoS policy with an egress SAP for a mirrored service.

By default, no specific QoS policy is associated with the SAP for egress, so the default QoS policy is used.

The **no** form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

Default

qos 1

Parameters

policy-id

Specifies the QoS policy ID to associate with SAP for the mirrored service. The policy ID must already exist.

Values 1 to 65535

queue-group-name

Specifies the queue group redirect list policy name.

instance-id

Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

Platforms

7705 SAR Gen 2

qos

Syntax

qos *network-policy-id*

qos *network-policy-id* **egress-port-redirect-group** *queue-group-name* **egress-instance** *instance-id* **ingress-fp-redirect-group** *queue-group-name* **ingress-instance** *instance-id*

qos *network-policy-id* **egress-port-redirect-group** *queue-group-name* **egress-instance** *instance-id*

qos *network-policy-id* **ingress-fp-redirect-group** *queue-group-name* **ingress-instance** *instance-id*

no qos

Context

[Tree] (config>router>if qos)

Full Context

configure router interface qos

Description

This command associates a network Quality of Service (QoS) policy with a network IP interface. Only one network QoS policy can be associated with an IP interface at one time. Attempts to associate a second QoS policy return an error.

Associating a network QoS policy with a network interface is useful for the following purposes:

- To apply classification rules for determining the forwarding-class and profile of ingress packets on the interface.
- To associate ingress packets on the interface with a queue-group instance applied to the ingress context of the interface's forwarding plane (FP). The referenced ingress queue-group instance may have policers defined in order to rate limit ingress traffic on a per-forwarding class (and forwarding type: unicast vs. multicast) basis.
- To perform 802.1p, DSCP, IP precedence and/or MPLS EXP re-marking of egress packets on the interface.
- To associate egress packets on the interface with a queue-group instance applied to the egress context of the interface's port. The referenced egress queue-group instance may have policers and/or queues defined in order to rate limit egress traffic on a per-forwarding class basis.

The **no** form of this command removes the network QoS policy association from the network IP interface, and the QoS policy reverts to the default.

Default

no qos

Parameters

network-policy-id

Specifies an existing network policy ID to associate with the IP interface.

Values 1 to 65535

egress-port-redirect-group *queue-group-name*

This optional parameter specifies the egress queue-group used for all egress forwarding-class redirections specified within the network QoS policy ID. The specified *queue-group-name* must exist as an egress queue group applied to the egress context of the port associated with the IP interface.

egress-instance *instance-id*

Since multiple instances of the same egress queue-group can be applied to the same port this optional parameter is used to specify which instance to associate with this specific network IP interface.

Values 1 to 16384

ingress-fp- redirect-group *queue-group-name*

This optional parameter specifies the ingress queue-group used for all ingress forwarding-class redirections specified within the network QoS policy ID. The specified queue-group-name must exist as an ingress queue group applied to the ingress context of the forwarding plane associated with the IP interface.

ingress-instance *instance-id*

Since multiple instances of the same ingress queue-group can be applied to the same forwarding plane this parameter is required to specify which instance to associate with this specific network IP interface.

Values 1 to 16384

Platforms

7705 SAR Gen 2

24.4 qos-policy-id-range

qos-policy-id-range

Syntax

qos-policy-id-range start *policy-id* **end** *policy-id*

no qos-policy-id-range

Context

[Tree] (config>qos>md-auto-id qos-policy-id-range)

Full Context

configure qos md-auto-id qos-policy-id-range

Description

This command specifies the range of IDs used by SR OS to automatically assign an ID to QoS policies that are created in model-driven interfaces without an ID explicitly specified by the user or client.

A QoS policy created with an explicitly-specified ID cannot use an ID in this range. In classic CLI and SNMP, the ID range cannot be changed while objects exist inside the previous or new range. In MD interfaces, the range can be changed which will cause any previously existing objects in the previous ID range to be deleted and re-created using a new ID in the new range.

The **no** form of this command removes the range values.

See the **config>eth-cfm md-auto-id** command for further details.

Default

no qos-policy-id-range

Parameters

start ***policy-id***

Specifies the lower value of the ID range. The value must be less than or equal to the **end** value.

Values 1 to 65535

end ***policy-id***

Specifies the upper value of the ID range. The value must be greater than or equal to the **start** value.

Values 1 to 65535

Platforms

7705 SAR Gen 2

24.5 query-interval

query-interval

Syntax

query-interval *seconds*

no query-interval

Context

[Tree] (config>service>vpls>spoke-sdp>mld-snooping query-interval)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping query-interval)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping query-interval)

[Tree] (config>service>vpls>sap>mld-snooping query-interval)

[Tree] (config>service>vpls>sap>igmp-snooping query-interval)

[Tree] (config>service>vpls>igmp-snooping query-interval)

[Tree] (config>service>vpls>mesh-sdp>mld-snooping query-interval)

[Tree] (config>service>vpls>mld-snooping query-interval)

Full Context

configure service vpls spoke-sdp mld-snooping query-interval

configure service vpls mesh-sdp igmp-snooping query-interval

configure service vpls spoke-sdp igmp-snooping query-interval

configure service vpls sap mld-snooping query-interval

configure service vpls sap igmp-snooping query-interval

configure service vpls igmp-snooping query-interval

configure service vpls mesh-sdp mld-snooping query-interval

configure service vpls mld-snooping query-interval

Description

This command configures the IGMP query interval. If the **send-queries** command is enabled, this parameter specifies the interval between two consecutive general queries sent by the system on this SAP or SDP.

The configured query-interval must be greater than the configured query-response-interval.

If send-queries is not enabled on this SAP or SDP, the configured query-interval value is ignored.

Default

query-interval 125

Parameters

seconds

Specifies the time interval, in seconds, that the router transmits general host-query messages

Values 2 to 1024

Values **config>service>vpls>igmp-snooping**: 1 - 65535
config>service>vpls>sap>igmp-snooping: 2 - 1024

Platforms

7705 SAR Gen 2

query-interval

Syntax

query-interval *seconds*

no query-interval

Context

[Tree] (config>router>mld query-interval)

[Tree] (config>router>mld>if query-interval)

Full Context

configure router mld query-interval

configure router mld interface query-interval

Description

This command specifies the frequency that the querier router transmits general host-query messages. The host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.

The **no** form of this command reverts to the default value.

Default

query-interval 125

Parameters

seconds

The time frequency, in seconds, that the router transmits general host-query messages.

Values 2 to 1024

Platforms

7705 SAR Gen 2

query-interval

Syntax

query-interval *seconds*

no query-interval

Context

[Tree] (config>service>vprn>igmp query-interval)

Full Context

configure service vprn igmp query-interval

Description

This command specifies the frequency that the querier router transmits general host-query messages. The host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.

Default

query-interval 125

Parameters

seconds

The time frequency, in seconds, that the router transmits general host-query messages.

Values 2 to 1024

Platforms

7705 SAR Gen 2

query-interval

Syntax

query-interval *seconds*

no query-interval

Context

[Tree] (config>service>vprn>mld>if query-interval)

[Tree] (config>service>vprn>mld query-interval)

Full Context

configure service vprn mld interface query-interval

configure service vprn mld query-interval

Description

This command specifies the frequency that the querier router transmits general host-query messages. The host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.

Default

query-interval 125

Parameters

seconds

The time frequency, in seconds, that the router transmits general host-query messages.

Values 2 to 1024

Platforms

7705 SAR Gen 2

query-interval

Syntax

query-interval *seconds*

no query-interval

Context

[Tree] (config>router>igmp>if query-interval)

Full Context

configure router igmp interface query-interval

Description

This command specifies the frequency that the querier router transmits general host-query messages. The host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.

Default

query-interval 125

Parameters

seconds

Specifies the frequency, in seconds, that the router transmits general host-query messages.

Values 2 to 1024

Platforms

7705 SAR Gen 2

query-interval

Syntax

query-interval *seconds*

no query-interval

Context

[Tree] (config>service>pw-template>igmp-snooping query-interval)

Full Context

configure service pw-template igmp-snooping query-interval

Description

This command configures the IGMP query interval. If the **send-queries** command is enabled, this parameter specifies the interval between two consecutive general queries sent by the system on this SAP or SDP.

The configured query-interval must be greater than the configured query-response-interval.

If send-queries is not enabled on this SAP or SDP, the configured query-interval value is ignored.

Default

query-interval 125

Parameters

seconds

Specifies the time interval, in seconds, that the router transmits general host-query messages.

Values 2 to 1024

Platforms

7705 SAR Gen 2

24.6 query-last-listener-interval

query-last-listener-interval

Syntax

query-last-listener-interval *seconds*

no query-last-listener-interval

Context

[Tree] (config>router>mld>if query-last-listener-interval)

Full Context

configure router mld interface query-last-listener-interval

Description

This command configures the frequency at which the querier router sends a group-specific query messages, including the messages sent in response to leave-group messages and is only applicable when the group interface is configured with the **no sub-hosts-only** command. The shorter the interval, the faster the loss of the last listener of a group can be detected. If nothing is configured, by default, the **query-last-listener-interval** takes the value defined in the **config>router>mld** context or in the **config>service>vprn>mld** context.

The **no** form of this command reverts to the default value.

Default

query-last-listener-interval 1

Parameters

seconds

Specifies the frequency, in seconds, at which query messages are sent.

Values 1 to 1023

Platforms

7705 SAR Gen 2

query-last-listener-interval

Syntax

query-last-listener-interval *seconds*

no query-last-listener-interval

Context

[Tree] (config>service>vprn>mld query-last-listener-interval)

[Tree] (config>service>vprn>mld>if query-last-listener-interval)

Full Context

configure service vprn mld query-last-listener-interval

configure service vprn mld interface query-last-listener-interval

Description

This command specifies the Max Response Time inserted into Group-Specific Queries sent in response to Leave Group messages, and is also the amount of time between Group-Specific Query messages. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.

The **no** form of this command reverts to the default value.

Default

query-last-listener-interval 1

Parameters

seconds

Specifies the frequency, in seconds, at which Group-Specific-Query packets are transmitted.

Values 1 to 1023

Platforms

7705 SAR Gen 2

24.7 query-last-member-interval

query-last-member-interval

Syntax

query-last-member-interval *seconds*

no query-last-member-interval

Context

[Tree] (config>router>igmp query-last-member-interval)

[Tree] (config>router>igmp>if query-last-member-interval)

Full Context

configure router igmp query-last-member-interval

configure router igmp interface query-last-member-interval

Description

This command configures the frequency at which the querier sends group-specific query messages including messages sent in response to leave-group messages. The lower the interval, the faster the detection of the loss of the last member of a group.

Default

query-last-member-interval 1

Parameters

seconds

Specifies the frequency, in seconds, at which query messages are sent.

Values 1 to 1023

Platforms

7705 SAR Gen 2

query-last-member-interval

Syntax

query-last-member-interval *seconds*

Context

[Tree] (config>service>vprn>igmp query-last-member-interval)

Full Context

configure service vprn igmp query-last-member-interval

Description

This command configures the frequency at which the querier sends group-specific query messages including messages sent in response to leave-group messages. The lower the interval, the faster the detection of the loss of the last member of a group.

Default

query-last-member-interval 1

Parameters

seconds

Specifies the frequency, in seconds, at which query messages are sent.

Values 1 to 1023

Platforms

7705 SAR Gen 2

24.8 query-response-interval

query-response-interval

Syntax

query-response-interval *seconds*

Context

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping query-response-interval)

[Tree] (config>service>vpls>sap>mld-snooping query-response-interval)

[Tree] (config>service>vpls>sap>igmp-snooping query-response-interval)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping query-response-interval)

[Tree] (config>service>vpls>spoke-sdp>mld-snooping query-response-interval)

[Tree] (config>service>vpls>mesh-sdp>mld-snooping query-response-interval)

Full Context

configure service vpls spoke-sdp igmp-snooping query-response-interval

configure service vpls sap mld-snooping query-response-interval

configure service vpls sap igmp-snooping query-response-interval

configure service vpls mesh-sdp igmp-snooping query-response-interval

configure service vpls spoke-sdp mld-snooping query-response-interval

configure service vpls mesh-sdp mld-snooping query-response-interval

Description

This command configures the IGMP query response interval. If the **send-queries** command is enabled, this parameter specifies the maximum response time advertised in IGMPv2 or IGMPv3 queries.

The configured query response interval must be smaller than the configured query interval.

If **send-queries** is not enabled on this SAP or SDP, the configured query response interval value is ignored.

The **no** form of this command reverts to the default value.

Default

query-response-interval 10

Parameters

seconds

Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 to 1023

Platforms

7705 SAR Gen 2

query-response-interval

Syntax

query-response-interval *seconds*

no query-response-interval

Context

[Tree] (config>service>vprn>igmp query-response-interval)

Full Context

configure service vprn igmp query-response-interval

Description

This command configures the query response interval on when the group interface is configured with the **no sub-hosts-only** command. If nothing is configured, by default, the **query-response-interval** takes the value defined in the **config>router>igmp** (or **mld**) context or in the **config>service>vprn>igmp** (or **mld**) context.

The **no** form of this command reverts to the default value.

Default

query-response-interval 10

Parameters

seconds

Specifies the length of time to wait to receive a host-query message response from the host.

Values 1 to 1023

Platforms

7705 SAR Gen 2

query-response-interval

Syntax

query-response-interval *seconds*

no query-response-interval

Context

[Tree] (config>router>mld query-response-interval)

[Tree] (config>router>mld>if query-response-interval)

Full Context

configure router mld query-response-interval

configure router mld interface query-response-interval

Description

This command specifies how long the querier router waits to receive a response to a host-query message from a host.

The **no** form of this command reverts to the default value.

Default

query-response-interval 10

Parameters

seconds

Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 to 1023

Platforms

7705 SAR Gen 2

query-response-interval

Syntax

query-response-interval *seconds*

no query-response-interval

Context

[Tree] (config>router>igmp>if query-response-interval)

[Tree] (config>router>igmp query-response-interval)

Full Context

configure router igmp interface query-response-interval

configure router igmp query-response-interval

Description

This command specifies how long the querier router waits to receive a response to a host-query message from a host.

Default

query-response-interval 10

Parameters

seconds

Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 to 1023

Platforms

7705 SAR Gen 2

query-response-interval

Syntax

query-response-interval *seconds*

Context

[Tree] (config>service>vprn>mld>if query-response-interval)

[Tree] (config>service>vprn>mld query-response-interval)

Full Context

configure service vprn mld interface query-response-interval

configure service vprn mld query-response-interval

Description

This command specifies how long the querier router waits to receive a response to a host-query message from a host.

Default

query-response-interval 10

Parameters

seconds

Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 to 1023

Platforms

7705 SAR Gen 2

query-response-interval

Syntax

query-response-interval *seconds*

no query-response-interval

Context

[Tree] (config>service>pw-template>igmp-snooping query-response-interval)

Full Context

configure service pw-template igmp-snooping query-response-interval

Description

This command configures the IGMP query response interval. If the **send-queries** command is enabled, this parameter specifies the maximum response time advertised in IGMPv2/v3 queries.

The configured **query-response-interval** must be smaller than the configured **query-interval**.

If **send-queries** is not enabled on this SAP or SDP, the configured **query-response-interval** value is ignored.

Default

query-response-interval 10

Parameters

seconds

Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 to 1023

Platforms

7705 SAR Gen 2

24.9 query-src-ip

query-src-ip

Syntax

query-src-ip *ip-address*

no query-src-ip

Context

[Tree] (config>service>vpls>igmp-snooping query-src-ip)

Full Context

configure service vpls igmp-snooping query-src-ip

Description

This command configures the IP source address used in IGMP or MLD queries.

The **no** form of this command removes the IP address from this configuration.

Parameters

ip-address

Specifies an IPv4 address in the form of a.b.c.d or an IPv6 address in the following form:

x:x:x:x:x:x:x:x:x:x:x:d.d.d.d

where:

x - [0 to FF]

d - [0 to 255]

Platforms

7705 SAR Gen 2

query-src-ip

Syntax

query-src-ip *ipv6-address*

no query-src-ip

Context

[Tree] (config>service>vpls>mld-snooping query-src-ip)

Full Context

configure service vpls mld-snooping query-src-ip

Description

This command configures the IP source address used in MLD queries.

Parameters

ipv6-address

Specifies an IPv6 address in the following form:

x:x:x:x:x:x:x:x (eight 16-bit pieces)

Platforms

7705 SAR Gen 2

24.10 queue

queue

Syntax

queue *queue-id* [**create**]

no queue *queue-id*

Context

[Tree] (config>service>ies>if>sap>ingress>queue-override queue)

[Tree] (config>service>vpls>sap>ingress>queue-override queue)

[Tree] (config>service>vpls>sap>egress>queue-override queue)

[Tree] (config>service>ies>if>sap>egress>queue-override queue)

Full Context

configure service ies interface sap ingress queue-override queue

configure service vpls sap ingress queue-override queue

configure service vpls sap egress queue-override queue

configure service ies interface sap egress queue-override queue

Description

This command specifies the ID of the queue whose parameters are to be overridden.

The **no** form of this command removes the queue ID from the configuration.

Parameters

queue-id

Specifies the queue ID whose parameters are to be overridden.

Values 1 to 8 for SAP egress

1 to 32 for SAP ingress

create

Keyword used to create the queue ID. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id* [**create**]

no queue *queue-id*

Context

[Tree] (config>port>ethernet>access>egr>qgrp>qover queue)

[Tree] (config>port>eth>network>egr>qgrp>qover queue)

Full Context

configure port ethernet access egress queue-group queue-overrides queue

configure port ethernet network egress queue-group queue-overrides queue

Description

This command associates a queue for use in a queue group template. The defined queue-id acts as a repository for the default parameters for the queue. The template queue is created on each queue-group object which is created with the queue group template name. Each queue is identified within the template by a queue-id number. The template ensures that all queue groups created with the template's name will have the same queue-ids providing a uniform structure for the forwarding class redirection commands in the SAP egress QoS policies. The parameters within the template queue will be used as the default settings for each queue in the actual queue group. The queue parameters may be individually changed for each queue in each queue group using per queue overrides.

The **no** form of this command removes the queue-id from the configuration.

Parameters

queue-id

Specifies the queue ID.

Values 1 to 8

create

Mandatory when creating an entry.

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id* [**create**]

no queue *queue-id*

Context

[Tree] (config>port>ethernet>access>ing>qgrp>qover queue)

Full Context

configure port ethernet access ingress queue-group queue-overrides queue

Description

This command associates a queue for use in a queue group template. The defined queue-id acts as a repository for the default parameters for the queue. The template queue is created on each queue-group object which is created with the queue group template name. Each queue is identified within the template by a queue-id number. The template ensures that all queue groups created with the template's name will have the same queue-ids providing a uniform structure for the forwarding class redirection commands in the SAP egress QoS policies. The parameters within the template queue will be used as the default settings for each queue in the actual queue group. The queue parameters may be individually changed for each queue in each queue group using per queue overrides.

The **no** form of this command removes the queue-id from the configuration.

Parameters

queue-id

Specifies the queue ID.

Values 1 to 32

create

Mandatory when creating an entry.

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id* [**create**]

no queue *queue-id*

Context

[Tree] (config>service>epipe>sap>ingress>queue-override queue)

[Tree] (config>service>epipe>sap>egress>queue-override queue)

Full Context

configure service epipe sap ingress queue-override queue

configure service epipe sap egress queue-override queue

Description

This command specifies the ID of the queue whose parameters are to be overridden.

Parameters

queue-id

The queue ID whose parameters are to be overridden.

Values 1 to 32

create

This keyword is mandatory when creating a queue.

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id* [**create**]

no queue *queue-id*

Context

[Tree] (config>service>vprn>if>sap>egress>queue-override queue)

[Tree] (config>service>vprn>if>sap>ingress>queue-override queue)

Full Context

configure service vprn interface sap egress queue-override queue

configure service vprn interface sap ingress queue-override queue

Description

This command specifies the ID of the queue whose parameters are to be overridden.

Parameters

queue-id

Specifies the queue ID whose parameters are to be overridden.

Values 1 to 32

create

Keyword used to create the group override instance.

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id* [**group** *queue-group-name*]

no queue

Context

[Tree] (config>qos>sap-ingress>fc queue)

Full Context

configure qos sap-ingress fc queue

Description

This command overrides the default queue mapping for **fc** fc-name. The specified queue-id must exist within the policy before the mapping can be made. When the forwarding class mapping is executed, all traffic classified to fc-name on a SAP using this policy.

The **no** form of this command sets the queue-id back to the default queue for the forwarding class.

Default

queue 1

Parameters

queue-id

Specifies the SAP egress queue-id to be associated with the forwarding class. The queue-id must be an existing queue defined in sap-egress policy-id.

Values 1 — 8

Default 1

group ***queue-group-name***

This optional parameter is used to redirect the forwarding type within the forwarding class to the specified *queue-id* within the *queue-group-name*. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the specified queue group. The queue-group-name are configured in the **config>qos>queue-group-templates** egress and ingress contexts. This parameter is used when policy-based queue group redirection is desired. That is, the specific queue group to redirect to is named in the QoS policy.

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id* [**multipoint**] [*queue-type*] [*queue-mode*] [**create**]

no queue *queue-id*

Context

[Tree] (config>qos>sap-ingress queue)

Full Context

configure qos sap-ingress queue

Description

This command creates the context to configure an ingress SAP QoS policy queue.

Explicit definition of an ingress queue's type is supported. A single ingress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or best effort nature of the queue based on the forwarding classes mapped to it. As long as all forwarding classes mapped to the queue are expedited (nc, ef, h1, or h2), the queue is treated as an expedited queue by the hardware schedulers. When any best effort forwarding classes are mapped to the queue (be, af, l1, or l2), the queue is treated as best effort (be) by the hardware schedulers. The queue type must be defined at the time of queue creation within the policy.

The **queue** command allows the creation of multipoint queues. Only multipoint queues can receive ingress packets that need flooding to multiple destinations. By separating the unicast for multipoint traffic at service ingress and handling the traffic on separate multipoint queues, special handling of the multipoint traffic is possible. Each queue acts as an accounting and (optionally) shaping device offering precise control over potentially expensive multicast, broadcast, and unknown unicast traffic. Only the back-end support of multipoint traffic (between the forwarding class and the queue based on forwarding type) needs to be defined. The individual classification rules used to place traffic into forwarding classes are not affected. Queues must be defined as multipoint at the time of creation within the policy.

The multipoint queues are for multipoint-destined service traffic. Within non-multipoint services, such as Epipe services, all traffic is considered unicast due to the nature of the service type. Multicast and broadcast-destined traffic in an Epipe service will not be mapped to a multipoint service queue.

When an ingress SAP QoS policy with multipoint queues is applied to an Epipe SAP, the multipoint queues are not created. When an ingress SAP QoS policy with multipoint queues is applied to an IES SAP, a multipoint queue will be created when PIM is enabled on the IES interface.

Any billing or statistical queries about a multipoint queue on a non-multipoint service returns zero values. Any queue parameter information requested about a multipoint queue on a non-multipoint service returns the queue parameters in the policy. Buffers will not be allocated for multipoint queues on non-multipoint services. Buffer pool queries return zero values for actual buffers allocated and current buffer utilization.

The **no** form of this command removes the *queue-id* from the SAP ingress QoS policy and from any existing SAPs using the policy. If any forwarding class forwarding types are mapped to the queue, they revert to their default queues. When a queue is removed, any pending accounting information for each SAP queue created due to the definition of the queue in the policy is discarded.

Parameters

queue-id

The *queue-id* for the queue, expressed as an integer. The *queue-id* uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.

Values 1 to 32

queue-type

The **expedite**, **best-effort**, and **auto-expedite** queue types are mutually exclusive. Each defines the method that the system uses to service the queue from a hardware perspective. A keyword can be specified at the time the queue is created. If an attempt to change the keyword after the queue is initially defined, an error is generated.

Values expedite, best-effort, auto-expedite

expedite — This keyword ensures that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.

best-effort — This keyword ensures that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.

auto-expedite — This keyword allows the system to auto-define the way the queue is serviced by the hardware. When **auto-expedite** is defined on the queue, the queue is treated in an expedited manner when all forwarding classes mapped to the queue are configured as expedited types nc, ef, h1 or h2. When a single non-expedited forwarding class is mapped to the queue (be, af, l1, and l2), the queue automatically falls back to non-expedited status.

Default auto-expedite

multipoint

This optional keyword specifies that this *queue-id* is for multipoint forwarded traffic only. This *queue-id* can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. If forwarding class unicast traffic is mapped to a multipoint queue, an error is generated and no changes are made to the current unicast traffic queue mapping.

A queue must be created as multipoint. The **multipoint** designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the **multipoint** keyword, an error is generated and the command will not execute.

The **multipoint** keyword can be entered in the command line on a preexisting multipoint queue to edit *queue-id* parameters.

Default non-multipoint (unicast queue)

queue-mode

Specifies the mode in which the queue is operating. This attribute is associated with the queue at the time of creation and cannot be modified thereafter.

Values **profile-mode**: When the queue is operating in the profile mode (or the color aware mode), the queue tries to provide the appropriate bandwidth to the packets with different profiles. The profiles are assigned according to the configuration of the forwarding class or the sub-forwarding class.

priority-mode: The queue is capable of handling traffic differently with two distinct priorities. These priorities are assigned by the stages preceding the queueing framework in the system. In priority mode, the queue does not have the functionality to support the profiled traffic and in such cases the queue will have a degraded performance. However, the converse is not valid and a queue in-profile mode should be capable of supporting the different priorities of traffic.

Default priority-mode

create

Keyword creates an ingress SAP QoS policy queue.

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id* [{**group** *queue-group-name* [**instance** *instance-id*] | **port-redirect-group-queue**}]

no queue

Context

[Tree] (config>qos>sap-egress>fc queue)

Full Context

configure qos sap-egress fc queue

Description

This command overrides the default queue mapping for **fc** fc-name. The specified queue ID must exist within the policy before the mapping can be made. When the forwarding class mapping is executed, all traffic is classified to fc-name on a SAP using this policy.

The **no** form of this command sets the queue-id back to the default queue for the forwarding class (queue 1).

Default

no queue

Parameters

queue-id

Specifies the SAP egress queue-id to be associated with the forwarding class. The queue-id must be an existing queue defined in sap-egress policy-id.

Values 1 to 8

Default 1

queue-group-name

This optional parameter is used to redirect the forwarding type within the forwarding class to the specified *queue-id* within the *queue-group-name*. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the specified queue group. The queue-group-name are configured in the **config>qos>queue-group-templates** egress and ingress contexts. This parameter is used when policy-based queue group redirection is desired. That is, the specific queue group to redirect to is named in the QoS policy.

instance-id

This parameter is used to specify the specific instance of a queue group with template queue-group-name to which this queue should be redirected. This parameter is only valid for queue groups on egress ports where policy-based redirection is required.

Values 1 to 40960

Default 1

port-redirect-group-queue

This keyword is used to mark a given forwarding class queue for redirection to an egress queue group queue. This is only used when the specific queue group instance is assigned at the time the QoS policy is applied to the SAP. This redirection model is known as SAP-based redirection.

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id* [*queue-type*] [**create**]

no queue *queue-id*

Context

[Tree] (config>qos>sap-egress queue)

Full Context

configure qos sap-egress queue

Description

This command creates the context to configure an egress service access point (SAP) QoS policy queue.

Explicit definition of an egress queue's type is supported. A single egress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or best effort nature of the queue based on the forwarding classes mapped to it. As long as all forwarding classes mapped to the queue are expedited (nc, ef, h1, or h2), the queue is treated as an expedited queue by the hardware schedulers. When any best effort forwarding classes are mapped to the queue (be, af, l1, or l2), the queue is treated as best effort by the hardware schedulers. The queue type must be defined at the time of queue creation within the policy.

The **no** form of this command removes the *queue-id* from the SAP egress QoS policy and from any existing SAPs using the policy. If any forwarding class forwarding types are mapped to the queue, they revert to their default queues. When a queue is removed, any pending accounting information for each SAP queue created due to the definition of the queue in the policy is discarded.

Parameters

queue-id

The ID for the queue, expressed as an integer. The *queue-id* uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.

Values 1 to 32

queue-type

Specifies the method that system uses to service the queue from a hardware perspective. A keyword can be specified at the time the queue is created. If an attempt is made to change the keyword after the queue is initially defined, an error is generated.

Values **expedite** - Specifies that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.

best-effort - Specifies that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.

auto-expedite - Allows the system to auto-define the way the queue is serviced by the hardware. When **auto-expedite** is defined on the queue, the queue is treated in an expedited manner when all

forwarding classes mapped to the queue are configured as expedited types nc, ef, h1, or h2. When a single non-expedited forwarding class is mapped to the queue (be, af, l1, and l2), the queue automatically falls back to non-expedited status.

Default auto-expedite

create

Creates an entry for the queue.

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id* [**multipoint**] [*queue-type*] [**create**]

no queue *queue-id*

Context

[Tree] (config>qos>network-queue>fc queue)

Full Context

configure qos network-queue fc queue

Description

Commands in this context configure a QoS network-queue policy queue.

Explicit definition of an ingress queue's type status is supported. A single ingress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or best effort nature of the queue based on the forwarding classes mapped to it. As long as all forwarding classes mapped to the queue are expedited (nc, ef, h1, or h2), the queue is treated as an expedited queue by the hardware schedulers. When any best effort forwarding classes are mapped to the queue (be, af, l1, or l2), the queue is treated as best effort (be) by the hardware schedulers. The queue type must be defined at the time of queue creation within the policy.

The **queue** command allows the creation of multipoint queues. Only multipoint queues can receive ingress packets that need flooding to multiple destinations. By separating the unicast for multipoint traffic at service ingress and handling the traffic on separate multipoint queues, special handling of the multipoint traffic is possible. Each queue acts as an accounting and (optionally) shaping device offering precise control over potentially expensive multicast, broadcast, and unknown unicast traffic. Only the back-end support of multipoint traffic (between the forwarding class and the queue based on forwarding type) needs to be defined. The individual classification rules used to place traffic into forwarding classes are not affected. Queues must be defined as multipoint at the time of creation within the policy.

The multipoint queues are for multipoint traffic.

The multipoint queues are for multipoint-destined service traffic. Within non-multipoint services, such as Epipe services, all traffic is considered unicast due to the nature of the service type. Multicast and broadcast-destined traffic in an Epipe service will not be mapped to a multipoint service queue.

When a QoS policy with multipoint queues is applied to an Epipe or IES SAP, the multipoint queues are not created. Any billing or statistical queries about a multipoint queue on a non-multipoint service returns zero values. Any queue parameter information requested about a multipoint queue on a non-multipoint service returns the queue parameters in the policy. Buffers will not be allocated for multipoint queues on non-multipoint services. Buffer pool queries return zero values for actual buffers allocated and current buffer utilization.

The **no** form of this command removes the *queue-id* from the network-queue policy and from any existing SAPs using the policy. If any forwarding class forwarding types are mapped to the queue, they revert to their default queues. When a queue is removed, any pending accounting information for each SAP queue created due to the definition of the queue in the policy is discarded.

Parameters

queue-id

The *queue-id* for the queue, expressed as an integer. The *queue-id* uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.

Values 1 to 32

multipoint

This optional keyword specifies that this *queue-id* is for multipoint forwarded traffic only. This *queue-id* can only be used to forward multicast, broadcast, or unknown unicast ingress traffic.

A queue must be created as multipoint. The **multipoint** designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the **multipoint** keyword, an error is generated, and the command will not execute.

The **multipoint** keyword can be entered in the command line on a preexisting multipoint queue to edit *queue-id* parameters.

Default Non-multipoint (unicast queue)

queue-type

The **expedite**, **best-effort**, and **auto-expedite** queue types are mutually exclusive. Each defines the method that the system uses to service the queue from a hardware perspective. While parental virtual schedulers can be defined for the queue, they only enforce how the queue interacts for bandwidth with other queues associated with the same scheduler hierarchy. An internal mechanism that provides access rules when the queue is vying for bandwidth with queues in other virtual schedulers is also needed. A keyword must be specified at the time the queue is created in the network-queue policy. If an attempt is made to change the keyword after the queue is initially defined, an error is generated.

Values expedite, best-effort, auto-expedite

expedite — This keyword ensures that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.

best-effort — This keyword ensures that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.

auto-expedite — This keyword allows the system to auto-define the way the queue is serviced by the hardware. When **auto-expedite** is defined on the queue, the queue is treated in an expedited manner when all forwarding classes mapped to the queue are configured as expedited types nc, ef, h1 or h2. When a single non-expedited forwarding class is mapped to the queue (be, af, l1, and l2), the queue automatically falls back to non-expedited status.

Default auto-expedite

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id* [**multipoint**] [*queue-type*] [*queue-mode*] [**create**]

no queue *queue-id*

Context

[Tree] (config>qos>qgrps>ing>queue-group queue)

Full Context

configure qos queue-group-templates ingress queue-group queue

Description

This command creates a queue for use in a queue group template. When created, the defined queue-id acts as a repository for the default parameters for the queue. The template queue is created on each queue-group object that is created with the queue group template name. Each queue is identified within the template by a queue-id number. The template ensures that all queue groups created with the template name will have the same queue-ids providing a uniform structure for the forwarding class redirection commands in the SAP ingress QoS policies. The parameters within the template queue will be used as the default settings for each queue in the actual queue group. The queue parameters may be individually changed for each queue in each queue group using per queue overrides.

When a queue within a template is mapped by a forwarding class on any object, the queue may be edited, but not deleted.

The **no** form of this command removes a template queue from the queue group template. If the queue is specified as a forwarding class redirection target in any SAP ingress QoS policy, the command will fail.

Parameters

queue-id

This required parameter identifies the queue that will either be created or edited within the queue group template.

Values 1 to 32

multipoint

This optional keyword creates an ingress multipoint queue. Multipoint queues in a queue group may be used by ingress VPLS for forwarding types multicast, broadcast or unknown within a forwarding class. For ingress IES and VPRN access SAPs, only multicast is supported. Multipoint queues are only supported on ingress queue group templates.

Default non-multipoint (unicast queue)

queue-type

The queue types are mutually exclusive.

Values **expedite —** This keyword ensures that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.

best-effort — This keyword ensures that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.

Default best-effort

queue-mode

These keywords are optional and mutually exclusive when creating a new template queue. The keywords specify how the queue manages ingress explicitly profiled packets.

Values **profile-mode** — Overrides the default priority mode of the queue and allows the adoption of color aware profiling within the queue. Forwarding classes and subclasses may be explicitly defined as in-profile or out-of-profile. Out-of-profile classified packets bypass the CIR rate associated with the queue, reserving it for the undefined or in-profile classified packets. If the template queue is not defined as profile-mode and the packet redirected to the queue is explicitly out-of-profile based on the classification rules, the queues within-CIR bandwidth may be consumed by the packet.

priority-mode — Defines that the SAP ingress QoS policy priority classification result will be honored by the queue. Priority mode is the default mode of the queue. High-priority packets are allowed into the queue up to the MBS defined for the queue. Low-priority packets are discarded at the low-priority MBS threshold that is derived from applying the low drop-tail percentage to the queue's MBS.

create

Keyword used to create the queue ID instance.

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id*

no queue

Context

[Tree] (config>qos>qgrps>egr>qgrp>fc queue)

Full Context

configure qos queue-group-templates egress queue-group fc queue

Description

This command is used to map the forwarding class to the specified *queue-id*. The specified *queue-id* must exist within the egress queue group template. When a queue is defined in a forwarding class mapping, that queue cannot be deleted unless the forwarding class mapping is moved to another queue within the template. Other criteria may also exist preventing the queue from being deleted from the template such as an applied SAP egress QoS policy mapping to the queue.

Parameters

queue-id

The specified *queue-id* must exist within the egress queue group template.

Values 1 to 8

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id* [*queue-type*] [**create**]

no queue *queue-id*

Context

[Tree] (cfg>qos>qgrps>egr>queue-group queue)

Full Context

configure qos queue-group-templates egress queue-group queue

Description

This command creates a queue for use in a queue group template. When created, the defined *queue-id* acts as a repository for the default parameters for the queue. The template queue is created on each queue group object that is created with the queue group template name. Each queue is identified within the template by a queue ID. The template ensures that all queue groups created with the template name will have the same queue-ids providing a uniform structure for the forwarding class redirection commands in the SAP egress QoS policies. The parameters within the template queue will be used as the default settings for each queue in the actual queue group. The queue parameters may be individually changed for each queue in each queue group using per queue overrides.

Parameters

queue-id

Specifies the queue ID. The specified *queue-id* must exist within the egress queue group template.

Values 1 to 8

queue-type

Specifies the method that the system uses to service the queue from a hardware perspective.

Values expedite, best-effort

expedite — This keyword ensures that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.

best-effort — This keyword ensures that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.

Default best-effort

Platforms

7705 SAR Gen 2

queue

Syntax

[**no**] **queue** *queue-id*

Context

[Tree] (config>log>acct-policy>cr queue)

Full Context

configure log accounting-policy custom-record queue

Description

This command specifies the queue-id for which counters will be collected in this custom record. The counters that will be collected are defined in egress and ingress counters.

The **no** form of this command reverts to the default value.

Parameters

queue-id

Specifies the queue-id for which counters will be collected in this custom record.

Platforms

7705 SAR Gen 2

queue

Syntax

queue *queue-id* [**multipoint**] [*queue-type*] [**create**]

no queue *queue-id*

Context

[Tree] (config>qos>network-queue queue)

Full Context

configure qos network-queue queue

Description

This command enters the context to configure a QoS network-queue policy queue.

Explicit definition of an ingress queue's type status is supported. A single ingress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or best effort nature of the queue based on the forwarding classes mapped to it. As long as all forwarding classes mapped to the queue are expedited (nc, ef, h1, or h2), the queue is treated as an expedited queue by the hardware schedulers. When any best effort forwarding classes are mapped to the queue (be, af, l1, or l2), the queue is treated as best effort (be) by the hardware schedulers. The queue type must be defined at the time of queue creation within the policy.

The **queue** command allows the creation of multipoint queues. Only multipoint queues can receive ingress packets that need flooding to multiple destinations. By separating the unicast for multipoint traffic at service ingress and handling the traffic on separate multipoint queues, special handling of the multipoint traffic is possible. Each queue acts as an accounting and (optionally) shaping device offering precise control over potentially expensive multicast, broadcast, and unknown unicast traffic. Only the back-end support of multipoint traffic (between the forwarding class and the queue based on forwarding type) needs to be defined. The individual classification rules used to place traffic into forwarding classes are not affected. Queues must be defined as multipoint at the time of creation within the policy.

The multipoint queues are for multipoint traffic.

The multipoint queues are for multipoint-destined service traffic. Within non-multipoint services, such as Epipe services, all traffic is considered unicast due to the nature of the service type. Multicast and broadcast-destined traffic in an Epipe service will not be mapped to a multipoint service queue.

When a QoS policy with multipoint queues is applied to an Epipe or IES SAP, the multipoint queues are not created. Any billing or statistical queries about a multipoint queue on a non-multipoint service returns zero values. Any queue parameter information requested about a multipoint queue on a non-multipoint service returns the queue parameters in the policy. Buffers will not be allocated for multipoint queues on non-multipoint services. Buffer pool queries return zero values for actual buffers allocated and current buffer utilization.

The **no** form of this command removes the *queue-id* from the network-queue policy and from any existing SAPs using the policy. If any forwarding class forwarding types are mapped to the queue, they revert to their default queues. When a queue is removed, any pending accounting information for each SAP queue created due to the definition of the queue in the policy is discarded.

Parameters

queue-id

The *queue-id* for the queue, expressed as an integer. The *queue-id* uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.

Values 1 to 32

multipoint

This optional keyword specifies that this *queue-id* is for multipoint forwarded traffic only. This *queue-id* can only be used to forward multicast, broadcast, or unknown unicast ingress traffic.

A queue must be created as multipoint. The **multipoint** designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the **multipoint** keyword, an error is generated, and the command will not execute.

The **multipoint** keyword can be entered in the command line on a preexisting multipoint queue to edit *queue-id* parameters.

Default Non-multipoint (unicast queue)

queue-type

The **expedite**, **best-effort**, and **auto-expedite** queue types are mutually exclusive. Each defines the method that the system uses to service the queue from a hardware perspective. While parental virtual schedulers can be defined for the queue, they only enforce how the queue interacts for bandwidth with other queues associated with the same scheduler hierarchy. An internal mechanism that provides access rules when the queue is vying for bandwidth with queues in other virtual schedulers is also needed. A keyword must be specified at the time the queue is created in the network-queue policy. If an attempt is made to change the keyword after the queue is initially defined, an error is generated.

Values expedite, best-effort, auto-expedite

expedite — This keyword ensures that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.

best-effort — This keyword ensures that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.

auto-expedite — This keyword allows the system to auto-define the way the queue is serviced by the hardware. When **auto-expedite** is defined on the queue, the queue is treated in an expedited manner when all forwarding classes mapped to the queue are configured as expedited types nc, ef, h1 or h2. When a single non-expedited forwarding class is mapped to the queue (be, af, l1, and l2), the queue automatically falls back to non-expedited status.

Default auto-expedite

Platforms

7705 SAR Gen 2

24.11 queue-delay

queue-delay

Syntax

queue-delay *delay*

no queue-delay

Context

[Tree] (config>qos>qgrps>egr>qgrp>queue queue-delay)

Full Context

configure qos queue-group-templates egress queue-group queue queue-delay

Description

This command configures the target queue delay for packets forwarded through the queue. It is used to determine the related queue parameters based on the administrative PIR of the queue. This command and the **mbs** command are mutually exclusive.

In order to change between the **mbs** and **queue-delay** parameters, the current parameter must be removed before adding the new parameter; that is, changing from **mbs** to **queue-delay** requires a **no mbs** before the **queue-delay** is configured and changing from **queue-delay** to **mbs** requires a **no queue-delay** before the **mbs** is configured.

If **queue-delay** is configured for an egress queue group queue, it is not possible to override the MBS for that queue.

The **no** form of this command disables the determination of the queue parameters based on the queue delay.

Default

no queue-delay

Parameters

delay

Specifies the target queue delay in ms.

Values 0 to 5000 (decimal)

Platforms

7705 SAR Gen 2

24.12 queue-frame-based-accounting

queue-frame-based-accounting

Syntax

[**no**] **queue-frame-based-accounting**

Context

[Tree] (config>service>ies>if>sap>egress>agg-rate queue-frame-based-accounting)

Full Context

configure service ies interface sap egress agg-rate queue-frame-based-accounting

Description

This command enables frame-based accounting on all queues associated with the **agg-rate** context. Only supported on Ethernet ports. Packet byte offset settings are not included in the applied rate when queue frame based accounting is configured, however the offsets are applied to the statistics.

The **no** form of this command disables frame-based accounting.

Platforms

7705 SAR Gen 2

queue-frame-based-accounting

Syntax

[**no**] **queue-frame-based-accounting**

Context

[Tree] (config>service>epipe>sap>egress>agg-rate queue-frame-based-accounting)

Full Context

configure service epipe sap egress agg-rate queue-frame-based-accounting

Description

This command is used to enable (or disable) frame based accounting on all policers and queues associated with the agg-rate context.

The command is supported on Ethernet ports only.

Packet byte offset settings are not included in the applied rate when queue frame based accounting is configured; however the offsets are applied to the statistics.

Platforms

7705 SAR Gen 2

queue-frame-based-accounting

Syntax

[**no**] **queue-frame-based-accounting**

Context

[Tree] (config>service>vpls>sap>egress>agg-rate queue-frame-based-accounting)

Full Context

configure service vpls sap egress agg-rate queue-frame-based-accounting

Description

This command is used to enabled frame-based accounting on all policers and queues associated with the agg-rate context. Only supported on Ethernet ports. Packet byte offset settings are not included in the applied rate when queue frame based accounting is configured; however the offsets are applied to the statistics.

The **no** form of this command disables the-frame based accounting.

Platforms

7705 SAR Gen 2

queue-frame-based-accounting

Syntax

[**no**] **queue-frame-based-accounting**

Context

[Tree] (config>service>vprn>if>sap>egress>agg-rate queue-frame-based-accounting)

Full Context

configure service vprn interface sap egress agg-rate queue-frame-based-accounting

Description

This command is used to enabled (or disable) frame based accounting on all policers and queues associated with the agg-rate context. Only supported on Ethernet ports. Packet byte offset settings are not included in the applied rate when queue frame-based accounting is configured; the offsets are applied to the statistics.

Platforms

7705 SAR Gen 2

24.13 queue-group

queue-group

Syntax

queue-group *queue-group-name* **instance** *instance-id* [**create**]

no queue-group *queue-group-name* **instance** *instance-id*

Context

[Tree] (config>card>fp>ingress>access queue-group)

Full Context

configure card fp ingress access queue-group

Description

This command creates an instance of a named queue group template on the ingress forwarding plane of a given IOM/IMM. The queue-group-name and **instance** *instance-id* are mandatory parameters when executing the command.

The named queue group template can contain only policers. If it contains queues, then the command will fail.

The **no** form of this command deletes a specific instance of a queue group.

Parameters

queue-group-name

Specifies the name of the queue group template to be instantiated on the forwarding plane of the IOM/IMM, up to 32 characters. The queue-group-name must correspond to a valid ingress queue group template name, configured under **config>qos>queue-group-templates**.

instance-id

Specifies the instance of the named queue group to be created on the IOM/IMM ingress forwarding plane.

Values 1 to 65535

create

Keyword used to associate the queue group. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

queue-group

Syntax

queue-group *queue-group-name* **instance** *instance-id* [**create**]

no queue-group *queue-group-name* **instance** *instance-id*

Context

[Tree] (config>card>fp>ingress>network queue-group)

Full Context

configure card fp ingress network queue-group

Description

This command creates a queue-group instance in the network ingress context of a forwarding plane.

Only a queue-group containing policers can be instantiated. If the queue-group template contains policers and queues, the queues are not instantiated. If the queue-group contains queues only, the instantiation in the data path is failed.

One or more instances of the same policer queue-group name and/or a different policer queue-group name can be created on the network ingress context of a forwarding plane.

The queue-group-name must be unique within all network ingress and access ingress queue groups in the system. The queue-group instance-id must be unique within the context of the forwarding plane.

The **no** form of this command deletes the queue-group instance from the network ingress context of the forwarding plane.

Parameters

queue-group-name

Specifies the name of the queue group template up to 32 characters.

instance-id

Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

create

Keyword used to create the queue-group instance.

Platforms

7705 SAR Gen 2

queue-group

Syntax

queue-group *queue-group-name* **instance** *instance-id*

no queue-group

Context

[Tree] (config>port>ethernet>network>egress queue-group)

Full Context

configure port ethernet network egress queue-group

Description

This command configures a queue-group instance in the network egress context of a port.

Queue-groups containing queues only or policers and queues can be instantiated. When a port is a LAG, one instance of the queue-group is instantiated on each member link.

One or more instances of the same queue-group name and/or a different queue-group name can be created in the network egress context of a port.

The queue-group-name must be unique within all network egress and access egress queue groups in the system. The queue-group instance-id must be unique within the context of the port.

The **no** version of this command deletes the queue-group instance from the network egress context of the port.

Parameters

queue-group-name

Specifies the name of the queue group template up to 32 characters.

instance-id

Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

Platforms

7705 SAR Gen 2

queue-group

Syntax

[**no**] **queue-group** *queue-group-name* [**instance** *instance-id*] [**create**]

Context

[Tree] (config>port>ethernet>access>egr queue-group)

[Tree] (config>port>ethernet>access>ing queue-group)

Full Context

configure port ethernet access egress queue-group

configure port ethernet access ingress queue-group

Description

This command creates an ingress or egress queue group on an Ethernet port. A queue group is a collection of queues identified by a group name. Queue groups created on access ports are used as an alternative queue destination for SAPs.

Within a SAP, a forwarding class may be redirected from the local SAP queue to a port queue group queue. The forwarding classes from multiple SAPs may be redirected to the same queue group which can be used to minimize the number of per-SAP queues.

Queue groups may be created on both access and network oriented ports. When the port is in access mode, the queue groups must be created within the port access node.

Within the access node, queue groups are also configured as ingress or egress. Access ingress queue groups can only be used by ingress SAP forwarding classes and only a single ingress queue group per port is supported. Multiple access egress queue groups may be created on a single port and are used by egress SAP forwarding classes. The instance-id parameter identifies different instances of the same queue group template. Creating multiple queue groups with a different instance ID but the same queue group name results in separate queue groups being created on the port. The instance-id parameter is only valid for egress queue groups on access ports.

When the queue group is created in an ingress port context, the group-name must be an existing ingress queue group template. Similarly, queue groups created in an egress port context must have a group-name of an existing egress queue group template. Two ingress queue groups with the same name cannot be created on the same port. Two egress queue groups can only be created on the same port with the same queue group template name if they have different instance-id values.

The queues defined in the template are created on the queue group. The queue parameters within the template are used as the default queue parameters for each queue in the queue group. The default queue parameters for each queue may be overridden on the queue group with specific queue parameters.

Each queue group supports the application of a scheduler-policy for the purpose of managing the queues within the group into an aggregate SLA. The queues defined within the template may be configured with parent scheduler defining the mapping of a queue to one of the schedulers within the scheduler policy. Egress queue groups also support the **agg-rate** parameter and the queues in the egress template support the port-parent command. Each command is used for configuring egress port virtual scheduling behavior.

Each queue group allows the application of an accounting policy and the ability to enable and disable collecting statistics. The statistics are derived from the queue counters on each queue within the queue

group. The accounting policy defines which queue counters are collected and to which accounting file they will be written.

A queue group does not have an administrative shutdown or no shutdown command. A queue group is considered to be always on once created.

When creating a queue group, the system will attempt to allocate queue resources based on the queues defined in the queue group template. If the appropriate queue resources do not currently exist, the queue group will not be created. Ingress port queue groups do not support the shared-queuing or multipoint-shared queuing behavior.

When the queue group is created on a LAG (Link Aggregation Group), it must be created on the primary port member. The primary port member is the port with the lowest port ID based on the slot, MDA position and port number on the MDA. A queue group created on the primary LAG port will be automatically created on all other port members. If a new port is being added to a LAG with an existing queue group, the queue group must first be created on the port prior to adding the port to the LAG. If the LAG queue group has queue overrides, the queue overrides must also be defined on the port queue group prior to adding the port to the LAG.

A port queue group cannot be removed from the port when a forwarding class is currently redirected to the group. All forwarding class redirections must first be removed prior to removing the queue group.

Parameters

queue-group-name

The group-name parameter is required when executing the port queue-group command. The specified group-name must exist as an ingress or egress queue group template depending on the ingress or egress context of the port queue group. Only a single queue group may be created on an ingress port. Multiple queue groups may be created on an egress port.

instance-id

Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

create

Keyword used to associate the queue group. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

queue-group

Syntax

queue-group *queue-group-name* [**create**]

no queue-group *queue-group-name*

Context

[Tree] (config>qos>qgrps>egress queue-group)

[Tree] (config>qos>qgrps>ingress queue-group)

Full Context

configure qos queue-group-templates egress queue-group

configure qos queue-group-templates ingress queue-group

Description

This command creates a queue group template. The system does not maintain default queue groups or queue group templates. Each queue group template used in the system must be explicitly created.

The **no** form of this command removes the specified queue group template from the system. If the queue group template is currently in use by an ingress port, the command will fail. If *queue-group-name* does not exist, the command has no effect and does not return an error.

Parameters

queue-group-name

Specifies the name of the queue group template up to 32 characters. Each ingress queue group template must be uniquely named within the system. Multiple ingress queue group templates may not share the same name. An ingress and egress queue group template may share the same name.

create

Keyword used to create the queue group instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

24.14 queue-group-templates

queue-group-templates

Syntax

queue-group-templates

Context

[Tree] (config>qos queue-group-templates)

Full Context

configure qos queue-group-templates

Description

Commands in this context define ingress and egress queue group templates.

Platforms

7705 SAR Gen 2

24.15 queue-override

queue-override

Syntax

[**no**] **queue-override**

Context

[Tree] (config>service>vprn>if>sap>egress queue-override)

[Tree] (config>service>ies>if>sap>ingress queue-override)

[Tree] (config>service>vpls>sap>ingress queue-override)

[Tree] (config>service>vpls>sap>egress queue-override)

[Tree] (config>service>vprn>if>sap>ingress queue-override)

[Tree] (config>service>ies>if>sap>egress queue-override)

Full Context

configure service vprn interface sap egress queue-override

configure service ies interface sap ingress queue-override

configure service vpls sap ingress queue-override

configure service vpls sap egress queue-override

configure service vprn interface sap ingress queue-override

configure service ies interface sap egress queue-override

Description

Commands in this context configure override values for the specified SAP egress or ingress QoS queue. These values override the corresponding ones specified in the associated SAP egress QoS policy.

Platforms

7705 SAR Gen 2

queue-override

Syntax

[**no**] **queue-override**

Context

[Tree] (config>service>epipe>sap>egress queue-override)

[Tree] (config>service>epipe>sap>ingress queue-override)

Full Context

configure service epipe sap egress queue-override

configure service epipe sap ingress queue-override

Description

Commands in this context configure override values for the specified SAP egress or ingress QoS queue. These values override the corresponding ones specified in the associated SAP egress or ingress QoS policy. If the policy was created as a template policy, this command overrides the parameter and its description and queue parameters in the policy.

Platforms

7705 SAR Gen 2

24.16 queue-overrides

queue-overrides

Syntax

queue-overrides

Context

[Tree] (config>port>ethernet>access>ing>qgrp queue-overrides)

[Tree] (config>port>ethernet>network>egr>qgrp queue-overrides)

[Tree] (config>port>ethernet>access>egr>qgrp queue-overrides)

Full Context

configure port ethernet access ingress queue-group queue-overrides

configure port ethernet network egress queue-group queue-overrides

configure port ethernet access egress queue-group queue-overrides

Description

Commands in this context define optional queue parameter overrides for each queue within the queue group.

Platforms

7705 SAR Gen 2

24.17 queue-policy

queue-policy

Syntax

queue-policy *name*

no queue-policy

Context

[Tree] (config>card>fp>ingress>network queue-policy)

Full Context

configure card fp ingress network queue-policy

Description

This command specifies the network-queue policy which defines queue parameters such as CBS, high priority only burst size, MBS, CIR and PIR rates, as well as forwarding-class to queue mappings. The network-queue policy is defined in the **config>qos>network-queue** context.

Default

queue-policy default

Parameters

name

Specifies an existing network-queue policy name, up to 32 characters long.

Platforms

7705 SAR Gen 2

queue-policy

Syntax

queue-policy *name*

no queue-policy

Context

[Tree] (config>port>ethernet>network queue-policy)

Full Context

configure port ethernet network queue-policy

Description

This command specifies the existing network queue policy which defines queue parameters such as CBS, high priority only burst size, MBS, CIR and PIR rates, as well as forwarding-class to queue mappings. The network-queue policy is defined in the **config>qos>network-queue** context.

Default

queue-policy default

Parameters

name

Specifies an existing network-queue policy name. The name can be up to 32 characters.

Platforms

7705 SAR Gen 2

24.18 quiet-period

quiet-period

Syntax

quiet-period *seconds*

Context

[Tree] (config>port>ethernet>dot1x quiet-period)

Full Context

configure port ethernet dot1x quiet-period

Description

This command configures the period between two authentication sessions during which no EAPOL frames are sent by the router.

The **no** form of this command returns the value to the default.

Default

quiet-period 60

Parameters

seconds

Specifies the quiet period in seconds.

Values 1 to 3600

Platforms

7705 SAR Gen 2

24.19 quit

quit

Syntax

quit

Context

[Tree] (candidate quit)

Full Context

candidate quit

Description

This command exits the **edit-cfg** mode. The contents of the current candidate will not be deleted and the operator can continue editing the candidate later.

Platforms

7705 SAR Gen 2

25 r Commands – Part I

25.1 radius

radius

Syntax

[**no**] **radius**

Context

[Tree] (debug>router radius)

Full Context

debug router radius

Description

This command enables the debug router RADIUS context.

Platforms

7705 SAR Gen 2

radius

Syntax

radius [**create**]

no radius

Context

[Tree] (config>service>vprn>aaa>rmt-srv radius)

Full Context

configure service vprn aaa remote-servers radius

Description

This command creates the context to configure RADIUS authentication on the VPRN.

Implement redundancy by configuring multiple server addresses for each VPRN.

The **no** form of this command removes the RADIUS configuration.

Parameters

create

Keyword used to create the RADIUS context.

Platforms

7705 SAR Gen 2

radius

Syntax

radius [**detail**] [**hex**]

no radius

Context

[Tree] (debug radius)

Full Context

debug radius

Description

This command enables debugging for RADIUS connections.

The **no** form of the command disables the debug output.

Parameters

detail

Displays detailed output.

hex

Displays the packet dump in hex format.

Platforms

7705 SAR Gen 2

radius

Syntax

[**no**] **radius**

Context

[Tree] (config>system>security radius)

Full Context

configure system security radius

Description

This command creates the context to configure RADIUS authentication on the router.

Implement redundancy by configuring multiple server addresses for each router.

The **no** form of this command removes the RADIUS configuration.

Platforms

7705 SAR Gen 2

25.2 radius-accounting-policy

radius-accounting-policy

Syntax

radius-accounting-policy *policy-name*

no radius-accounting-policy

Context

[Tree] (config>service>vprn>if>sap>ipsec-gw radius-accounting-policy)

[Tree] (config>service>ies>if>sap>ipsec-gw radius-accounting-policy)

Full Context

configure service vprn interface sap ipsec-gw radius-accounting-policy

configure service ies interface sap ipsec-gw radius-accounting-policy

Description

This command configures the RADIUS accounting policy.

The **no** form of this command reverts to the default value.

Default

no radius-accounting-policy

Parameters

policy-name

Specifies the policy name, up to 32 characters.

Platforms

7705 SAR Gen 2

radius-accounting-policy

Syntax

radius-accounting-policy *name* [**create**]

no radius-accounting-policy *name*

Context

[Tree] (config>ipsec radius-accounting-policy)

Full Context

configure ipsec radius-accounting-policy

Description

This command specifies an existing RADIUS accounting policy to use to collect accounting statistics on this subscriber profile by RADIUS. This command is used independently of the **collect-stats** command.

Parameters

name

Specifies an existing RADIUS based accounting policy.

Platforms

7705 SAR Gen 2

25.3 radius-attr

radius-attr

Syntax

radius-attr type *attribute-type* [**extended-type** *attribute-ext-type*] [**transaction**]

radius-attr type *attribute-type* [**transaction**] {**address** | **hex** | **integer** | **string**} **value** *attribute-value*

radius-attr vendor *vendor-id* **type** *attribute-type* [**extended-type** *attribute-ext-type*] [**transaction**] [**encoding** *encoding-type*]

radius-attr vendor *vendor-id* **type** *attribute-type* [**extended-type** *attribute-ext-type*] [**transaction**] [**encoding** *encoding-type*] {**address** | **hex** | **integer** | **string**} **value** *attribute-value*

no radius-attr type *attribute-type* [**extended-type** *attribute-ext-type*]

no radius-attr type *attribute-type* [**extended-type** *attribute-ext-type*] {**address** | **hex** | **integer** | **string**} **value** *attribute-value*

no radius-attr vendor *vendor-id* **type** *attribute-type* [**extended-type** *attribute-ext-type*]

no radius-attr vendor *vendor-id* **type** *attribute-type* [**extended-type** *attribute-ext-type*] {**address** | **hex** | **integer** | **string**} [**value**] *attribute-value*

Context

[Tree] (debug>router>radius radius-attr)

Full Context

debug router radius radius-attr

Description

This command specifies the RADIUS attribute filter of command **debug router radius**.

Parameters

attribute-type

Specifies the RADIUS attribute type.

Values 1 to 255

attribute-ext-type

Specifies the RADIUS attribute extended type (RFC 6929).

Values 1 to 255

address

Specifies the value is a IPv4 or IPv6 address/prefix/subnet.

string

Specifies the value is a ASCII string.

integer

Specifies the value is a integer.

hex

Specifies the value is a binary string in hex format, such as "\0xAB01FE”.

attribute-value

Specifies the value of the RADIUS attribute.

Values

address <ipv4-address> | <ipv6-address> | <ipv6-prefix/prefix-length>

ipv4-address a.b.c.d

ipv6-address x:x:x:x:x:x:x:x (eight 16-bit pieces)

ipv6-prefix x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

ipv6-prefix-length [0 to 128]

hex	[0x0 to 0xFFFFFFFF (up to 506 hex nibbles)]
integer	[0 to 4294967295]
string	ascii-string (up to 253 characters)

transaction

Specifies that the system outputs both request and response packets in the same session even if the response packet does not include the filter attribute.

vendor-id

Specifies the vendor ID for the vendor specific attribute.

Values 0 to 16777215

encoding-type

Specifies the size of the vendor-type and vendor-length in bytes. It is a two digitals string: "xy", x is the size of vendor-type, range from 1 to 4; y is the size of vendor-length, range from 0 to 2; it is "11" by default.

Values type-size:1 to 4, length-size: 0 to 2

Platforms

7705 SAR Gen 2

25.4 radius-authentication-policy

radius-authentication-policy

Syntax

radius-authentication-policy *name*

no radius-authentication-policy

Context

[Tree] (config>service>ies>if>sap>ipsec-gw radius-authentication-policy)

[Tree] (config>service>vprn>if>sap>ipsec-gw radius-authentication-policy)

Full Context

configure service ies interface sap ipsec-gw radius-authentication-policy

configure service vprn interface sap ipsec-gw radius-authentication-policy

Description

This command configures the policy used for the IKEv2 remote-access tunnels terminated on the IPsec gateway. The **radius-authentication-policy** is defined under **config>ipsec** context.

Parameters

name

Specifies the name of an existing RADIUS authentication policy.

Platforms

7705 SAR Gen 2

radius-authentication-policy

Syntax

radius-authentication-policy *name* [**create**]

no radius-authentication-policy *name*

Context

[Tree] (config>ipsec radius-authentication-policy)

Full Context

configure ipsec radius-authentication-policy

Description

This command specifies the RADIUS authentication policy associated with this IPsec gateway.

Parameters

name

Specifies an existing RADIUS authentication policy.

Platforms

7705 SAR Gen 2

25.5 radius-coa-port

radius-coa-port

Syntax

radius-coa-port {*port-number*}

no radius-coa-port

Context

[Tree] (config>aaa radius-coa-port)

Full Context

configure aaa radius-coa-port

Description

This command configures the system-wide UDP port number that RADIUS is listening on for CoA and Disconnect messages.

The **no** form of this command reverts to the default.

Default

radius-coa-port 3799

Parameters

port-number

Specifies the UDP port number for RADIUS CoA and disconnect messages.

Values 1647, 1700, 1812, 3799

Platforms

7705 SAR Gen 2

25.6 radius-plcy

radius-plcy

Syntax

radius-plcy *name*

no radius-plcy

Context

[Tree] (config>port>ethernet>dot1x radius-plcy)

Full Context

configure port ethernet dot1x radius-plcy

Description

This command references the RADIUS policy to be used for 802.1x authentication. An 802.1x RADIUS policy must be configured (**config**>**system**>**security**>**dot1x**) before it is associated to a port. If the RADIUS policy ID does not exist, an error is returned. Only one 802.1x RADIUS policy can be associated with a port at a time.

The **no** form of this command removes the RADIUS policy association.

Default
no radius-plcy

Parameters
name
Specifies an existing 802.1x RADIUS policy name, up to 32 characters.

Platforms
7705 SAR Gen 2

25.7 radius-server

radius-server

Syntax
radius-server

Context
[Tree] (config>router radius-server)
[Tree] (config>service>vprn radius-server)

Full Context
configure router radius-server
configure service vprn radius-server

Description
Commands in this context configure the RADIUS server under router or VPRN service.

Platforms
7705 SAR Gen 2

25.8 radius-server-policy

radius-server-policy

Syntax
radius-server-policy *policy-name* [**create**]
no radius-server-policy *policy-name*

Context

[Tree] (config>aaa radius-server-policy)

Full Context

configure aaa radius-server-policy

Description

This command creates a radius-server-policy.

A RADIUS server policy can be used in

- radius-proxy, for application like EAP authentication for WIFI access
- authentication policy, for Enhanced Subscriber Management authentication
- RADIUS accounting policy, for Enhanced Subscriber Management accounting
- dynamic data service RADIUS accounting
- AAA route downloader

The **no** form of this command removes the policy name from the configuration.

Parameters

policy-name

Specifies the name of the radius-server-policy up to 32 characters.

create

Keyword used to create a radius-server-policy name. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

radius-server-policy

Syntax

radius-server-policy *radius-server-policy-name*

no radius-server-policy

Context

[Tree] (config>ipsec>rad-auth-plcy radius-server-policy)

[Tree] (config>ipsec>rad-acct-plcy radius-server-policy)

Full Context

configure ipsec radius-authentication-policy radius-server-policy

configure ipsec radius-accounting-policy radius-server-policy

Description

This command references an existing **radius-server-policy** (available under the **config>aaa** context) for use in subscriber management authentication and accounting.

When configured in an authentication-policy, following CLI commands are ignored in the policy to avoid conflicts:

- all commands in the radius-authentication-server context
- accept-authorization-change
- coa-script-policy
- accept-script-policy
- request-script-policy

When configured in a radius-accounting-policy, following CLI commands are ignored in the policy to avoid conflicts:

- all commands in the radius-accounting-server context
- acct-request-script-policy

The **no** form of this command removes the radius-server-policy reference from the configuration.

Default

no radius-server-policy

Parameters

radius-server-policy-name

Specifies the RADIUS server policy.

Platforms

7705 SAR Gen 2

radius-server-policy

Syntax

radius-server-policy *policy-name*

radius-server-policy auth *policy-name-auth*

radius-server-policy acct *policy-name-acct*

radius-server-policy auth *policy-name-auth* **acct** *policy-name-acct*

no radius-server-policy

Context

[Tree] (config>port>ethernet>dot1x radius-server-policy)

Full Context

configure port ethernet dot1x radius-server-policy

Description

This command configures the RADIUS policy with IPv4/IPv6 in base routing and VPRN. The current RADIUS policy can be found under the **configure>aaa>radius-server-policy** context.

The RADIUS servers for the policy are configured under **configure>router>radius-server** or **configure>service>vprn>radius-server** context.

The RADIUS policy is assigned under dot1x using the **radius-server-policy** command. When the RADIUS policy is configured, both authorization and accounting are performed via the same server.

The **no** form of this command allows authorization and accounting via different servers.

Default

no radius-server-policy

Parameters

policy-name

Specifies the RADIUS server policy, up to 32 characters.

The policy is configured under **configure>aaa>radius-server-policy**. When the policy name is configured, both authorization and accounting are done via this server.

policy-name-auth

Specifies the AAA RADIUS server policy for dot1x authorization only; up to 32 characters.

The policy is configured under **configure>aaa>radius-server-policy**. The policy name authorization is used if the user needs a different server for authorization.

policy-name-acct

Specifies the AAA RADIUS server policy for dot1x accounting only; up to 32 characters.

The policy is configured under **configure>aaa>radius-server-policy**. The policy name accounting is used if the user needs a different server for accounting.

Platforms

7705 SAR Gen 2

25.9 range

range

Syntax

range *encap-range* **sync-tag** *sync-tag*

no range *encap-range*

Context

[Tree] (config>redundancy>multi-chassis>peer>sync>port range)

Full Context

configure redundancy multi-chassis peer sync port range

Description

This command configures a range of encapsulation values.

Parameters

encap-range

Specifies a range of encapsulation values on a port to be synchronized with a multi-chassis peer.

Values

Dot1Q	start-tag-end-tag
start-tag	0 to 4094
end-tag	0 to 4094
QinQ	qtag1.start-qtag2-qtag1.end-qtag2-start-qtag1.*-end-qtag1.*
qtag1	1 to 4094
start-qtag1	1 to 4094
en-qtag1	1 to 4094
start-qtag2	0 to 4094
end-qtag2	0 to 4094

sync-tag

Specifies a synchronization tag up to 32 characters to be used while synchronizing this encapsulation value range with the multi-chassis peer.

Platforms

7705 SAR Gen 2

range

Syntax

[**no**] **range** *vlan-range*

Context

[Tree] (config>service>vpls>sap>managed-vlan-list range)

Full Context

configure service vpls sap managed-vlan-list range

Description

This command configures a range of VLANs on an access port that are to be managed by an existing management VPLS.

This command is only valid when the VPLS in which it is entered was created as a management VPLS, and when the SAP in which it was entered was created on an Ethernet port with encapsulation type of dot1q or qinq, or on a SONET/SDH port with encapsulation type of bcp-dot1q.

To modify the range of VLANs, first the new range should be entered and afterwards the old range removed.

The **no** form of this command removes the VLAN range from this configuration.

Parameters

vlan-range

Specifies the VLAN start value and VLAN end value. The end-vlan must be greater than start-vlan. The format is <start-vlan>-<end-vlan>.

Values start-vlan: 1 to 4094

end-vlan: 1 to 4094

Platforms

7705 SAR Gen 2

25.10 rapid-retransmit-time

rapid-retransmit-time

Syntax

rapid-retransmit-time *hundred-milliseconds*

no rapid-retransmit-time

Context

[Tree] (config>router>rsvp rapid-retransmit-time)

Full Context

configure router rsvp rapid-retransmit-time

Description

This command defines the value of the Rapid Retransmission Interval. It is used in the re-transmission mechanism to handle unacknowledged message_id objects and is based on an exponential back-off timer.

Re-transmission interval of a RSVP message with the same message_id = 2 * rapid-retransmit-time interval of time.

The node stops re-transmission of unacknowledged RSVP messages:

- If the updated back-off interval exceeds the value of the regular refresh interval.
- If the number of re-transmissions reaches the value of the **rapid-retry-limit** parameter, whichever comes first.

The Rapid Retransmission Interval must be smaller than the regular refresh interval configured in **config>router>rsvp>refresh-time**.

The **no** form of this command reverts to the default value.

Default

rapid-retransmit-time 5

Parameters

hundred-milliseconds

Specifies the rapid retransmission interval, in hundred-milliseconds (for example, enter "6" for a 600 millisecond retransmit time).

Values 1 to 100, in units of 100 ms.

Platforms

7705 SAR Gen 2

25.11 rapid-retry-limit

rapid-retry-limit

Syntax

rapid-retry-limit *number*

no rapid-retry-limit

Context

[Tree] (config>router>rsvp rapid-retry-limit)

Full Context

configure router rsvp rapid-retry-limit

Description

This command defines the value of the Rapid Retry Limit. This is used in the retransmission mechanism based on an exponential backoff timer in order to handle unacknowledged message_id objects. The RSVP message with the same message_id is retransmitted every 2 * rapid-retransmit-time interval of time. The node stops retransmission of unacknowledged RSVP messages whenever the updated backoff interval exceeds the value of the regular refresh interval or the number of retransmissions reaches the value of the **rapid-retry-limit** parameter, whichever comes first.

The **no** form of this command reverts to the default value.

Default

rapid-retry-limit 3

Parameters

number

Specifies the value of the Rapid Retry Limit.

Values 1 to 6, integer values

Platforms

7705 SAR Gen 2

25.12 rapid-update

rapid-update

Syntax

rapid-update [**l2-vpn**] [**mvpn-ipv4**] [**mvpn-ipv6**] [**mdt-safi**] [**evpn**] [**label-ipv4**] [**label-ipv6**] [**vpn-ipv4**] [**vpn-ipv6**] [**mcast-vpn-ipv4**] [**mcast-vpn-ipv6**]

no rapid-update

Context

[Tree] (config>router>bgp rapid-update)

Full Context

configure router bgp rapid-update

Description

This command enables and disables BGP rapid update for specified address families.

If rapid update is enabled for a set of address families, and a route belonging to a family in that set is received by the router and chosen for propagation to certain BGP peers, the remaining time on the MRAI timer of these peers is ignored and the route is transmitted immediately, along with all other pending routes for these peers (including routes of address families not specified in the **rapid-update** command).

The **rapid-update** command overrides the peer-level **min-route-advertisement** (**config>router>bgp min-route-advertisement**, **config>router>bgp>group min-route-advertisement**, **config>router>bgp>group>neighbor min-route-advertisement**) time and applies the minimum setting (0 seconds) to routes belonging to specified address families; routes of other address families continue to be advertised according to the session-level MRAI setting.

The **no** form of this command disables rapid update for all address families.

Default

no rapid-update

Parameters

l2-vpn

Specifies the BGP rapid update for the 12-byte Virtual Switch Instance identifier (VSI-ID) value consisting of the 8-byte route distinguisher (RD) followed by a 4-byte value.

mvpn-ipv4

Specifies BGP rapid update for the mvpn-ipv4 address family. The mvpn-pv4 address is a variable size value consisting of the 1-byte route type, 1-byte length and variable size that is route type specific. Route type defines encoding for the route type specific field. Length indicates the length in octets of the route type specific field.

mdt-safi

Specifies BGP rapid update for the mdt-safi address family. The address is a 16-byte value consisting of 12-byte route distinguisher (RD) followed by a 4-byte group address.

mvpn-ipv6

Specifies BGP rapid update for the mvpn-ipv6 address family.

evpn

Specifies BGP rapid update for the evpn address family by including or removing EVPN routes from the set of routes that can trigger rapid update.

label-ipv4

Includes or removes label-ipv4 routes from the set of routes that can trigger rapid update.

label-ipv6

Includes or removes label-ipv6 routes from the set of routes that can trigger rapid update.

vpn-ipv4

Includes or removes vpn-ipv4 routes from the set of routes that can trigger rapid update.

vpn-ipv6

Includes or removes vpn-ipv6 routes from the set of routes that can trigger rapid update.

mcast-vpn-ipv4

Includes or removes mcast-vpn-ipv4 routes from the set of routes that can trigger rapid update.

mcast-vpn-ipv6

Includes or removes mcast-vpn-ipv6 routes from the set of routes that can trigger rapid update.

Platforms

7705 SAR Gen 2

25.13 rapid-withdrawal

rapid-withdrawal

Syntax

[**no**] **rapid-withdrawal**

Context

[Tree] (config>service>vprn>bgp rapid-withdrawal)

Full Context

configure service vprn bgp rapid-withdrawal

Description

This command disables the delay (Minimum Route Advertisement) on sending BGP withdrawals. Normal route withdrawals may be delayed up to the minimum route advertisement to allow for efficient packing of BGP updates.

The **no** form of this command removes this command from the configuration and returns withdrawal processing to the normal behavior.

Default

no rapid-withdrawal

Platforms

7705 SAR Gen 2

rapid-withdrawal

Syntax

[**no**] **rapid-withdrawal**

Context

[Tree] (config>router>bgp rapid-withdrawal)

Full Context

configure router bgp rapid-withdrawal

Description

This command disables the delay (Minimum Route Advertisement) on sending BGP withdrawals. Normal route withdrawals may be delayed up to the minimum route advertisement to allow for efficient packing of BGP updates.

The **no** form of this command removes this command from the configuration and returns withdrawal processing to the normal behavior.

Default

no rapid-withdrawal

Platforms

7705 SAR Gen 2

25.14 rate

rate

Syntax

rate *kilobits-per-second*

no rate

Context

[Tree] (config>service>ies>if>sap>egress>agg-rate rate)

Full Context

configure service ies interface sap egress agg-rate rate

Description

This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered to be active on the context's object (SAP, subscriber, Vport, and so on).

The **no** form of this command removes an explicit rate value from the aggregate rate therefore returning it to its default value.

Parameters

kilobits-per-second

Specifies the rate limit for the SAP, in kilobits per second.

Values 1 to 6400000000, **max**

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>service>ies>if>sap>egress>queue-override>queue rate)

[Tree] (config>service>vpls>sap>ingress>queue-override>queue rate)

[Tree] (config>service>ies>if>sap>ingress>queue-override>queue rate)

[Tree] (config>service>vpls>sap>egress>queue-override>queue rate)

Full Context

configure service ies interface sap egress queue-override queue rate

configure service vpls sap ingress queue-override queue rate

configure service ies interface sap ingress queue-override queue rate

configure service vpls sap egress queue-override queue rate

Description

This command overrides specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile, then out-of-profile, packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (max, 0).

Default

rate max cir 0

Parameters

pir-rate

Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be configured as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 6400000000, **max**

Default max

cir-rate

Overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and must be configured as a positive integer.

Values 0 to 6400000000, **max**

Default 0

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>service>vpls>sap>egress>sched-override>scheduler rate)

Full Context

configure service vpls sap egress scheduler-override scheduler rate

Description

This command overrides specific attributes of the specified scheduler rate. The **rate** command defines the maximum bandwidth that the scheduler can offer its policers, child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler because of insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler assumes that an infinite amount of bandwidth is available and allow all child queues and schedulers to operate at their maximum rates.

The **no** form of this command returns the scheduler's PIR and CIR parameters to the value configured in the applied scheduler policy.

Default

rate max cir sum

Parameters

pir-rate

Specifies the PIR rates. The **pir** parameter accepts a step multiplier value that specifies the multiplier used to determine the PIR rate at which the queue operates. A value of 0 to 100000000 or the keyword **max** is accepted. Any other value results in an error without modifying the current PIR rate.

To calculate the actual PIR rate, the rate described by the queue's **rate** is multiplied by the *pir-rate*.

The SAP ingress context for PIR is independent of the defined forwarding class (fc) for the queue. The default **pir** and definable range is identical for each class. The PIR in effect for a queue defines the maximum rate at which the queue is allowed to forward packets in a given second, shaping the queue's output.

The PIR parameter for SAP ingress queues do not have a negate (**no**) function. To return the queue's PIR rate to the default value, that value must be specified as the PIR value.

Values 1 to 6400000000, max

Default max

cir-rate

Specifies the CIR rate. The **cir** parameter accepts a step-multiplier value that specifies the multiplier used to determine the CIR rate at which the queue operate. A value of 0 to 250 or the keyword max is accepted. Any other value results in an error without modifying the current CIR rate.

To calculate the actual CIR rate, the rate described by the **rate pir** *pir-rate* is multiplied by the *cir-rate*. If the **cir** is set to max, then the CIR rate is set to infinity.

The context for CIR is dependent on the defined forwarding class (fc) for the queue. The default CIR and definable range is different for each class. The CIR in effect for a policer or queue defines both its profile (in or out) marking level as well as the relative importance compared to other queues for scheduling purposes during congestion periods.

Values 0 to 6400000000, max, sum

Default sum

Platforms

7705 SAR Gen 2

rate

Syntax

rate {*rate* | **max**} [**cir** {**max** | *rate*}]

no rate

Context

[Tree] (config>card>fp>ingress>network>qgrp>policer-over>plcr rate)

[Tree] (config>card>fp>ingress>access>qgrp>policer-over>plcr rate)

Full Context

configure card fp ingress network queue-group policer-override policer rate

configure card fp ingress access queue-group policer-override policer rate

Description

This command configures the policer's metering and optional profiling rates. The metering rate is used by the system to configure the policer's PIR leaky bucket's decrement rate while the profiling rate configures the policer's CIR leaky bucket's decrement rate. The decrement function empties the bucket while packets applied to the bucket attempt to fill it based on its packet size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches it's exceeded (CIR) or violate (PIR) threshold. The **cbs**, **mbs**, and **high-prio-only** commands are used to configure the policer's PIR and CIR thresholds.

If a packet arrives at the policer while the bucket's depth is less than the threshold associated with the packet, the packet is considered to be conforming to the bucket's rate. If the bucket depth is equal to or greater than the threshold, the packet is considered to be in the exception state. For the CIR bucket, the exception state is exceeding the CIR rate while the PIR bucket's exception state is violating the PIR bucket rate. If the packet is violating the PIR, the packet is marked red and will be discarded. If the packet is not red, it may be green or yellow based on the conforming or exceeding state from the CIR bucket.

When a packet is red neither the PIR nor the CIR bucket depths are incremented by the packets size. When the packet is yellow the PIR bucket is incremented by the packet size, but the CIR bucket is not. When the packet is green, both the PIR and CIR buckets are incremented by the packet size. This ensures that conforming packets impact the bucket depth while exceeding or violating packets do not.

The policer's **adaptation-rule** command settings are used by the system to convert the specified rates into hardware timers and decrement values for the policer's buckets.

By default, the policer's metering rate is **max** and the profiling rate is 0 kb/s (all packets out-of-profile).

The **rate** settings defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command reverts to the default metering and profiling rate of a policer.

Parameters

{*rate* | max}

Specifying the keyword **max** or an explicit *rate* parameter directly following the rate command is required and identifies the policer's metering rate for the PIR leaky bucket. When the policer is first created, the metering rate defaults to max. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When max is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to max.

Values **max** or 1 to 2000000000

cir {max | *rate*}

The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit *rate* parameter directly following the cir keyword is required and identifies the policer's profiling rate for the CIR leaky bucket. When the policer is first created, the profiling rate defaults to 0 kb/s. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When max is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the CIR used is equivalent to max.

Values **max** or 0 to 2000000000

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>port>ethernet>access>ing>qgrp>qover>q rate)

[Tree] (config>port>ethernet>access>egr>qgrp>qover>q rate)

Full Context

configure port ethernet access ingress queue-group queue-overrides queue rate

configure port ethernet access egress queue-group queue-overrides queue rate

Description

This command specifies the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can

transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile then out-of-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the **rate** is performed under the **hs-wrr-group** within the egress queue group template.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default

rate max cir 0 - The max default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The max value is mutually exclusive to the pir-rate value.

Parameters

pir-rate

Defines the administrative PIR rate, in kilobits per second, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed. Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 200000000, max

Default max

cir-rate

The **cir** parameter overrides the default administrative CIR used by the queue, in kilobits per second. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be given as a positive integer.

Values 0 to 200000000, max

Default 0

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>port>ethernet>access>egr>qgrp>sched-override>scheduler rate)

[Tree] (config>port>ethernet>access>ing>qgrp>sched-override>scheduler rate)

Full Context

configure port ethernet access egress queue-group scheduler-override scheduler rate

configure port ethernet access ingress queue-group scheduler-override scheduler rate

Description

This command can be used to override specific attributes of the specified scheduler rate. The rate command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler because of insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler assumes that an infinite amount of bandwidth is available and allow all child policers, queues, and schedulers to operate at their maximum rates.

The **no** form of this command returns the scheduler's PIR and CIR parameters to the value configured in the applied scheduler policy.

Parameters

pir-rate

Specifies the PIR rate, in kilobits per second. Any other value results in an error without modifying the current PIR rate.

Values 1 to 6400000000, **max**

cir-rate

Specifies the CIR rate, in kilobits per second. If the CIR is set to **max**, then the CIR rate is set to infinity. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers, policers, or queues.

Values 0 to 6400000000, **sum**, **max**

Platforms

7705 SAR Gen 2

rate

Syntax

rate *kilobits-per-second*

no rate

Context

[Tree] (config>service>epipe>sap>egress>agg-rate rate)

Full Context

configure service epipe sap egress agg-rate rate

Description

This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered to be active on the context's object (SAP, subscriber, Vport, and so on).

The **no** form of this command removes an explicit rate value from the aggregate rate returning it to its default value.

Parameters

kilobits-per-second

The enforced aggregate rate for all queues associated with the agg-rate context, in kilobits per second.

Values 1 to 6400000000, **max**

Platforms

7705 SAR Gen 2

rate

Syntax

rate {*rate* | **max**} [**cir** {*rate* | **max**}]

Context

[Tree] (config>service>epipe>sap>egress>policer-over>plcr rate)

[Tree] (config>service>epipe>sap>ingress>policer-over>plcr rate)

Full Context

configure service epipe sap egress policer-override policer rate

configure service epipe sap ingress policer-override policer rate

Description

This command within the SAP ingress and egress policer-overrides contexts is used to override the sap-ingress and sap-egress QoS policy configured rate parameters for the specified policer-id.

The **no** rate command is used to restore the policy defined metering and profiling rate to a policer.

Parameters

rate *rate*

Specifies the policer instance metering rate for the PIR leaky bucket, in kilobits per second. The integer value is multiplied by 1000 to derive the actual rate in bits per second.

Values 1 to 6400000000

cir *rate*

Specifies the overriding value for the policy-derived profiling rate of the policer, in kilobits per second. The integer value is multiplied by 1000 to derive the actual rate in bits per second.

Values 0 to 6400000000

max

Uses the maximum policer rate, equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR or CIR used is equivalent to **max**.

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>service>epipe>sap>egress>queue-override>queue rate)

[Tree] (config>service>epipe>sap>ingress>queue-override>queue rate)

Full Context

configure service epipe sap egress queue-override queue rate

configure service epipe sap ingress queue-override queue rate

Description

This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile and then out-of-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the **rate** is performed under the **hs-wrr-group** within the SAP egress QoS policy.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default

rate max cir 0

Parameters

pir-rate

Defines the administrative PIR rate, in kilobits per second, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 6400000000, **max**

Default max

cir-rate

The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has

not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and must be given as a positive integer. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers, policers, or queues.

Values 0 to 6400000000, **max**, **sum**

Default 0

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>service>epipe>sap>egress>sched-override>scheduler rate)

[Tree] (config>service>epipe>sap>ingress>sched-override>scheduler rate)

Full Context

configure service epipe sap egress scheduler-override scheduler rate

configure service epipe sap ingress scheduler-override scheduler rate

Description

This command can be used to override specific attributes of the specified scheduler rate. The **rate** command defines the maximum bandwidth that the scheduler can offer its child policers, queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child policers or queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child policers, queues, and schedulers to operate at their maximum rates.

The **no** form of this command returns the scheduler's PIR and CIR parameters to the values configured in the applied scheduler policy.

Parameters

pir-rate

The **pir** parameter accepts the **max** keyword or a value in kilobits per second. Any other value will result in an error without modifying the current PIR rate.

Values 1 to 6400000000, **max**

cir cir-rate

The **cir** parameter accepts a value in kilobits per second or the **max** keyword. Any other value will result in an error without modifying the current CIR rate.

If the **cir** parameter is set to **max**, then the CIR rate is set to infinity but bounded by the PIR rate.

The **sum** keyword specifies that the CIR will be used as the summed CIR values of the children schedulers, policers, or queues.

Values 0 to 6400000000, **max**, **sum**

Platforms

7705 SAR Gen 2

rate

Syntax

rate *kilobits-per-second*

no rate

Context

[Tree] (config>service>vpls>sap>egress>agg-rate rate)

Full Context

configure service vpls sap egress agg-rate rate

Description

This command defines the enforced aggregate rate for all queues associated with the **agg-rate** context. A rate must be specified for the **agg-rate** context to be considered active on the context's object (SAP, subscriber, Vport, and so on.).

The **no** form of this command removes an explicit rate value from the aggregate rate returning it to its default value.

Parameters

kilobits-per-second

The enforced aggregate rate for all queues associated with the **agg-rate** context, in kilobits per second.

Values 1 to 6400000000, **max**

Platforms

7705 SAR Gen 2

rate

Syntax

rate {*rate* | **max**} [**cir** {**max** | *rate*}]

Context

[Tree] (config>service>vpls>sap>ingress>policer-override>plcr rate)

[Tree] (config>service>vpls>sap>egress>policer-override>plcr rate)

Full Context

configure service vpls sap ingress policer-override policer rate

configure service vpls sap egress policer-override policer rate

Description

This command within the SAP ingress and egress policer-overrides contexts is used to override the sap-ingress and sap-egress QoS policy configured rate parameters for the specified policer-id.

The **no** form of this command removes the **rate** override so that the **rate** configured for the policer in the applied SAP egress QoS policy is used.

Parameters

{*rate* | max}

Specifying the keyword **max** or an explicit kilobits per second parameter directly following the rate override command is required and identifies the policer instance's metering rate for the PIR leaky bucket. The kilobits per second value must be expressed as an integer and defines the rate in kilobits per second. The integer value is multiplied by 1,000 to derive the actual rate in bits per second. When **max** is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to **max**.

Values 1 to 6400000000, **max**

cir {max | *rate*}

The optional **cir** keyword is used to override the policy derived profiling rate of the policer. Specifying the keyword **max** or an explicit kilobits per second parameter directly following

the **cir** keyword is required. The kilobits per second value must be expressed as an integer and defines the rate in kilobits per second. The integer value is multiplied by 1,000 to derive the actual rate in bits per second. When **max** is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the CIR used is equivalent to **max**.

Values 0 to 6400000000, **max**

Platforms

7705 SAR Gen 2

rate

Syntax

rate {*rate* | **max**} [**cir** {**max** | *rate*}]

Context

[Tree] (config>service>ies>if>sap>ingress>policer-override>plcr rate)

[Tree] (config>service>ies>if>sap>egress>policer-override>plcr rate)

Full Context

configure service ies interface sap ingress policer-override policer rate

configure service ies interface sap egress policer-override policer rate

Description

This command within the SAP ingress and egress policer-overrides contexts is used to override the sap-ingress and sap-egress QoS policy configured rate parameters for the specified policer-id.

The **no** form of the command removes the **rate** override so that the **rate** configured for the policer in the applied SAP egress QoS policy is used.

Parameters

{*rate* | max}

Specifying the keyword **max** or an explicit kilobits per second parameter directly following the rate override command is required and identifies the policer instance's metering rate for the PIR leaky bucket. The kilobits per second value must be expressed as an integer and defines the rate in kilobits per second. The integer value is multiplied by 1,000 to derive the actual rate in bits per second. When **max** is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to **max**.

Values 1 to 6400000000, **max**

cir {max | *rate*}

The optional **cir** keyword is used to override the policy derived profiling rate of the policer. Specifying the keyword **max** or an explicit kilobits per second parameter directly following the **cir** keyword is required. The kilobits per second value must be expressed as an integer and defines the rate in kilobits per second. The integer value is multiplied by 1,000 to derive the actual rate in bits per second. When **max** is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the CIR used is equivalent to **max**.

Values 0 to 6400000000, **max**

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>service>ies>if>sap>egress>sched-override>scheduler rate)

[Tree] (config>service>ies>if>sap>ingress>sched-override>scheduler rate)

Full Context

configure service ies interface sap egress scheduler-override scheduler rate

configure service ies interface sap ingress scheduler-override scheduler rate

Description

This command can be used to override specific attributes of the specified scheduler rate. The **rate** command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child queues and schedulers to operate at their maximum rates.

The **no** form of this command returns the scheduler's PIR and CIR parameters to the value configured in the applied scheduler policy.

Parameters

pir-rate

The **pir** parameter accepts a value in kilobits per second, or the keyword **max**. Any other value will result in an error without modifying the current PIR rate.

Values 1 to 6400000000, **max**

cir-rate

This parameter accepts a step-multiplier value that specifies the multiplier used to determine the CIR rate at which the queue will operate. A value in kilobits per second or the keywords **max** or **sum** is accepted. Any other value will result in an error without modifying the current CIR rate.

If the **cir** is set to max, then the CIR rate is set to infinity but is restricted by the PIR rate.

The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers, policers, or queues.

For **egress>sched-override>scheduler** and **ingress>sched-override>scheduler**:

Values 0 to 6400000000, **max**, **sum**

Platforms

7705 SAR Gen 2

rate

Syntax

rate *kilobits-per-second*

no rate

Context

[Tree] (config>service>vprn>if>sap>egress>agg-rate rate)

Full Context

configure service vprn interface sap egress agg-rate rate

Description

This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered to be active on the context's object.

The **no** form of this command removes an explicit rate value from the aggregate rate returning it to its default value.

Parameters

kilobits-per-second

Specifies the rate limit for the SAP, in kilobits per second.

Values 1 to 6400000000, max

Platforms

7705 SAR Gen 2

rate

Syntax

rate {*rate* | **max**} [**cir** {**max** | *rate*}]

Context

[Tree] (config>service>vprn>if>sap>egress>policer-override>plcr rate)

[Tree] (config>service>vprn>if>sap>ingress>policer-override>plcr rate)

Full Context

configure service vprn interface sap egress policer-override policer rate

configure service vprn interface sap ingress policer-override policer rate

Description

This command within the SAP ingress and egress policer-overrides contexts is used to override the sap-ingress and sap-egress QoS policy configured rate parameters for the specified policer-id.

The **no** form of this command restores the policy defined metering and profiling rate to a policer.

Parameters

{*rate* | max}

Specifying the keyword **max** or an explicit kilobits per second parameter directly following the rate override command is required and identifies the policer instance's metering rate for the PIR leaky bucket. The kilobits per second value must be expressed as an integer and defines the rate in kilobits per second. The integer value is multiplied by 1,000 to derive the actual rate in bits per second. When **max** is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to **max**.

Values 1 to 6400000000, **max**

cir {max | *rate*}

The optional **cir** keyword is used to override the policy derived profiling rate of the policer. Specifying the keyword **max** or an explicit kilobits per second parameter directly following the **cir** keyword is required. The kilobits per second value must be expressed as an integer and defines the rate in kilobits per second. The integer value is multiplied by 1,000 to

derive the actual rate in bits per second. When **max** is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the CIR used is equivalent to **max**.

Values 0 to 6400000000, **max**

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>service>vprn>if>sap>ingress>queue-override>queue rate)

[Tree] (config>service>vprn>if>sap>egress>queue-override>queue rate)

Full Context

configure service vprn interface sap ingress queue-override queue rate

configure service vprn interface sap egress queue-override queue rate

Description

This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the rate is performed under the **hs-wrr-group** within the SAP egress QoS policy.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default

rate max cir 0

Parameters

pir-rate

Defines the administrative PIR rate, in kb/s, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 6400000000, **max**

Default max

cir-rate

Defines the administrative CIR rate, in kb/s, for the queue. The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and must be given as a positive integer.

Values 0 to 6400000000, **max**

Default 0

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>service>vprn>if>sap>egress>sched-override>scheduler rate)

[Tree] (config>service>vprn>if>sap>ingress>sched-override>scheduler rate)

Full Context

configure service vprn interface sap egress scheduler-override scheduler rate

configure service vprn interface sap ingress scheduler-override scheduler rate

Description

This command can be used to override specific attributes of the specified scheduler rate. The **rate** command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child policers and queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child queues and schedulers to operate at their maximum rates.

The **no** form of this command returns the scheduler's PIR and CIR parameters to the value configured in the applied scheduler policy.

Parameters

pir-rate

Specifies the PIR rate for the scheduler. The **pir** parameter accepts a value in kb/s, or the **max** keyword. Any other value will result in an error without modifying the current PIR rate.

Values 1 to 6400000000, **max**

cir-rate

Specifies the CIR rate for the scheduler. The **cir** parameter accepts a value in kb/s, or the **max** or **sum** keywords. Any other value will result in an error without modifying the current CIR rate.

If the **cir** is set to **max**, then the CIR rate is set to infinity, but is limited by the *pir-rate*.

If the **cir** is set to **sum**, then the CIR rate is set to the summed CIR values of the children schedulers, policers, or queues.

Values 0 to 6400000000, **max**, **sum**

Platforms

7705 SAR Gen 2

rate

Syntax

rate *rate*

no rate

Context

[Tree] (config>qos>plcr-ctrl-plcy>tier>arbiter rate)

Full Context

configure qos policer-control-policy tier arbiter rate

Description

This command is used to define the maximum bandwidth an instance of the arbiter can receive from its parent tier 1 arbiter or the root arbiter. The arbiter instance enforces this limit by calculating the bandwidth each of its child policers should receive relative to their offered loads, parenting parameters, and individual rate limits, and using that derived rate as a child PIR decrement rate override. The override will not exceed the child policer's administrative rate limit and the aggregate of all the child PIR decrement rates will not exceed the specified arbiter rate limit.

The arbiter's policy defined rate value may be overridden at the SAP or sub-profile where the **policer-control-policy** is applied. Specifying an override prevents the arbiter from being removed from the policer control policy until the override is removed.

The **no** form of this command is used to remove a rate limit from the arbiter at the policer control policy level. The policy level rate limit for the arbiter will return to the default value of **max**. The **no rate** command has no effect on instances of the arbiter where a rate limit override has been defined.

Default

rate max

Parameters

rate

Enter an integer representing the rate limit in kilobits per second.

Values 1 to 6400000000, **max**

max

When **max** is specified, the arbiter does not enforce a rate limit on its child policers or arbiters other than the individual rate limits enforced at the child level.

Platforms

7705 SAR Gen 2

rate

Syntax

rate {**max** | *pir-rate*} [**cir** {**max** | *cir-rate*}]

Context

[Tree] (config>qos>sap-ingress>policer rate)

[Tree] (config>qos>sap-egress>policer rate)

Full Context

configure qos sap-ingress policer rate

configure qos sap-egress policer rate

Description

This command is used to configure the policer's metering and optional profiling rates. The metering rate is used by the system to configure the policer's PIR leaky bucket's decrement rate while the profiling rate configures the policer's CIR leaky bucket's decrement rate. The decrement function empties the bucket while packets applied to the bucket attempt to fill it based on each packet's size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches its exceed (CIR) or violate (PIR) threshold. The **cbs**, **mbs**, and **high-prio-only** commands are used to configure the policer's PIR and CIR thresholds.

If a packet arrives at the policer while the bucket's depth is less than the threshold associated with the packet, the packet is considered to be conforming to the bucket's rate. If the bucket depth is equal to or greater than the threshold, the packet is considered to be in the exception state. For the CIR bucket, the exception state is exceeding the CIR rate while the PIR bucket's exception state is violating the PIR bucket rate. If the packet is violating the PIR, the packet is marked red and will be discarded. If the packet is not red, it may be green or yellow, based on the conforming or exceeding state from the CIR bucket.

When a packet is red, neither the PIR nor CIR bucket depths are incremented by the packets size. When the packet is yellow, the PIR bucket is incremented by the packet size, but the CIR bucket is not. When the packet is green, both the PIR and CIR buckets are incremented by the packet size. This ensures that conforming packets impact the bucket depth while exceeding or violating packets do not.

The policer's **adaptation-rule** command settings are used by the system to convert the specified rates into hardware timers and decrement values for the policer's buckets.

By default, the policer's metering rate is **max** and the profiling rate is 0 kb/s (all packets out-of-profile).

The **rate** settings defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command is used to restore the default metering and profiling rate to a policer.

Parameters

{max | *pir-rate*}

Specifying the keyword **max** or an explicit *pir-rate* parameter directly following the rate command is required and identifies the policer's metering rate for the PIR leaky bucket. When the policer is first created, the metering rate defaults to max. The *pir-rate* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When max is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to max.

Values 1 to 6400000000, **max**

cir {max | *cir-rate*}

The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit *cir-rate* parameter directly following the cir keyword is required and identifies the policer's profiling rate for the CIR leaky bucket. When the policer is first created, the profiling rate defaults to 0 kb/s. The *cir-rate* value must be expressed

as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When **max** is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the CIR used is equivalent to max.

Values 0 to 6400000000, **max**

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*] [**fir** *fir-rate*]

rate *pir-rate* **police**

no rate

Context

[Tree] (config>qos>sap-ingress>queue rate)

Full Context

configure qos sap-ingress queue rate

Description

This command defines the administrative Peak Information Rate (PIR), the administrative Committed Information Rate (CIR), and the administrative Fair Information Rate (FIR) parameters for the queue.

The PIR defines the maximum rate that the queue can transmit packets through the switch fabric (for SAP ingress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. For SAP ingress, the CIR also defines the rate that packets are considered in-profile by the system, unless **cir-non-profiling** is configured. In-profile, then out-of-profile, packets are preferentially queued by the system at egress and at subsequent next-hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The FIR defines an additional rate at which the system prioritizes the queue over other queues competing for the same bandwidth above that used by the CIR.

The **rate** command can be executed at any time, altering the PIR, CIR, and FIR for all queues created through the association of the SAP ingress QoS policy with the *queue-id*.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0, 0).

Default

rate max cir 0 fir 0

Parameters

pir-rate

Defines the administrative PIR, in kilobits per second, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and the value must be given as a positive integer.

The actual PIR is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 6400000000, **max**

Default max

cir-rate

The **cir** parameter overrides the default administrative CIR, in kilobits per second, used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and the value must be given as a positive integer. The actual CIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 0 to 6400000000, **max**

Default 0

fir-rate

The **fir** parameter overrides the default administrative FIR, in kilobits per second, used by the queue. When the **rate** command is executed, an FIR setting is optional. When the **rate** command has not been executed or the **fir** parameter is not explicitly specified, the default FIR (0) is assumed.

Fractional values are not allowed and the value must be given as a positive integer. The actual FIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.

Values 0 to 6400000000, **max**

Default 0

police

Specifies that traffic feeding into the queue instance above the specified PIR rate will be dropped. When the **police** keyword is defined, only the PIR rate may be overridden.

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>qos>sap-egress>queue rate)

Full Context

configure qos sap-egress queue rate

Description

This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

When configured on an egress HSQ queue group queue, the **cir** keyword is ignored.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the rate is performed under the **hs-wrr-group** within the SAP egress QoS policy.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default

rate max cir 0

Parameters

pir-rate

Defines the administrative PIR rate, in kilobits per second, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 6400000000, **max**

Default max

cir-rate

The **cir** parameter overrides the default administrative CIR, in kilobits per second, used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and must be given as a positive integer.

Values 0 to 6400000000, **max**

Default 0

Platforms

7705 SAR Gen 2

rate

Syntax

rate *percent* [**cir** *percent*] [**fir** *percent*]

no rate

Context

[Tree] (config>qos>network-queue>queue rate)

Full Context

configure qos network-queue queue rate

Description

This command defines the administrative Peak Information Rate (PIR), the administrative Committed Information Rate (CIR), and the administrative Fair Information Rate (FIR) parameters for the queue.

The PIR defines the percentage that the queue can transmit packets through the switch fabric (for ingress queues) or out of an egress port (for egress queues). Defining a PIR does not necessarily guarantee that

the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available bandwidth.

The CIR defines the percentage at which the system prioritizes the queue over other queues competing for the same bandwidth.

The CIR can be used by the queue's **port-parent** commands **cir-level** and **cir-weight** parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent port scheduler.

The FIR defines an additional percentage at which the system prioritizes the queue over other queues competing for the same bandwidth above that used by the CIR percentage.

The **rate** command can be executed at any time, altering the PIR, CIR, and FIR for all queues created through the association of the network queue policy with the *queue-id*.

When configured on an egress HSQ queue group queue, the **cir** keyword is ignored.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the rate is performed under the **hs-wrr-group** within the network queue policy.

The **no** form of the command returns all queues created with the *queue-id* by association with the network queue policy to the default PIR, CIR, and FIR parameters.

Default

rate 100 cir 0 fir 0

Parameters

percent

Defines the percentage of the sum of the capacities of network and hybrid ports on that FP (taking into account any **ingress-rate** configuration) or egress port speed for the rate allowed for the queue. When the **rate** command is executed, a valid *percent* (PIR setting) must be explicitly defined. When the **rate** command has not been executed, the default PIR of **100** is assumed. Fractional values are not allowed, and the value must be given as a positive integer.

The actual PIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 100

Default 100

cir ***percent***

Defines the percentage of the sum of the capacities of network and hybrid ports on that FP (taking into account any **ingress-rate** configuration) or egress port speed for the CIR allowed for the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed, and the value must be given as a positive integer. The actual CIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 0 to 100

Default 0

fir *percent*

Defines the percentage of the sum of the capacities of network and hybrid ports on that FP (taking into account any **ingress-rate** configuration) or egress port speed for the FIR allowed for the queue. When the **rate** command is executed, a FIR setting is optional. When the **rate** command has not been executed or the **fir** parameter is not explicitly specified, the default FIR (0) is assumed. Fractional values are not allowed, and the value must be given as a positive integer. The actual FIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned. FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.

Values 0 to 100

Default 0

Platforms

7705 SAR Gen 2

rate

Syntax

rate {**max** | *pir-rate*} [**cir** {**max** | *cir-rate*}]

no rate

Context

[Tree] (config>qos>qgrps>ing>qgrp>policer rate)

[Tree] (config>qos>qgrps>egr>qgrp>policer rate)

Full Context

configure qos queue-group-templates ingress queue-group policer rate

configure qos queue-group-templates egress queue-group policer rate

Description

This command is used to configure the policer's metering and optional profiling rates. The metering rate is used by the system to configure the policer's PIR leaky bucket's decrement rate while the profiling rate configures the policer's CIR leaky bucket's decrement rate. The decrement function empties the bucket while packets applied to the bucket attempt to fill it based on each packet's size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches its exceed (CIR) or violate (PIR) threshold. The **cbs**, **mbs**, and **high-prio-only** commands are used to configure the policer's PIR and CIR thresholds.

If a packet arrives at the policer while the bucket's depth is less than the threshold associated with the packet, the packet is considered to be conforming to the bucket's rate. If the bucket depth is equal to or greater than the threshold, the packet is considered to be in the exception state. For the CIR bucket, the

exception state is exceeding the CIR rate while the PIR bucket's exception state is violating the PIR bucket rate. If the packet is violating the PIR, the packet is marked red and will be discarded. If the packet is not red, it may be green or yellow, based on the conforming or exceeding state from the CIR bucket.

When a packet is red, neither the PIR nor CIR bucket depths are incremented by the packets size. When the packet is yellow, the PIR bucket is incremented by the packet size, but the CIR bucket is not. When the packet is green, both the PIR and CIR buckets are incremented by the packet size. This ensures that conforming packets impact the bucket depth while exceeding or violating packets do not.

The policer's **adaptation-rule** command settings are used by the system to convert the specified rates into hardware timers and decrement values for the policer's buckets.

By default, the policer's metering rate is **max** and the profiling rate is 0 kb/s (all packets out-of-profile).

The **no** form of this command is used to restore the default metering and profiling rate to a policer.

Parameters

{max | *pir-rate*}

Specifying the keyword **max** or an explicit *pir-rate* parameter directly following the rate command is required and identifies the policer's metering rate for the PIR leaky bucket. When the policer is first created, the metering rate defaults to max. The *pir-rate* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When max is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to max.

Values **max**, 1 to 2000000000

cir {max | *cir-rate*}

The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit *cir-rate* parameter directly following the cir keyword is required and identifies the policer's profiling rate for the CIR leaky bucket. When the policer is first created, the profiling rate defaults to 0 kb/s. The *cir-rate* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When **max** is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the CIR used is equivalent to max.

Values **max,** 0 to 2000000000

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*] [**fir** *fir-rate*]

rate *pir-rate* **police**

no rate

Context

[Tree] (config>qos>queue-group-templates>ingress>queue-group>queue rate)

Full Context

configure qos queue-group-templates ingress queue-group queue rate

Description

This command defines the administrative Peak Information Rate (PIR), the administrative Committed Information Rate (CIR), and the administrative Fair Information Rate (FIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the switch fabric (for SAP ingress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. For SAP ingress, the CIR also defines the rate that packets are considered in-profile by the system, unless **cir-non-profiling** is configured. In-profile, then out-of-profile, packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The FIR defines an additional rate at which the system prioritizes the queue over other queues competing for the same bandwidth above that used by the CIR.

The **rate** command can be executed at any time, altering the PIR, CIR, and FIR for all queues created through the association of the ingress queue group template with the *queue-id*.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR, CIR, and FIR parameters (**max**, 0, 0).

Default

rate max cir 0 fir 0

Parameters

pir-rate

Defines the administrative PIR, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed. Fractional values are not allowed and the value must be given as a positive integer.

The actual PIR is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 2000000000 kb/s, **max**

Default max

cir-rate

The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and the value must be given as a positive integer. The actual CIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 2000000000 kb/s, **max**

Default 0

fir-rate

The **fir** parameter overrides the default administrative FIR used by the queue. When the rate command is executed, an FIR setting is optional. When the rate command has not been executed or the **fir** parameter is not explicitly specified, the default FIR (0) is assumed.

Fractional values are not allowed and the value must be given as a positive integer. The actual FIR used is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

FIR is only supported on FP4 hardware and is ignored when the related policy is applied to FP2- or FP3-based hardware.

Values 1 to 2000000000 kb/s, **max**

Default 0

police

Specifies that traffic feeding into the queue instance above the specified rate is dropped.

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>qos>queue-group-templates>egress>queue-group>queue rate)

Full Context

configure qos queue-group-templates egress queue-group queue rate

Description

This command defines the administrative PIR and the administrative CIR parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets out an egress port. Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR for all queues created through the association of the egress queue group template with the *queue-id*.

When configured on an egress HSQ queue group queue, the **cir** keyword is ignored.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the rate is performed under the **hs-wrr-group** within the SAP egress QoS policy.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default

rate max cir 0

Parameters

pir-rate

Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 200000000 kb/s, **max**

Default max

cir-rate

The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and must be given as a positive integer.

Values 0 to 200000000 kb/s, **max**

Default 0

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>qos>scheduler-policy>tier>scheduler rate)

Full Context

configure qos scheduler-policy tier scheduler rate

Description

The **rate** command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's within-CIR distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child queues and schedulers to operate at their maximum rates.

The **no** form of this command returns the scheduler's PIR and CIR parameters to the value configured in the applied scheduler policy.

Parameters

pir ***pir***

Specifies the PIR rate of the scheduler in kb/s or it can be set to the maximum using the max keyword.

Values 1 to 6400000000, **max**

Default max

cir *cir*

Specifies the CIR rate of the scheduler in kb/s or it can be set to the maximum using the max keyword. The sum keyword can also be used, which sets the CIR to the sum of child CIR values.

Values 0 to 6400000000, **max, sum**

Default sum

Platforms

7705 SAR Gen 2

rate

Syntax

rate *pir-rate* [**cir** *cir-rate*]

no rate

Context

[Tree] (config>service>cust>multi-service-site>ingress>sched-override>scheduler rate)

[Tree] (config>service>cust>multi-service-site>egress>sched-override>scheduler rate)

Full Context

configure service customer multi-service-site ingress scheduler-override scheduler rate

configure service customer multi-service-site egress scheduler-override scheduler rate

Description

This command overrides specific attributes of the specified scheduler rate.

The **rate** command defines the maximum bandwidth that the scheduler can offer its child policers, queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the scheduler's amount of bandwidth to be considered during the parent schedulers 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child policers or queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child queues and schedulers to operate at their maximum rates.

The **no** form of this command returns the scheduler's to the PIR and CIR parameters to the value configured in the applied scheduler policy.

Parameters

pir-rate

Specifies the PIR rate.

Values 1 to 6400000000, **max**

Default **max**

cir-rate

Specifies the CIR rate.

If the *cir-rate* is set to **max**, then the CIR rate is set to infinity. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers, policers or queues.

Values 0 to 6400000000, **max**, **sum**

Default **sum**

Platforms

7705 SAR Gen 2

rate

Syntax

rate kbps {*kilobits-per-second* | **max**} [**mbs** *size*] [**bytes** | **kilobytes**]

rate packets {*ppi* | **max**} **within** *seconds* [**initial-delay** *packets*]

no rate

Context

[Tree] (config>sys>security>dist-cpu-protection>policy>local-monitoring-policer rate)

[Tree] (config>sys>security>dist-cpu-protection>policy>protocol>dynamic-parameters rate)

Full Context

configure system security dist-cpu-protection policy local-monitoring-policer rate

configure system security dist-cpu-protection policy protocol dynamic-parameters rate

Description

This command configures the rate and burst tolerance for the policer in either a packet rate or a bit rate.

The actual hardware may not be able to perfectly rate limit to the exact configured parameters. In this case, the configured parameters will be adapted to the closest supported rate. The actual (operational) parameters can be seen in CLI, for example, **show service id** *33* **sap** 1/1/3:33 **dist-cpu-protection detail**.

If the *kilobits-per-second* parameter value is configured as max, then the policer is effectively disabled (always conforming).

If the *size* parameter value is configured as 0, then all packets are considered as nonconforming.

Default

rate packets max within 1 initial-delay 0

Parameters

packets | kbps

specifies that the rate is either in units of packets per interval or in units of kilobits per second. The packets option would typically be used for lower rates (for example, for per-subscriber DHCP rate limiting) while the kbps option would typically be used for higher rates (for example, per-interface BGP rate limiting).

ppi

Specifies packets per interval.

Values 0 to 255, max

max = disable the policer (always conforming)

packets 0 = all packets considered nonconforming

seconds

Specifies the length of the ppi rate measurement interval.

Values 1 to 32767

packets

Specifies the number of packets allowed (even at line rate) in an initial burst (or a burst after the policer bucket has drained to zero) in addition to the normal *ppi*. This would typically be set to a value that is equal to the number of received packets in several full handshakes/negotiations of the particular protocol.

Values 0 to 255

kilobits-per-second

Specifies the kilobits per second.

Values 1 to 20000000, max

size

Specifies the tolerance for the kbps rate.

Values 0 to 4194304

Default 10

bytes | kilobytes

Specifies that the units of the mbs size parameter are either in bytes or kilobytes.

Platforms

7705 SAR Gen 2

rate

Syntax

rate kbps {*kilobits-per-second* | **max**} [**mbs** *size*] [**bytes** | **kilobytes**]

rate packets {*ppi* | **max**} **within** *seconds* [**initial-delay** *packets*]

no rate

Context

[Tree] (config>sys>security>dist-cpu-protection>policy>static-policer rate)

Full Context

configure system security dist-cpu-protection policy static-policer rate

Description

This command configures the rate and burst tolerance for the policer in either a packet rate or a bit rate.

The actual hardware may not be able to perfectly rate limit to the exact configured parameters. In this case, the configured parameters will be adapted to the closest supported rate. The actual (operational) parameters can be seen in CLI, for example, **show service id** *33* **sap** 1/1/3:33 **dist-cpu-protection detail**.

If the *kilobits-per-second* parameter value is configured as max, then the policer is effectively disabled (always conforming).

If the *size* parameter is configured as 0, then all packets are considered as nonconforming.

Default

rate packets max within 1 initial-delay 0

Parameters

packets | kbps

specifies that the rate is either in units of packets per interval or in units of kilobits per second. The packets option would typically be used for lower rates (for example, for per-subscriber DHCP rate limiting) while the kbps option would typically be used for higher rates (for example, per-interface BGP rate limiting).

ppi

Specifies packets per interval.

Values 0 to 8000, max

max = disable the policer (always conforming)

packets 0 = all packets considered nonconforming

seconds

Specifies the length of the ppi rate measurement interval.

Values 1 to 32767

packets

Specifies the number of packets allowed (even at line rate) in an initial burst (or a burst after the policer bucket has drained to zero) in addition to the normal *ppi*. This would typically be set to a value that is equal to the number of received packets in several full handshakes/negotiations of the particular protocol.

Values 0 to 255

kilobits-per-second

Specifies the kilobits per second.

Values 1 to 20000000, max

size

Specifies the tolerance for the kbps rate.

Values 0 to 4194304

Default 10

bytes | kilobytes

Specifies that the units of the mbs size parameter are either in bytes or kilobytes.

Platforms

7705 SAR Gen 2

25.15 rate-limit

rate-limit

Syntax

rate-limit *value* [**kbps** | **pps**] [**mbs** *mbs-value*]

rate-limit *value* [**kbps** | **pps**] **extracted-traffic**

rate-limit *value* [**kbps** | **pps**] **packet-length** {**lt** | **gt** | **eq**} *packet-length-value*

rate-limit *value* [**kbps** | **pps**] **packet-length range** *packet-length-value packet-length-value*

rate-limit *value* [**kbps** | **pps**] **pattern expression** *expression* **mask** *mask* **offset-type** *offset-type* **offset-value** *offset-value*

rate-limit *value* [**kbps** | **pps**] **ttl** {**lt** | **gt** | **eq**} *ttl-value*

rate-limit *value* [**kbps** | **pps**] **ttl range** *ttl-value ttl-value*

Context

[Tree] (config>filter>ip-filter>entry>action rate-limit)

Full Context

configure filter ip-filter entry action rate-limit

Description

This command configures the rate-limit value for traffic matching this filter entry. Rate-limit policers are configured with MBS equals CBS equals 10 ms of the rate and high-prio-only equals 0.

Traffic can also be rate limited based on **extracted-traffic**, **packet-length**, **ttl**, or a pattern of conditional match criteria.

Packets that match the filter entry match criteria, but do not match the conditional match criteria value, are implicitly forwarded with no further match in the following filter entries.

For pattern match:

- the expression is left-aligned for the odd number bytes, for example, the expression 0xABC is programmed 0x0ABC in the line card
- the 'data' offset requires protocol UDP or TCP to be selected in the filter entry match criteria.

Parameters

value

Specifies the **rate-limit** *value* in kb/s (default) or packets per second (pps). A rate of 0 results in all traffic being dropped. A rate of **max** results in all traffic being forwarded.

Values 0 to 2000000000 kb/s, max

0 to 100000000 pps, max

mbs-value

Specifies the maximum burst size in bytes. This parameter can only be specified when the **rate-limit** *value* unit is **kbps**.

Values 0 to 268435456

extracted-traffic

Specifies rate-limit packets both extracted to the CPM and matching the filter entry match criteria.

packet-length

Specifies rate-limit packets matching both the filter entry match criteria and the *packet-length value* defined in the **rate-limit** action statement. Packets matching the filter entry match criteria and not matching the *packet-length* value, as defined in the **rate-limit** action statement, are implicitly forwarded with no further match in the following filter entries.

Values **lt** — Specifies "less than”. The **lt** parameter cannot be used with the lowest possible numerical value for the parameter.

gt — Specifies "greater than”. The **gt** parameter cannot be used with the highest possible numerical value for the parameter.

eq — Specifies "equal to”.

packet-length-value

Specifies the packet length value for the rate limit action.

Values 0 to 65535

range

Specifies an inclusive range. When **range** is used, the start of the range (the first value entered) must be smaller than the end of the range (the second value entered).

expression

Specifies the hexadecimal pattern to match; up to eight bytes.

Values 0x0000000000000001 to 0xffffffffffffffff

mask

Specifies the mask for the pattern expression, up to eight bytes.

Values 0x0000000000000001 to 0xffffffffffffffff

offset-type

Specifies the starting point reference for the offset-value of this pattern.

Values layer-3, layer-4, data, dns-qtype

offset-value

Specifies the offset value for the pattern expression. Dns-qtype supports offset value of 0.

Values 0 to 255

ttl-value

Specifies rate-limit packets matching both the filter entry match criteria and the TTL value defined in the *rate-limit* action statement. Packets matching the filter entry match criteria and not matching the TTL value, as defined in the *rate-limit* action statement, are implicitly forwarded with no further match in the following filter entries.

Values 0 to 255

Platforms

7705 SAR Gen 2

rate-limit

Syntax

rate-limit *value* [**kbps** | **pps**] [**mbs** *mbs-value*]

rate-limit *value* [**kbps** | **pps**] **extracted-traffic**

rate-limit *value* [**kbps** | **pps**] **hop-limit** {**lt** | **gt** | **eq**} *hop-limit-value*

rate-limit *value* [**kbps** | **pps**] **hop-limit range** *hop-limit-value hop-limit-value*

rate-limit *value* [**kbps** | **pps**] **pattern expression** *expression* **mask** *mask* **offset-type** *offset-type* **offset-value** *offset-value*

rate-limit *value* [**kbps** | **pps**] **payload-length** {**lt** | **gt** | **eq**} *payload-length-value*

rate-limit *value* [**kbps** | **pps**] **payload-length range** *payload-length-value payload-length-value*

Context

[Tree] (config>filter>ipv6-filter>entry>action rate-limit)

Full Context

configure filter ipv6-filter entry action rate-limit

Description

This command configures the rate-limit value for traffic matching this filter entry.

Traffic can also be rate-limited based on **extracted-traffic**, **payload-length**, **hop-limit**, or a pattern of conditional match criteria.

Packets that match the filter entry match criteria, but do not match the conditional match criteria value, are implicitly forwarded with no further match in the following filter entries.

For pattern match:

- the expression is left-aligned for the odd number bytes, for example, the expression 0xABC is programmed 0x0ABC in the line card.
- the 'data' offset requires protocol UDP or TCP to be selected in the filter entry match criteria.

Parameters

value

Specifies the **rate-limit** *value* in kb/s (default) or packets per second (pps). A rate of 0 results in all traffic being dropped. A rate of **max** results in all traffic being forwarded.

Values 0 to 2000000000 kb/s, max

0 to 100000000 pps, max

mbs-value

Specifies the maximum burst size in bytes. This parameter can only be specified when the **rate-limit** *value* unit is **kbps**.

Values 0 to 268435456

extracted-traffic

Specifies packets extracted to the CPM.

hop-limit

Specifies the hop limit value for the rate limit action.

Values **lt** — Specifies "less than”. The **lt** parameter cannot be used with the lowest possible numerical value for the parameter.

gt — Specifies "greater than”. The **gt** parameter cannot be used with the highest possible numerical value for the parameter.

eq — Specifies "equal to”.

hop-limit-value

Specifies the hop limit value for the rate limit action.

Values 0 to 255

range

Specifies an inclusive range. When the **range** parameter is used, the start of the range (the first value entered) must be smaller than the end of the range (the second value entered).

expression

Specifies the hexadecimal pattern to match; up to eight bytes.

Values 0x0000000000000001 to 0xffffffffffffffff

mask

Specifies the mask for the pattern expression, up to eight bytes.

Values 0x0000000000000001 to 0xffffffffffffffff

offset-type

Specifies the starting point reference for the offset-value of this pattern.

Values layer-3, layer-4, data, dns-qtype

offset-value

Specifies the offset value for the pattern expression. Dns-qtype supports offset value of 0.

Values 0 to 255

payload-length

Specifies rate-limit packets matching both the filter entry match criteria and the *payload-length-value* defined in the **rate-limit** action statement. Packets matching the filter entry match criteria and not matching the *payload-length-value,* as defined in the **rate-limit** action statement, are implicitly forwarded with no further match in the following filter entries.

Values **lt** — Specifies "less than”. The **lt** parameter cannot be used with the lowest possible numerical value for the parameter.

gt — Specifies "greater than”. The **gt** parameter cannot be used with the highest possible numerical value for the parameter.

eq — Specifies "equal to”.

payload-length-value

Specifies the payload length value for the rate limit action.

Values 0 to 65535

Platforms

7705 SAR Gen 2

25.16 rd

rd

Syntax

rd *file-url* **rf**

rd *file-url* [**force**]

Context

[Tree] (file rd)

Full Context

file rd

Description

If the directory is empty, the **rd** command is used to remove it. The **force** option executes the command without prompting the user to confirm the action.

If the directory contains files and/or subdirectories, the **rf** parameter must be used to remove the directory.

Example:

```
A:nE1>file cf1:\ # rd test
Are you sure (y/n)? y
Deleting directory cf1:\test ..MINOR: CLI Cannot delete cf1:\test.
A:nE1>file cf1:\ # rd test force
Deleting directory cf1:\test .MINOR: CLI Cannot delete cf1:\test.


A:nE1>file cf1:\ # rd testbase rf
Deleting all subdirectories and files in specified directory. y/n ?y
Deleting directory cf1:\testbase\testbase1 ..OK
Deleting directory cf1:\test .OK
```

Parameters

file-url

Specifies the directory to be removed.

Values		
	local-url	[*cflash-id/*][*file-path*] up to 200 characters, including cflash-id directory length up to 99 each
	remote-url	[{ftp:// | tftp://}*login*:*pswd*@*remote-locn*/][*file-path*] up to 247 characters directory length up to 99 characters each
	remote-locn	[hostname | ipv4-address | [ipv6-address]]

ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[*-interface*] *x:x:x:x:x:x:d.d.d.d*[*-interface*] *x* - [0 to FFFF]H *d* - [0 to 255]D interface - up to 32 characters, for link local addresses 255
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

rf

Forces a recursive delete.

force

Forces an immediate deletion of the specified directory. The **rd** *file-url* **force** command executes the command without displaying a user prompt message.

Platforms

7705 SAR Gen 2

25.17 rd-entry

rd-entry

Syntax

rd-entry *rd*

no rd-entry *rd*

Context

[Tree] (config>router>policy-options>route-distinguisher-list rd-entry)

Full Context

configure router policy-options route-distinguisher-list rd-entry

Description

This command creates a route distinguisher (RD) entry in the RD list, containing an IPv4 address or ASN and the assigned number.

The **no** form of the command deletes the RD entry from the list.

Parameters

rd

Specifies a route distinguisher matching an entry in one of the following formats:

- *a.b.c.d*/*m:** – RD in IPv4 format with a wildcard character (such as 10.0.0.0/16:*)
- *a.b.c.d*/*m:n* – RD in IPv4 format with a specific number (such as 10.0.0.2/32:535)
- *asn*:* – RD in ASN format with a wildcard character (such as 65000:*)
- *asn*:*n* – RD in ASN format with a specific number (such as 65000:535)

See the "Route distinguishers" section of the *7705 SAR Gen 2 Layer 3 Services Guide: IES and VPRN* for information about Type values.

Platforms

7705 SAR Gen 2

25.18 rdnss-lifetime

rdnss-lifetime

Syntax

rdnss-lifetime {*seconds* | **infinite**}

no rdnss-lifetime

Context

[Tree] (config>service>vprn>router-advert>if>dns-options rdnss-lifetime)

[Tree] (config>service>vprn>router-advert>dns-options rdnss-lifetime)

Full Context

configure service vprn router-advertisement interface dns-options rdnss-lifetime

configure service vprn router-advertisement dns-options rdnss-lifetime

Description

This command specifies the maximum time that the RDNSS address may be used for name resolution by the client. The RDNSS Lifetime must be no more than twice MaxRtrAdvLifetime with a maximum of 3600 seconds.

Default

rdnss-lifetime infinite

Parameters

infinite

Specifies an infinite RDNSS lifetime.

seconds

Specifies the time in seconds.

Values 4to 3600

Platforms

7705 SAR Gen 2

rdnss-lifetime

Syntax

rdnss-lifetime *seconds*

rdnss-lifetime infinite

no rdnss-lifetime

Context

[Tree] (config>router>router-advert>dns-opt rdnss-lifetime)

[Tree] (config>router>router-advert>if>dns-opt rdnss-lifetime)

Full Context

configure router router-advertisement dns-options rdnss-lifetime

configure router router-advertisement interface dns-options rdnss-lifetime

Description

This command specifies the maximum time that the RDNSS address may be used for name resolution by the client.

Default

rdnss-lifetime infinite

Parameters

seconds

Specifies the time in seconds.

Values 4 to 3600

infinite

Specifies an infinite RDNSS lifetime.

Platforms

7705 SAR Gen 2

25.19 re-auth-period

re-auth-period

Syntax

re-auth-period *seconds*

no re-auth-period

Context

[Tree] (config>port>ethernet>dot1x re-auth-period)

Full Context

configure port ethernet dot1x re-auth-period

Description

This command configures the period after which re-authentication is performed. This value is only relevant if **re-authentication** is enabled.

The **no** form of this command returns the value to the default.

Default

re-auth-period 3600

Parameters

seconds

Specifies the re-authentication delay period in seconds.

Values 1 to 9000

Platforms

7705 SAR Gen 2

25.20 re-authentication

re-authentication

Syntax

[**no**] **re-authentication**

Context

[Tree] (config>port>ethernet>dot1x re-authentication)

Full Context

configure port ethernet dot1x re-authentication

Description

This command enables/disables periodic 802.1x re-authentication.

When **re-authentication** is enabled, the router re-authenticates clients on the port every **re-auth-period**.

The **no** form of this command returns the value to the default.

Default

no re-authentication

Platforms

7705 SAR Gen 2

25.21 reachable-time

reachable-time

Syntax

reachable-time *milli-seconds*

no reachable-time

Context

[Tree] (config>router>router-advert>if reachable-time)

[Tree] (config>service>vprn>router-advert>if reachable-time)

Full Context

configure router router-advertisement interface reachable-time

configure service vprn router-advertisement interface reachable-time

Description

This command configures how long this router should be considered reachable by other nodes on the link after receiving a reachability confirmation.

The configured value is placed in the reachable time field in router advertisement messages sent from this interface.

The **no** form of this command reverts to the default.

Default

reachable-time 0

Parameters

milli-seconds

Specifies the reachable time, in seconds, for advertisements from this interface.

Values 0 to 3600000

Platforms

7705 SAR Gen 2

reachable-time

Syntax

reachable-time *seconds*

no reachable-time

Context

[Tree] (config>service>vprn>if>ipv6 reachable-time)

[Tree] (config>service>vprn>ipv6 reachable-time)

Full Context

configure service vprn interface ipv6 reachable-time

configure service vprn ipv6 reachable-time

Description

This command configures how long this router should be considered reachable by other nodes on the link after receiving a reachability confirmation.

Default

no reachable-time

Parameters

seconds

Specifies the length of time, in seconds the router should be considered reachable.

Values 30 to 3600

Platforms

7705 SAR Gen 2

reachable-time

Syntax

reachable-time *seconds*

no reachable-time

Context

[Tree] (config>router>ipv6 reachable-time)

Full Context

configure router ipv6 reachable-time

Description

This command configures the neighbor reachability detection timer.

The **no** form of this command reverts to the default value.

Default

reachable-time 30

Parameters

seconds

Specifies the length of time the router should be considered reachable.

Values 30 to 3600

Platforms

7705 SAR Gen 2

reachable-time

Syntax

reachable-time *seconds*

no reachable-time

Context

[Tree] (config>router>if>ipv6 reachable-time)

Full Context

configure router interface ipv6 reachable-time

Description

This command configures the neighbor reachability detection timer.

The **no** form of this command reverts to the default value.

Default

no reachable-time

Parameters

seconds

Specifies the length of time the router should be considered reachable.

Values 30 to 3600

Platforms

7705 SAR Gen 2

25.22 read-algorithm

read-algorithm

Syntax

read-algorithm {**hash** | **hash2** | **custom**| **all-hash**}

no read-algorithm

Context

[Tree] (config>system>security>management-interface>classic-cli read-algorithm)

Full Context

configure system security management-interface classic-cli read-algorithm

Description

This command specifies how encrypted configuration secrets are interpreted, and which encryption types are accepted, when secrets are input into the system or read from a configuration file (for example at system bootup time).

The **no** form of this command reverts to the default value.

Default

read-algorithm all-hash

Parameters

hash

Specifies hash. Use this option to transport a phrase between modules and nodes. In this case the write-algorithm should be **hash** as well.

hash2

Specifies hash2 which is module-specific.

custom

Specifies the custom encryption to management interface.

all-hash

Specifies that the system accepts hash or hash2.

Platforms

7705 SAR Gen 2

25.23 reassemble

reassemble

Syntax

reassemble

Context

[Tree] (config>filter>ip-filter>entry>action reassemble)

Full Context

configure filter ip-filter entry action reassemble

Description

This command sets the filter entry action to reassemble.

Platforms

7705 SAR Gen 2

25.24 reassembly

reassembly

Syntax

reassembly [*wait-msecs*]

no reassembly

Context

[Tree] (config>service>vprn>if>sap>ip-tunnel reassembly)

[Tree] (config>service>ies>if>sap>ip-tunnel reassembly)

Full Context

configure service vprn interface sap ip-tunnel reassembly

configure service ies interface sap ip-tunnel reassembly

Description

This command configures the maximum number of seconds to wait to receive all fragments of a particular IPsec or GRE packet for reassembly.

The **no** form of this commands removes the wait time from the configuration.

Default

no reassembly

Parameters

wait-msecs

Specifies the reassembly wait time in 100 increments.

Values 1 to 5000 ms

Platforms

7705 SAR Gen 2

reassembly

Syntax

reassembly [*wait-msecs*]

no reassembly

Context

[Tree] (config>isa>tunnel-group reassembly)

Full Context

configure isa tunnel-group reassembly

Description

This command configures IP packet reassembly for IPsec and GRE tunnels supported by an MS-ISA. The **reassembly** command at the tunnel-group level configures IP packet reassembly for all IPsec and GRE tunnels associated with the tunnel-group. The **reassembly** command at the GRE tunnel level configures IP packet reassembly for that one specific GRE tunnel, overriding the tunnel-group configuration.

The **no** form of this command disables IP packet reassembly.

Default

no reassembly (tunnel-group level)

reassembly (gre-tunnel level)

Parameters

wait

Specifies the maximum number of milliseconds that the ISA tunnel application will wait to receive all fragments of a particular IPsec or GRE packet. If one or more fragments are still missing when this limit is reached the partially reassembled datagram is discarded and an ICMP time exceeded message is sent to the source host (if allowed by the ICMP configuration of the sending interface). Internally, the configured value is rounded up to the nearest multiple of 100 ms.

Values 1 to 5000

Default 2000 (tunnel-group level)

Platforms

7705 SAR Gen 2

25.25 reassembly-group

reassembly-group

Syntax

reassembly-group *nat-group-id* [**to-base-network**]

no reassembly-group

Context

[Tree] (config>service>vprn reassembly-group)

[Tree] (config>router reassembly-group)

Full Context

configure service vprn reassembly-group

configure router reassembly-group

Description

This command associates a reassembly-group consisting of multiple ISAs with the routing context in which the application requiring reassembly service resides.

Default

no reassembly-group

Parameters

nat-group-id

Specifies the NAT group ID; the NAT group contains up to 10 active ISAs.

Values 1 to 4

to-base-network

Enables the reassembly context to use network interfaces in the base routing context.

Platforms

7705 SAR Gen 2

25.26 rebind-timer

rebind-timer

Syntax

rebind-timer [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no rebind-timer

Context

[Tree] (config>service>vprn>dhcp6>server>pool>prefix rebind-timer)

[Tree] (config>router>dhcp6>server>pool>prefix rebind-timer)

Full Context

configure service vprn dhcp6 local-dhcp-server pool prefix rebind-timer

configure router dhcp6 local-dhcp-server pool prefix rebind-timer

Description

This command configures the lease rebind timer (T2) via LUDB.

The T2 time is the time at which the client contacts any available addressing authority to extend the lifetimes of DHCPv6 leases. T2 is a time duration relative to the current time expressed in units of seconds.

The IP addressing authority controls the time at which the client contacts the addressing authority to extend the lifetimes on assigned addresses/prefixes through the T1 and T2 parameters assigned to an IA. At time T1 for an IA, the client initiates a Renew/Reply message exchange to extend the lifetimes on any addresses in the IA. The client includes an IA option with all addresses/prefixes currently assigned to the IA in its Renew message. Recommended values for T1 and T2 are .5 and .8 times the shortest preferred lifetime of the addresses/prefixes in the IA that the addressing authority is willing to extend, respectively.

The configured rebind timer should always be longer than or equal to the renew timer.

The T1 and T2 are carried in the IPv6 address option that is within the IA.

The **no** form of this command reverts to the default.

Default

rebind-timer min 48

Parameters

rebind-timer

Specifies the preferred lifetime.

Values	**days** *days*	0 to 14
	hrs *hours*	0 to 23
	min *minutes*	0 to 59
	sec *seconds*	0 to 9

Platforms

7705 SAR Gen 2

25.27 reboot

reboot

Syntax

reboot [**active** | **standby** | **upgrade**] [**now**]

Context

[Tree] (admin reboot)

Full Context

admin reboot

Description

This command reboots the router or one CPM and can also be used to force an upgrade of the system boot ROMs.

If no options are specified, the user is prompted to confirm the reboot operation. Answering yes (y) will result in both CPMs and all IOMs rebooting.

```
ALA-1>admin# reboot
Are you sure you want to reboot (y/n)?
```

Parameters

active

Reboots the active CPM.

Default active

standby

Reboots the standby CPM.

Default active

upgrade

Forces card firmware to be upgraded during chassis reboot. This option should only be used if it has been indicated as required in the Release Notes or by Nokia technical support. Normally, the SR OS automatically performs firmware upgrades on CPMs and XCM/IOM cards without the need for the **upgrade** keyword.

When the **upgrade** keyword is specified, a chassis flag is set for the BOOT Loader (boot.ldr) and on the subsequent boot of the OS on the chassis, firmware images on CPMs, XCMs, and IOMs will be upgraded automatically.

Firmware on CPMs, XCMs, or IOMs that are installed in a running chassis will be upgraded automatically. For example, if a card is inserted as the result of a hot swap, and the card has a firmware version that is no longer compatible with the SR OS image running on the chassis, then the firmware on the card will be automatically upgraded before the card is brought online.

If the card firmware is upgraded, a chassis cardUpgraded (event 2032) log event is generated. The corresponding SNMP trap for this log event is tmnxEqCardFirmwareUpgraded.

During any firmware upgrade, automatic or manual, it is imperative that during the upgrade procedure:

- Power must not be switched off or interrupted.
- The system must not be reset.

- No cards are inserted or removed.

Any of the above conditions may render cards inoperable requiring a return of the card for resolution.

The time required to upgrade the firmware on the cards in the chassis depends on the number of cards to be upgraded. The progress of a firmware upgrade can be monitored at the console.

now

Forces a reboot of the router immediately without an interactive confirmation.

Platforms

7705 SAR Gen 2

25.28 recall

recall

Syntax

[**no**] **recall**

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>history recall)

Full Context

configure system management-interface cli md-cli environment history recall

Description

This command configures command history recall and search execution. When enabled, command history recall (**!**), substitution (**!$**), display (**:p**, **Esc+.**), and backward search (**Ctrl-R**) is enabled.

The **no** form of this command disables history recall and search execution.

Default

no recall

Platforms

7705 SAR Gen 2

25.29 receive

receive

Syntax

receive {**both** | **none** | **version-1** | **version-2**}

no receive

Context

[Tree] (config>service>vprn>rip>group>neighbor receive)

[Tree] (config>service>vprn>ripng>group receive)

[Tree] (config>service>vprn>rip>group receive)

[Tree] (config>service>vprn>rip receive)

[Tree] (config>service>vprn>ripng receive)

[Tree] (config>service>vprn>ripng>group>neighbor receive)

Full Context

configure service vprn rip group neighbor receive

configure service vprn ripng group receive

configure service vprn rip group receive

configure service vprn rip receive

configure service vprn ripng receive

configure service vprn ripng group neighbor receive

Description

This command configures the type(s) of RIP updates that will be accepted and processed.

If **both** or **version-2** is specified, the RIP instance listens for and accepts packets sent to the broadcast and multicast (224.0.0.9) addresses.

If **version-1** is specified, the router only listens for and accepts packets sent to the broadcast address.

This control can be issued at the global, group or interface level. The default behavior accepts and processes both RIPv1 and RIPv2 messages.

The **no** form of this command resets the type of messages accepted to both.

Default

no receive

Parameters

both

Accept RIP updates in either Version 1 or Version 2 format.

none

Do not accept and RIP updates.

version-1

Router should only accept RIP updates in Version 1 format.

version-2

Router should only accept RIP updates in Version 2 format.

Platforms

7705 SAR Gen 2

receive

Syntax

receive

Context

[Tree] (config>system>security>keychain>direction>uni receive)

Full Context

configure system security keychain direction uni receive

Description

This command enables the receive nodal context. Entries defined under this context are used to authenticate TCP segments that are being received by the router.

Platforms

7705 SAR Gen 2

receive

Syntax

receive *option-number*

no receive

Context

[Tree] (config>system>security>keychain>tcp-option-number receive)

Full Context

configure system security keychain tcp-option-number receive

Description

This command configures the TCP option number accepted in TCP packets received.

The **no** form of this command reverts to the default value.

Default

receive 254

Parameters

option-number

Specifies an enumerated integer that indicates the TCP option number to be used in the TCP header.

Values 253, 254, 253&254, tcp-ao

Platforms

7705 SAR Gen 2

receive

Syntax

receive {**both** | **none** | **version-1** | **version-2**}

no receive

Context

[Tree] (config>router>rip receive)

[Tree] (config>router>rip>group receive)

[Tree] (config>router>ripng>group receive)

[Tree] (config>router>ripng>group>neighbor receive)

[Tree] (config>router>rip>group>neighbor receive)

[Tree] (config>router>ripng receive)

Full Context

configure router rip receive

configure router rip group receive

configure router ripng group receive

configure router ripng group neighbor receive

configure router rip group neighbor receive

configure router ripng receive

Description

This command configures the types of RIP updates that will be accepted and processed.

If **both** or **version-2** is specified, the RIP instance listens for and accepts packets sent to the broadcast and multicast (224.0.0.9) addresses.

If **version-1** is specified, the router only listens for and accept packets sent to the broadcast address.

This control can be issued at the global, group or interface level. The default behavior is to accept and process both RIPv1 and RIPv2 messages.

The **no** form of the command reverts to the default value.

Default

receive both – in the config>router>rip context

receive version-1 – in the config>router>ripng context

Parameters

both

Specifies that RIP updates in either version 1 or version 2 format will be accepted.

none

Specifies that RIP updates will not be accepted.

version-1

Specifies that RIP updates in version 1 format only will be accepted.

version-2

Specifies that RIP updates in version 2 format only will be accepted.

Platforms

7705 SAR Gen 2

25.30 receive-interval

receive-interval

Syntax

receive-interval *receive-interval*

no receive-interval

Context

[Tree] (config>router>bfd>bfd-template receive-interval)

Full Context

configure router bfd bfd-template receive-interval

Description

This command specifies the receive timer used for BFD packets. If the template is used for a BFD session on an MPLS-TP LSP, then this timer is used for CC packets.

The **no** form of this command reverts to the default value.

Default

receive-interval 100

Parameters

receive-interval

Specifies the receive interval. The minimum interval that can be configured is hardware dependent.

Values 10 ms to 100,000 ms in 1 ms intervals

Default 10 ms for CPM3 or higher; 1 second for other hardware

Platforms

7705 SAR Gen 2

25.31 received-garp-flood

received-garp-flood

Syntax

[**no**] **received-garp-flood**

Context

[Tree] (config>service>vpls>proxy-arp received-garp-flood)

Full Context

configure service vpls proxy-arp received-garp-flood

Description

This command configures flooding of GARP requests and replies received on a SAP (or SDP-bind) to the service flood list (which includes EVPN destinations and other SAPs and SDP-binds).

The **no** form of this command does not flood GARPs.

Default

received-garp-flood

Platforms

7705 SAR Gen 2

25.32 received-host-unsolicited-na-flood

received-host-unsolicited-na-flood

Syntax

[**no**] **received-host-unsolicited-na-flood**

Context

[Tree] (config>service>vpls>proxy-nd received-host-unsolicited-na-flood)

Full Context

configure service vpls proxy-nd received-host-unsolicited-na-flood

Description

This command configures the system to flood received unsolicited NAs into the VPLS service (to EVPN destinations and SAPs or SDP-binds).

The impacted NA messages contain the following flags: [S=0 and R=0].

The **no** form of this command does not flood unsolicited NAs.

Default

received-host-unsolicited-na-flood

Platforms

7705 SAR Gen 2

25.33 received-router-unsolicited-na-flood

received-router-unsolicited-na-flood

Syntax

[**no**] **received-router-unsolicited-na-flood**

Context

[Tree] (config>service>vpls>proxy-nd received-router-unsolicited-na-flood)

Full Context

configure service vpls proxy-nd received-router-unsolicited-na-flood

Description

This command configures the system to flood received unsolicited router NAs into the VPLS service (to EVPN destinations and SAPs or SDP-binds).

The impacted NA messages contain the following flags: [S=0 and R=1].

The **no** form of this command does not flood unsolicited NAs.

Default

received-router-unsolicited-na-flood

Platforms

7705 SAR Gen 2

25.34 received-unknown-arp-request-flood

received-unknown-arp-request-flood

Syntax

[**no**] **received-unknown-arp-request-flood**

Context

[Tree] (config>service>vpls>proxy-arp received-unknown-arp-request-flood)

Full Context

configure service vpls proxy-arp received-unknown-arp-request-flood

Description

This command configures flooding of unknown ARP requests received on a SAP (or SDP-bind) to the service flood list (which includes EVPN destinations and other SAPs and SDP-binds).

By default, if there is no active proxy ARP entry for the requested IP address, the system floods ARP requests, including EVPN (with source squelching).

The **no** form of this command does not flood unknown ARP requests.

Default

received-unknown-arp-request-flood

Platforms

7705 SAR Gen 2

25.35 received-unknown-ns-flood

received-unknown-ns-flood

Syntax

[**no**] **received-unknown-ns-flood**

Context

[Tree] (config>service>vpls>proxy-nd received-unknown-ns-flood)

Full Context

configure service vpls proxy-nd received-unknown-ns-flood

Description

This command configures the system to flood received unknown NS messages into the VPLS service (to EVPN destinations and SAPs or SDP-binds).

The **no** form of this command does not flood unknown NS messages.

Default

received-unknown-ns-flood

Platforms

7705 SAR Gen 2

25.36 reclassify-using-qos

reclassify-using-qos

Syntax

reclassify-using-qos *policy-id*

no reclassify-using-qos

Context

[Tree] (config>service>ies>if>vpls>egress reclassify-using-qos)

Full Context

configure service ies interface vpls egress reclassify-using-qos

Description

The reclassify-using-qos command is used to specify a sap-egress QoS policy that will be used to reclassify the forwarding class and profile of egress routed packets on the VPLS or I-VPLS service. When routed packets associated with the IP interface egress a VPLS SAP, the reclassification rules within the sap-egress QoS policy applied to the SAP are always ignored (even when reclassify-using-qos is not defined).

Any queues or policers defined within the specified QoS policy are ignored and are not created on the VPLS egress SAPs. Instead, the routed packets continue to use the forwarding class mappings, queues and policers from the sap-egress QoS policy applied to the egress VPLS SAP.

While the specified sap-egress policy ID is applied to an IP interface it cannot be deleted from the system.

The **no** form of this command removes the sap-egress QoS policy used for reclassification from the egress IP interface. When removed, IP routed packets will not be reclassified on the egress SAPs of the VPLS service attached to the IP interface.

Parameters

policy-id

Specifies the SAP egress QoS policy ID. This parameter is required when executing the reclassify-using-qos command. The specified SAP egress QoS ID must exist within the system or the command fails.

Platforms

7705 SAR Gen 2

reclassify-using-qos

Syntax

reclassify-using-qos *policy-id*

no reclassify-using-qos

Context

[Tree] (config>service>vprn>if>vpls>egress reclassify-using-qos)

Full Context

configure service vprn interface vpls egress reclassify-using-qos

Description

This command specifies a SAP egress QoS policy that is used to reclassify the forwarding class and profile of egress routed packets on the VPLS service. When routed packets associated with the IP interface egress a VPLS SAP, the reclassification rules within the sap-egress QoS policy applied to the SAP are always ignored (even when reclassify-using-qos is not defined).

Any queues or policers defined within the specified QoS policy are ignored and are not created on the VPLS egress SAPs. Instead, the routed packets continue to use the forwarding class mappings, queues and policers from the SAP egress QoS policy applied to the egress VPLS SAP.

While the specified SAP egress policy ID is applied to an IP interface it cannot be deleted from the system.

The **no** form of this command removes the SAP egress QoS policy used for reclassification from the egress IP interface. When removed, IP routed packets is not reclassified on the egress SAPs of the VPLS service attached to the IP interface.

Parameters

policy-id

Specifies the SAP egress QoS policy ID This parameter is required when executing the **reclassify-using-qos** command. The specified SAP egress QoS ID must exist within the system or the command fails.

Platforms

7705 SAR Gen 2

25.37 record

record

Syntax

[**no**] **record**

Context

[Tree] (config>router>mpls>lsp>secondary record)

[Tree] (config>router>mpls>lsp>primary record)

[Tree] (config>router>mpls>lsp-template record)

Full Context

configure router mpls lsp secondary record

configure router mpls lsp primary record

configure router mpls lsp-template record

Description

This command enables recording of all the hops that an LSP path traverses. Enabling **record** increases the size of the PATH and RESV refresh messages for the LSP since this information is carried end-to-end along the path of the LSP. The increase in control traffic per LSP may impact scalability.

The **no** form of this command disables the recording of all the hops for the given LSP. There are no restrictions as to when the **no** command can be used. The **no** form of this command also disables the **record-label** command.

Default

record

Platforms

7705 SAR Gen 2

record

Syntax

[**no**] **record** *record-name*

Context

[Tree] (config>log>accounting-policy record)

Full Context

configure log accounting-policy record

Description

This command adds the accounting record type to the accounting policy that is forwarded to the configured accounting file. A record name can only be used in one accounting policy. To obtain a list of all record types that can be configured, use the **show log accounting-records** command.

To configure an accounting policy for access ports, select a service record (for example, service-ingress-octets). To change the record name to another service record, enter the **record** command with the new record name and it replaces the old record name.

When configuring an accounting policy for network ports, select a network record. To change the record name to another network record, enter the **record** command with the new record name and it replaces the old record name.

If the change required modifies the record from network to service or from service to network, then the old record name must be removed using the **no** form of this command.

Only one record can be configured in a single accounting policy. For example, if an accounting-policy is configured with an **access-egress-octets** record, to change it to a **service-ingress-octets** record, use the **no record** command under the accounting-policy to remove the old record first, and then enter the **service-ingress-octets** record.

Note:
Collecting excessive statistics can adversely affect the CPU utilization and take up large amounts of storage space.

The **no** form of this command removes the record type from the policy.

Default

no record

Parameters

record-name

Specifies the accounting record name.

Platforms

7705 SAR Gen 2

25.38 record-label

record-label

Syntax

[**no**] **record-label**

Context

[Tree] (config>router>mpls>lsp>secondary record-label)

[Tree] (config>router>mpls>lsp-template record-label)

[Tree] (config>router>mpls>lsp>primary record-label)

Full Context

configure router mpls lsp secondary record-label

configure router mpls lsp-template record-label

configure router mpls lsp primary record-label

Description

This command enables recording of all the labels at each node that an LSP path traverses. Enabling the **record-label** command will also enable the **record** command if it is not already enabled.

The **no** form of this command disables the recording of the hops that an LSP path traverses.

Default

record-label

Platforms

7705 SAR Gen 2

25.39 record-stats

record-stats

Syntax

record-stats {**delay** | **loss** | **delay-and-loss**}

no record-stats

Context

[Tree] (config>oam-pm>session>ip>twamp-light record-stats)

Full Context

configure oam-pm session ip twamp-light record-stats

Description

This option provides the ability to determine which statistics are recorded. The TWAMP-Light PDU can report on both delay and loss using a single packet. The operator may choose which statistics they would like to report. Only delay recording is on by default. All other metrics are ignored. In order to change what is being recorded and reported, the TWAMP-Light session must be shutdown. This is required because the single packet approach means the base statistics are shared between the various datasets. Issuing a **no shutdown** command clears previous all non-volatile memory for the session and allocate new memory blocks. All the parameters under this context are mutually exclusive.

The **no** version of the command restores the default "delay" only.

Default

record-stats delay

Parameters

delay

Specifies report on delay using a single packet..

loss

Specifies to report on loss using a single packet..

delay-and-loss

Specifies to report on both delay and loss using a single packet.

Platforms

7705 SAR Gen 2

25.40 red

red

Syntax

[**no**] **red** [**detail**]

Context

[Tree] (debug>router>pim red)

Full Context

debug router pim red

Description

This command enables debugging for PIM redundancy messages to the standby CPM.

The **no** form of this command disables debugging for PIM redundancy messages to the standby CPM.

Parameters

detail

Displays detailed redundancy information.

Platforms

7705 SAR Gen 2

red

Syntax

red [**detail**]

no red

Context

[Tree] (debug>router>pcep>pcc>conn red)

[Tree] (debug>router>pcep>pcc red)

Full Context

debug router pcep pcc connection red

debug router pcep pcc red

Description

This command enables debugging for PCC or connection redundancy events.

The **no** form of this command disables debugging.

Parameters

detail

Keyword used to specify detailed information about PCC or connection redundancy events.

Platforms

7705 SAR Gen 2

25.41 redelegation-timer

redelegation-timer

Syntax

redelegation-timer *seconds*

no redelegation-timer

Context

[Tree] (config>router>pcep>pcc redelegation-timer)

Full Context

configure router pcep pcc redelegation-timer

Description

This command configures the redelegation timer for PCE-initiated LSPs.

The **no** form of the command sets this value to the default.

Default

redelegation-timer 90

Parameters

seconds

Specifies the number of seconds before the redelegation timer expires.

Values 1 to 3600

Platforms

7705 SAR Gen 2

25.42 redirect-policy

redirect-policy

Syntax

redirect-policy *redirect-policy-name* [**create**]

no redirect-policy *redirect-policy-name*

Context

[Tree] (config>filter redirect-policy)

Full Context

configure filter redirect-policy

Description

This command, creates a configuration context for the specified redirect policy.

The **no** form of the command removes the redirect policy from the filter configuration only if the policy is not referenced in a filter and the filter is not in use (applied to a service or network interface).

Parameters

redirect-policy-name

Specifies the redirect policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. There is no limit to the number of redirect policies that can be configured.

create

This keyword is required to create the configuration context. Once it is created, the context can be enabled with or without the **create** keyword.

Platforms

7705 SAR Gen 2

redirect-policy

Syntax

redirect-policy *redirect-policy-name* **destination** *ip-address*

no redirect-policy *redirect-policy-name* [**destination** *ip-address*]

Context

[Tree] (config>filter>redirect-policy-binding redirect-policy)

Full Context

configure filter redirect-policy-binding redirect-policy

Description

This command adds the destination (specified by its IP address) of a redirect-policy (specified by its name) to the binding. An error is thrown if either the destination does not exist for the specified redirect-policy or if the redirect-policy does not exist.

The **no** form of the command removes from the binding from all the destinations of the specified redirect-policy, or only the specified destination.

Parameters

redirect-policy-name

Specifies the name of the redirect-policy (up to 32 characters) as the destination that is to be added to the binding.

ip-address

The IP address of the destination. This can be an IPv4 or IPv6 address.

Values

ipv4-address:	a.b.c.d.
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

Platforms

7705 SAR Gen 2

25.43 redirect-policy-binding

redirect-policy-binding

Syntax

redirect-policy-binding *name* [**create**]

no redirect-policy-binding *name*

Context

[Tree] (config>filter redirect-policy-binding)

Full Context

configure filter redirect-policy-binding

Description

This command creates a redirect-policy binding (specified by its name) in case it does not exist and, enters the context associated with it. When a redirect-policy binding is created, no destination is associated to this binding by default and the binding operator is set to AND.

The **no** form of this command deletes the redirect-policy binding and all the associated configuration information.

Parameters

name

Specifies the name of the binding. Possible values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotations.

create

This keyword is required to create the binding if it does not exist. This has no effect when used with an existing binding.

Platforms

7705 SAR Gen 2

25.44 redirect-vprn

redirect-vprn

Syntax

redirect-vprn

Context

[Tree] (config>router>dns redirect-vprn)

Full Context

configure router dns redirect-vprn

Description

This command configures the DNS resolution to be resolved via VPRN. If configured, all packet URL resolution is done through a DNS server that is reachable in a VPRN. This includes packets in the global routing table.

Default

redirect-vprn

Platforms

7705 SAR Gen 2

25.45 redirection

redirection

Syntax

redirection *level*

no redirection

Context

[Tree] (config>system>file-trans-prof redirection)

Full Context

configure system file-transmission-profile redirection

Description

This command enables system to accept HTTP redirection response, along with the max level of redirection. The virtual router may send a new request to another server if the requested resources are not available (temporarily available to another server).

Default

no redirection

Parameters

level

Specifies the maximum level of redirection of the file transmission profile max level of HTTP redirection.

Values 1 to 8

Platforms

7705 SAR Gen 2

25.46 redirects

redirects

Syntax

redirects [*number seconds*]

no redirects

Context

[Tree] (config>service>ies>if>icmp redirects)

[Tree] (config>service>vprn>if>ipv6>icmp6 redirects)

[Tree] (config>service>vprn>nw-if>icmp redirects)

[Tree] (config>service>ies>if>ipv6>icmp6 redirects)

[Tree] (config>service>vprn>if>icmp redirects)

Full Context

configure service ies interface icmp redirects

configure service vprn interface ipv6 icmp6 redirects

configure service vprn network-interface icmp redirects

configure service ies interface ipv6 icmp6 redirects

configure service vprn interface icmp redirects

Description

This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.

When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.

The **redirects** command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.

The **no** form of this command disables the generation of ICMP redirects on the router interface.

Default

redirects 100 10

Parameters

number

Specifies the maximum number of ICMP redirect messages to send. This parameter must be specified with the *second* parameter.

Values 10 to 1000

seconds

Specifies the time frame in seconds used to limit the *number* of ICMP redirect messages that can be issued.

Values 1 to 60

Platforms

7705 SAR Gen 2

redirects

Syntax

redirects [*number seconds*]

no redirects

Context

[Tree] (config>router>if>icmp redirects)

Full Context

configure router interface icmp redirects

Description

This command enables and configures the rate for ICMP redirect messages issued on the router interface.

When routes are not optimal on this router, and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.

The **redirects** command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects are issued can be controlled with the optional *number* and *time* parameters by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of ICMP redirects on the router interface.

Default

redirects 100 10 — Maximum of 100 redirect messages in 10 seconds.

Parameters

number

The maximum number of ICMP redirect messages to send, expressed as a decimal integer. This parameter must be specified with the *time* parameter.

Values 10 to 1000

seconds

The time frame, in seconds, used to limit the *number* of ICMP redirect messages that can be issued, expressed as a decimal integer.

Values 1 to 60

Platforms

7705 SAR Gen 2

redirects

Syntax

redirects [*number seconds*]

no redirects

Context

[Tree] (config>router>if>ipv6>icmp6 redirects)

Full Context

configure router interface ipv6 icmp6 redirects

Description

This command configures the rate for ICMPv6 redirect messages. When configured, ICMPv6 redirects are generated when routes are not optimal on the router and another router on the same subnetwork has a better route to alert that node that a better route is available.

The **no** form of this command disables ICMPv6 redirects.

Default

redirects 100 10 (when IPv6 is enabled on the interface)

Parameters

number

Limits the number of redirects issued per the time frame specified in *seconds* parameter.

Values 10 to 1000

seconds

Determines the time frame, in seconds, that is used to limit the number of redirects issued per time frame.

Values 1 to 60

Platforms

7705 SAR Gen 2

25.47 redistribute-delay

redistribute-delay

Syntax

redistribute-delay *redistribute-delay*

no redistribute-delay

Context

[Tree] (config>router>ospf3>timers redistribute-delay)

[Tree] (config>router>ospf>timers redistribute-delay)

Full Context

configure router ospf3 timers redistribute-delay

configure router ospf timers redistribute-delay

Description

This command sets the internal OSPF hold down timer for external routes being redistributed into OSPF.

Shorting this delay can speed up the advertisement of external routes into OSPF but can result in additional OSPF messages if that source route is not yet stable.

The **no** form of this command resets the timer value back to the default value.

Note:
The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Default

redistribute-delay 1000

Parameters

redistribute-delay

Specifies the OSPF redistribution hold down time in milliseconds for external routes being advertised into OSPF.

Values 0 to 1000

Platforms

7705 SAR Gen 2

25.48 redistribute-external

redistribute-external

Syntax

[**no**] **redistribute-external**

Context

[Tree] (config>service>vprn>ospf>area>nssa redistribute-external)

[Tree] (config>service>vprn>ospf3>area>nssa redistribute-external)

Full Context

configure service vprn ospf area nssa redistribute-external

configure service vprn ospf3 area nssa redistribute-external

Description

This command enables the redistribution of external routes into the Not So Stubby Area (NSSA) or an NSSA area border router (ABR) that is exporting the routes into non-NSSA areas.

NSSA or Not So Stubby Areas are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an ABR to the entire OSPF domain.

The **no** form of this command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.

Default

redistribute-external — External routes are redistributed into the NSSA.

Platforms

7705 SAR Gen 2

redistribute-external

Syntax

[**no**] **redistribute-external**

Context

[Tree] (config>router>ospf3>area>nssa redistribute-external)

[Tree] (config>router>ospf>area>nssa redistribute-external)

Full Context

configure router ospf3 area nssa redistribute-external

configure router ospf area nssa redistribute-external

Description

This command enables the redistribution of external routes into the Not So Stubby Area (NSSA) or an NSSA area border router (ABR) that is exporting the routes into non-NSSA areas.

NSSA or Not So Stubby Areas are similar to stub areas in that no external routes are imported into the area from other OSPF or OSPF3 areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an Area Border Router to the entire OSPF or OSPF3 domain.

The **no** form of this command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.

Default

redistribute-external

Platforms

7705 SAR Gen 2

25.49 redo

redo

Syntax

redo [*count*]

Context

[Tree] (candidate redo)

Full Context

candidate redo

Description

This command reapplies the changes to the candidate that were removed using a previous undo. All undo or redo history is lost when the operator exits **edit-cfg** mode.

A **redo** command is blocked if another user has made changes in the same CLI branches that would be impacted during the redo.

Parameters

count

Specifies the number of previous changes to reapply.

Values 1 to 50

Default 1

Platforms

7705 SAR Gen 2

25.50 reduced-prompt

reduced-prompt

Syntax

reduced-prompt [*no-of-nodes-in-prompt*]

no reduced-prompt

Context

[Tree] (environment reduced-prompt)

Full Context

environment reduced-prompt

Description

This command configures the maximum number of higher CLI context levels to display in the CLI prompt for the current CLI session. This command is useful when configuring features that are several node levels deep, causing the CLI prompt to become too long. By default, the CLI prompt displays the system name and the complete context in the CLI.

The number of *nodes* specified indicates the number of higher-level contexts that can be displayed in the prompt. For example, if reduced prompt is set to 2, the two highest contexts from the present working context are displayed by name with the hidden (reduced) contexts compressed into a ellipsis ("…").

```
A:ALA-1>environment# reduced-prompt 2
A:ALA-1>config>router# interface to-103
A:ALA-1>...router>if#
```

The setting is not saved in the configuration. It must be reset for each CLI session or stored in an **exec** script file.

The **no** form of the command reverts to the default.

Default

no reduced-prompt

Parameters

no-of-nodes-in-prompt

Specifies the maximum number of higher-level nodes displayed by name in the prompt, expressed as a decimal integer.

Values 0 to 15

Default 2

Platforms

7705 SAR Gen 2

25.51 redundancy

redundancy

Syntax

redundancy

Context

[Tree] (config redundancy)

Full Context

configure redundancy

Description

This command allows the user to perform redundancy operations.

Associated commands include the following in the **admin>redundancy** context:

- **force-switchover** - Forces a switchover to the standby CPM card.
- **now** - Switch to standby CPM.

  Switching to the standby displays the following message.

  ```
  WARNING: Configuration and/or Boot options may have changed since the last save.

  Are you sure you want to switchover (y/n)?
  ```
- **synchronize** - Synchronizes the secondary CPM.

Platforms
7705 SAR Gen 2

redundancy

Syntax
redundancy

Context
[Tree] (admin redundancy)

Full Context
admin redundancy

Description
Commands in this context allow the user to perform redundancy operations.

Platforms
7705 SAR Gen 2

25.52 redundant-multicast

redundant-multicast

Syntax
[**no**] **redundant-multicast**

Context
[Tree] (config>router>igmp>if redundant-multicast)

Full Context
configure router igmp interface redundant-multicast

Description
This command configures the interface as a member of a redundant pair for multicast traffic.

The **no** form of the command removes the configuration.

Platforms
7705 SAR Gen 2

25.53 ref-policer

ref-policer

Syntax

ref-policer *policer-id*

ref-policer all

no ref-policer

Context

[Tree] (config>log>acct-policy>cr ref-policer)

Full Context

configure log accounting-policy custom-record ref-policer

Description

This command creates a policer context to configure reference policer counters for significant change only reporting. The custom record is only generated when the change in the sum of all queue and policer reference counters equals or exceeds the configured (non-zero) significant change value.

The **no** form of this command deletes all policer reference counters.

Default

no ref-policer

Parameters

policer-id

Specifies the policer for which reference counters are configured and to which **significant-change** is applied.

Values 1 to 63

all

Applies the **significant-change** to the specified counters for all policers.

Platforms

7705 SAR Gen 2

25.54 ref-queue

ref-queue

Syntax

ref-queue *queue-id*

ref-queue all

no ref-queue

Context

[Tree] (config>log>acct-policy>cr ref-queue)

Full Context

configure log accounting-policy custom-record ref-queue

Description

This command creates a queue context to configure reference queue counters for significant change only reporting. The custom record is only generated when the change in the sum of all queue and policer reference counters equals or exceeds the configured (non-zero) significant change value.

The **no** form of this command deletes all queue reference counters.

Default

no ref-queue

Parameters

queue-id

Specifies the queue for which reference counters are configured and to which the **significant-change** is applied.

Values 1 to 32

all

Applies the **significant-change** to the specified counters for all queues.

Platforms

7705 SAR Gen 2

25.55 reference-bandwidth

reference-bandwidth

Syntax

reference-bandwidth *bandwidth-in-kbps*

reference-bandwidth [**zbps** *Zetta-bps*] [**ebps** *Exa-bps*] [**pbps** *Peta-bps*] [**tbps** *Tera-bps*] [**gbps** *Giga-bps*] [**mbps** *Mega-bps*] [**kbps** *Kilo-bps*]

no reference-bandwidth

Context

[Tree] (config>service>vprn>isis reference-bandwidth)

Full Context

configure service vprn isis reference-bandwidth

Description

This command configures the reference bandwidth that provides the basis of bandwidth relative costing.

In order to calculate the lowest cost to reach a specific destination, each configured level on each interface must have a cost. If the reference bandwidth is defined, then the cost is calculated using the following formula:

cost = reference – bandwidth # bandwidth

If the reference bandwidth is configured as 10 Gigabits (10,000,000,000), a 100 M/bps interface has a default metric of 100. In order for metrics in excess of 63 to be configured, wide metrics must be deployed. (See **wide-metrics-only** in the **config>router>isis** context.)

If the reference bandwidth is not configured, all interfaces have a default metric of 10.

The **no** form of this command reverts to the default value.

Default

no reference-bandwidth — No reference bandwidth is defined. All interfaces have a metric of 10.

Parameters

Zetta-bps

Specifies the reference bandwidth in zettabits per second, expressed as a decimal integer.

Values 1 to 18

Exa-bps

Specifies the reference bandwidth in exabits per second, expressed as a decimal integer.

Values 1 to 999

Peta-bps

Specifies the reference bandwidth in petabits per second, expressed as a decimal integer.

Values 1 to 999

bandwidth-in-kbps

Specifies the reference bandwidth in kilobits per second, expressed as a decimal integer.

Values 1 to 18446744073709551615

Tera-bps

Specifies the reference bandwidth in terabits per second, expressed as a decimal integer.

Values 1 to 999

Giga-bps

Specifies the reference bandwidth in gigabits per second, expressed as a decimal integer.

Values 1 to 999

Mega-bps

Specifies the reference bandwidth in megabits per second, expressed as a decimal integer.

Values 1 to 999

Kilo-bps

Specifies the reference bandwidth in kilobits per second, expressed as a decimal integer.

Values 1 to 999

Platforms

7705 SAR Gen 2

reference-bandwidth

Syntax

reference-bandwidth *bandwidth-inkbps*

reference-bandwidth [**zbps** *Zetta-bps*] [**ebps** *Exa-bps*] [**pbps** *Peta-bps*] [**tbps** *Tera-bps*] [**gbps** *Giga-bps*] [**mbps** *Mega-bps*] [**kbps** *Kilo-bps*]

no reference-bandwidth

Context

[Tree] (config>service>vprn>ospf reference-bandwidth)

[Tree] (config>service>vprn>ospf3 reference-bandwidth)

Full Context

configure service vprn ospf reference-bandwidth

configure service vprn ospf3 reference-bandwidth

Description

This command configures the reference bandwidth in kilobits per second (kb/s) that provides the reference for the default costing of interfaces based on their underlying link speed.

The default interface cost is calculated as follows:

cost = reference–bandwidth # bandwidth

The default *reference-bandwidth* is 100,000,000 kb/s or 100 Gb/s, so the default auto-cost metrics for various link speeds are as follows:

- 10 Mb/s link default cost of 10000
- 100 Mb/s link default cost of 1000
- 1 Gb/s link default cost of 100
- 10 Gb/s link default cost of 10
- 40 Gb/s link default cost of 2
- 100 Gb/s link default cost of 1
- 400 Gb/s link default cost of 1

Note:
The default **reference-bandwidth** value must be manually configured to a higher value if interface speeds are greater than 100 Gb/s, and metrics based on link speed are used. When the default **reference-bandwidth** value is used, a metric of 1 is set on all interface speeds ≥ 100 Gb/s. For example, 100 GE, 100 GE LAG, 400 GE, and 400 GE LAG interfaces will all have a metric of 1.

If the reference bandwidth is configured as 10 Gb (reference-bandwidth 10000000000), a 100 Mb/s interface has a default metric of 100.

When a very large reference bandwidth value is configured, a metric calculation may result in a value higher than the supported protocol cost value. If this occurs, OSPF automatically reverts to the maximum configurable cost metric.

The reference-bandwidth command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the **metric** *metric* command configured in the **config>router>ospf>area>if** *ip-int-name* context.

The **no** form of this command reverts the reference bandwidth to the default value.

Default

reference-bandwidth 100000000

Parameters

bandwidth-in-kbps

Specifies the reference bandwidth in kilobits per second expressed as a decimal integer.

Values 1 to 4000000000

tbps Tera-bps

Specifies the reference bandwidth in terabits per second expressed as a decimal integer.

Values 1 to 4

gbps *Giga-bps*

Specifies the reference bandwidth in gigabits per second expressed as a decimal integer.

Values 1 to 999

mbps *Mega-bps*

Specifies the reference bandwidth in megabits per second expressed as a decimal integer.

Values 1 to 999

kbps *Kilo-bps*

Specifies the reference bandwidth in kilobits per second expressed as a decimal integer.

Values 1 to 999

Platforms

7705 SAR Gen 2

reference-bandwidth

Syntax

reference-bandwidth *bandwidth-in-kbps*

reference-bandwidth [**zbps** *Zetta-bps*] [**ebps** *Exa-bps*] [**pbps** *Peta-bps*] [**tbps** *Tera-bps*] [**gbps** *Giga-bps*] [**mbps** *Mega-bps*] [**kbps** *Kilo-bps*]

no reference-bandwidth

Context

[Tree] (config>router>isis reference-bandwidth)

Full Context

configure router isis reference-bandwidth

Description

This command configures the reference bandwidth that provides the basis of bandwidth relative costing.

To calculate the lowest cost to reach a specific destination, each configured level on each interface must have a cost. If the reference bandwidth is defined, then the cost is calculated using the following formula:

cost = reference-bandwidth # bandwidth

If the reference bandwidth is configured as 10 Gb (**reference-bandwidth** 10000000000), a 100 Mb/s interface has a default metric of 100. To configure metrics in excess of 63, wide metrics must be deployed (see **wide-metrics-only** in the **config>router>isis** context).

When a large **reference-bandwidth** value is configured, a metric calculation may result in a value higher than the supported protocol cost value. If this occurs, IS-IS automatically reverts to the maximum configurable cost metric.

If the reference bandwidth is not configured, then all interfaces have a default metric of 10.

The **no** form of this command reverts to the default value.

Default

no reference-bandwidth

Parameters

bandwidth-in-kbps

Specifies the reference bandwidth in kilobits per second, expressed as a decimal integer.

Values 1 to 18446744073709551615

Zetta-bps

Specifies the reference bandwidth in zettabits per second, expressed as a decimal integer.

Values 1 to 18

Exa-bps

Specifies the reference bandwidth in exabits per second, expressed as a decimal integer.

Values 1 to 999

Peta-bps

Specifies the reference bandwidth in petabits per second, expressed as a decimal integer.

Values 1 to 999

Tera-bps

Specifies the reference bandwidth in terabits per second, expressed as a decimal integer.

Values 1 to 999

Giga-bps

Specifies the reference bandwidth in gigabits per second, expressed as a decimal integer.

Values 1 to 999

Mega-bps

Specifies the reference bandwidth in megabits per second, expressed as a decimal integer.

Values 1 to 999

Kilo-bps

Specifies the reference bandwidth in kilobits per second, expressed as a decimal integer.

Values 1 to 999

Platforms

7705 SAR Gen 2

reference-bandwidth

Syntax

reference-bandwidth *bandwidth-in-kbps*

reference-bandwidth [**zbps** *Zetta-bps*] [**ebps** *Exa-bps*] [**pbps** *Peta-bps*] [**tbps** *Tera-bps*] [**gbps** *Giga-bps*] [**mbps** *Mega-bps*] [**kbps** *Kilo-bps*]

no reference-bandwidth

Context

[Tree] (config>router>ospf3 reference-bandwidth)

[Tree] (config>router>ospf reference-bandwidth)

Full Context

configure router ospf3 reference-bandwidth

configure router ospf reference-bandwidth

Description

This command configures the reference bandwidth in kilobits per second (kb/s) that provides the reference for the default costing of interfaces based on their underlying link speed.

The default interface cost is calculated as follows:

cost = reference-bandwidth # bandwidth

The default *reference-bandwidth* is 100,000,000 kb/s or 100 Gb/s, the default auto-cost metrics for various link speeds are as follows:

- 10 Mb/s link default cost of 10000
- 100 Mb/s link default cost of 1000
- 1 Gb/s link default cost of 100
- 10 Gb/s link default cost of 10
- 100 Gb/s link default cost of 1
- 400 Gb/s link default cost of 1

Note:
The default reference-bandwidth must be manually configured to a higher value if interface speeds are greater than 100 Gb/s, and metrics based on link speed are used. When the default reference-bandwidth is used, a metric of 1 is set on all interface speeds ≥ 100 Gb/s. For example, 100 GE, 100 GE LAG, 400 GE, and 400 GE LAG interfaces will all have a metric of 1.

If the reference bandwidth is configured as 10 Gb (reference-bandwidth 10000000000), a 100 Mb/s interface has a default metric of 100.

When a very large reference bandwidth value is configured, a metric calculation may result in a value higher than the supported protocol cost value. If this occurs, OSPF automatically reverts to the maximum configurable cost metric.

The **reference-bandwidth** command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the **metric** *metric* command configured in the **config>router>ospf>area>interface** *ip-int-name* context.

The **no** form of this command reverts to the default value.

Default

reference-bandwidth 100000000

Parameters

bandwidth-in-kbps

Specifies the reference bandwidth in kilobits per second, expressed as a decimal integer.

Values 1 to 18446744073709551615

Zetta-bps

Specifies the reference bandwidth in zettabits per second, expressed as a decimal integer.

Values 1 to 18

Exa-bps

Specifies the reference bandwidth in exabits per second, expressed as a decimal integer.

Values 1 to 999

Peta-bps

Specifies the reference bandwidth in petabits per second, expressed as a decimal integer.

Values 1 to 999

Tera-bps

Specifies the reference bandwidth in terabits per second, expressed as a decimal integer.

Values 1 to 999

Giga-bps

Specifies the reference bandwidth in gigabits per second, expressed as a decimal integer.

Values 1 to 999

Mega-bps

Specifies the reference bandwidth in megabits per second, expressed as a decimal integer.

Values 1 to 999

Kilo-bps

Specifies the reference bandwidth in kilobits per second, expressed as a decimal integer.

Values 1 to 999

Platforms

7705 SAR Gen 2

25.56 reflector

reflector

Syntax

reflector [**udp-port** *udp-port-number*] [**create**]

no reflector

Context

[Tree] (config>service>vprn>twamp-light reflector)

[Tree] (config>router>twamp-light reflector)

Full Context

configure service vprn twamp-light reflector

configure router twamp-light reflector

Description

This command configures a TWAMP Light session reflector parameters and to enable TWAMP Light functionality with the **no shutdow**n command. The **udp-port** keyword and value must be specified with the **create** keyword. An error message is generated if the specific UDP port is unavailable.

Parameters

udp-port-number

Specifies the UDP port number. A strictly enforced restricted range has been introduced. The TWAMP Light session reflector must be brought in line with this new restriction prior upgrading or rebooting from any previous release if there is an active TWAMP Light session reflector configured. Failure to do so prevents an ISSU operation from proceeding and fails to activate any reflector outside of the enforced range.

Note that in the Two-Way Active Measurement Protocol Light (TWAMP Light) section for a complete description. This parameter is required and specifies the destination udp-port that the session reflector uses to listen for TWAMP Light packets. The session controller launching the TWAMP Light packets must be configured with the same destination UDP port as part of the TWAMP Light test. The IES service uses the destination UDP port that is configured under the **router** context. Only one UDP port can be configured per unique context.

Values 862, 64364 to 64373

Platforms

7705 SAR Gen 2

25.57 refresh-reduction

refresh-reduction

Syntax

[**no**] **refresh-reduction**

Context

[Tree] (config>router>rsvp>interface refresh-reduction)

Full Context

configure router rsvp interface refresh-reduction

Description

This command enables the use of the RSVP overhead refresh reduction capabilities on this RSVP interface.

When this option is enabled, a node will enable support for three capabilities. It will accept bundles RSVP messages from its peer over this interface, it will attempt to perform reliable RSVP message delivery to its peer, and will use summary refresh messages to refresh path and resv states. The reliable message delivery must be explicitly enabled by the user after refresh reduction is enabled. The other two capabilities are enabled immediately.

A bundle message is intended to reduce overall message handling load. A bundle message consists of a bundle header followed by one or more bundle sub-messages. A sub-message can be any regular RSVP message except another bundle message. A node will only process received bundled RSVP messages but will not generate them.

When reliable message delivery is supported by both the node and its peer over the RSVP interface, an RSVP message is sent with a message_id object. A message_id object can be added to any RSVP message when sent individually or as a sub-message of a bundled message.

if the sender sets the ack_desired flag in the message_id object, the receiver acknowledges the receipt of the RSVP message by piggy-backing a message_ack object to the next RSVP message it sends to its peer. Alternatively, an ACK message can also be used to send the message_ack object. In both cases, one or many message_ack objects could be included in the same message.

The router supports the sending of separate ACK messages only but is capable of processing received message_ack objects piggy-backed to hop-by-hop RSVP messages, such as path and resv.

The router sets the ack_desired flag only in non-refresh RSVP messages and in refresh messages which contain new state information.

A retransmission mechanism based on an exponential backoff timer is supported in order to handle unacknowledged message_id objects. The RSVP message with the same message_id is retransmitted every 2 * rapid-retransmit-time interval of time. The rapid-retransmit-time is referred to as the rapid

retransmission interval as it must be smaller than the regular refresh interval configured in the **config>router>rsvp>refresh-time** context. There is also a maximum number of retransmissions of an unacknowledged RSVP message rapid-retry-limit. The node will stop retransmission of unacknowledged RSVP messages whenever the updated backoff interval exceeds the value of the regular refresh interval or the number of retransmissions reaches the value of the rapid-retry-limit parameter, whichever comes first. These two parameters are configurable globally on a system in the **config>router>rsvp** context.

Refresh summary consists of sending a summary refresh message containing a message_id list object. The fields of this object are populated each with the value of the message_identifier field in the message_id object of a previously sent individual path or resv message. The summary refresh message is sent every refresh regular interval as configured by the user using the refresh-time command in the **config>router>rsvp** context. The receiver checks each message_id object against the saved path and resv states. If a match is found, the state is updated as if a regular path or resv refresh message was received from the peer. If a specific message_identifier field does not match, then the node sends a message_id_nack object to the originator of the message.

The above capabilities are referred to collectively as "refresh overhead reduction extensions”. When the refresh-reduction is enabled on an RSVP interface, the node indicates this to its peer by setting a "refresh-reduction-capable” bit in the flags field of the common RSVP header. If both peers of an RSVP interface set this bit, all the above three capabilities can be used. Furthermore, the node monitors the settings of this bit in received RSVP messages from the peer on the interface. As soon as this bit is cleared, the router stops sending summary refresh messages. If a peer did not set the "refresh-reduction-capable” bit, a node does not attempt to send summary refresh messages.

However, if the peer did not set the "refresh-reduction-capable” bit, a node, with refresh reduction enabled and reliable message delivery enabled, will still attempt to perform reliable message delivery with this peer. If the peer does not support the message_id object, it returns an error message "unknown object class”. In this case, the node retransmits the RSVP message without the message_id object and reverts to using this method for future messages destined to this peer. The RSVP Overhead Refresh Reduction is supported with both RSVP P2P LSP path and the S2L path of an RSVP P2MP LSP instance over the same RSVP instance.

The **no** form of this command reverts to the default value.

Default

no refresh-reduction

Platforms

7705 SAR Gen 2

25.58 refresh-reduction-over-bypass

refresh-reduction-over-bypass

Syntax

refresh-reduction-over-bypass [**enable** | **disable**]

Context

[Tree] (config>router>rsvp refresh-reduction-over-bypass)

Full Context

configure router rsvp refresh-reduction-over-bypass

Description

This command enables the refresh reduction capabilities over all bypass tunnels originating on this PLR node or terminating on this Merge Point (MP) node.

By default, this is disabled. Since a bypass tunnel may merge with the primary LSP path in a node downstream of the next-hop, there is no direct interface between the PLR and the MP node and it is possible the latter will not accept summary refresh messages received over the bypass.

When disabled, the node as a PLR or MP will not set the "Refresh-Reduction-Capable" bit on RSVP messages pertaining to LSP paths tunneled over the bypass. It will also not send Message-ID in RSVP messages. This effectively disables summary refresh.

Default

refresh-reduction-over-bypass disable

Platforms

7705 SAR Gen 2

25.59 refresh-time

refresh-time

Syntax

refresh-time *seconds*

no refresh-time

Context

[Tree] (config>router>rsvp refresh-time)

Full Context

configure router rsvp refresh-time

Description

The **refresh-time** controls the interval (in s), between the successive Path and Resv refresh messages. RSVP declares the session down after it misses **keep-multiplier** *number* consecutive refresh messages.

The **no** form of this command reverts to the default value.

Default

refresh-time 30

Parameters

seconds

The refresh time in s.

Values 1 to 65535

Platforms

7705 SAR Gen 2

refresh-time

Syntax

refresh-time *seconds* **hold-time** *seconds*

no refresh-time

Context

[Tree] (config>router>origin-validation>rpki-session refresh-time)

Full Context

configure router origin-validation rpki-session refresh-time

Description

This command is used to configure the **refresh-time** and **hold-time** intervals that are used for liveness detection of the RPKI-Router session. The **refresh-time** defaults to 300 seconds and is reset whenever a Reset Query PDU or Serial Query PDU is sent to the cache server. When the timer expires, a new Serial Query PDU is sent with the last known serial number.

The **hold-time** specifies the length of time in seconds that the session is to be considered UP without any indication that the cache server is alive and reachable. The timer defaults to 600 seconds and must be at least 2x the refresh-time (otherwise the CLI command is not accepted). Reception of any PDU from the cache server resets the hold timer. When the **hold-time** expires, the session is considered to be DOWN and the stale timer is started.

Default

no refresh-time

Parameters

seconds

Specifies a time in seconds.

Values 30 to 32767

seconds

Specifies a time in seconds.

Values 60 to 65535

Platforms

7705 SAR Gen 2

25.60 refresh-timer

refresh-timer

Syntax

refresh-timer *value*

no refresh-timer

Context

[Tree] (config>service>vpls>spoke-sdp>control-channel-status refresh-timer)

[Tree] (config>service>epipe>spoke-sdp>control-channel-status refresh-timer)

Full Context

configure service vpls spoke-sdp control-channel-status refresh-timer

configure service epipe spoke-sdp control-channel-status refresh-timer

Description

This command configures the refresh timer for control channel status signaling packets. By default, no refresh packets are sent.

Default

no refresh-timer

Parameters

value

Specifies the refresh timer value, in seconds.

Values 10 to 65535

Default 0 (off)

Platforms

7705 SAR Gen 2

refresh-timer

Syntax

refresh-timer *value*

no refresh-timer

Context

[Tree] (config>service>ies>if>spoke-sdp>control-channel-status refresh-timer)

Full Context

configure service ies interface spoke-sdp control-channel-status refresh-timer

Description

This command configures the refresh timer for control channel status signaling packets. By default, no refresh packets are sent.

Default

no refresh-timer

Parameters

value

Specifies the refresh timer value.

Values 10 to 65535 seconds

Default 0 (off)

Platforms

7705 SAR Gen 2

refresh-timer

Syntax

refresh-timer *value*

no refresh-timer

Context

[Tree] (config>service>vprn>if>spoke-sdp>control-channel-status refresh-timer)

Full Context

configure service vprn interface spoke-sdp control-channel-status refresh-timer

Description

This command configures the refresh timer for control channel status signaling packets. By default, no refresh packets are sent.

Default

no refresh-timer

Parameters

value

Specifies the refresh timer value.

Values 10 to 65535 seconds

Default 0 (off)

Platforms

7705 SAR Gen 2

25.61 register

register

Syntax

register [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]

no register

Context

[Tree] (debug>router>pim register)

Full Context

debug router pim register

Description

This command enables debugging for PIM register mechanism.

The **no** form of this command disables debugging for PIM register mechanism.

Parameters

grp-ip-address

Debugs information associated with the specified PIM register.

Values multicast group address (ipv4, ipv6)

ip-address

Debugs information associated with the specified PIM register.

Values source address (ipv4, ipv6)

detail

Debugs detailed register information.

Platforms

7705 SAR Gen 2

25.62 register-message

register-message

Syntax

[**no**] **register-message** {*ip-address* | *ipv6-address*}

Context

[Tree] (config>router>pim>src-address register-message)

[Tree] (config>service>vprn>pim>src-address register-message)

Full Context

configure router pim source-address register-message

configure service vprn pim source-address register-message

Description

This command configures the source IP address for PIM register messages. The IP address can be set to any unicast address, regardless of whether it resides on the node. Ensure that the specified IP address is configured on the router as a loopback or interface IP address.

The **no** form of this command removes the IP address. By default, when no IP address is specified for the PIM instance, the source IP address for register messages is selected by choosing the smallest IP address from available interfaces on the node.

Parameters

ip-address | *ipv6-address*

Specifies the source IPv4 or IPv6 address, up to 64 characters.

Platforms

7705 SAR Gen 2

25.63 reinit-delay

reinit-delay

Syntax

reinit-delay *time*

no reinit-delay

Context

[Tree] (config>system>lldp reinit-delay)

Full Context

configure system lldp reinit-delay

Description

This command configures the time before re-initializing LLDP on a port.

The **no** form of this command reverts to the default value.

Default

no reinit-delay

Parameters

time

Specifies the time, in seconds, before re-initializing LLDP on a port.

Values 1 to 10

Default 2

Platforms

7705 SAR Gen 2

25.64 relay-plain-bootp

relay-plain-bootp

Syntax

[**no**] **relay-plain-bootp**

Context

[Tree] (config>service>vprn>if>dhcp relay-plain-bootp)

[Tree] (config>service>ies>if>dhcp relay-plain-bootp)

Full Context

configure service vprn interface dhcp relay-plain-bootp

configure service ies interface dhcp relay-plain-bootp

Description

This command enables the relaying of plain BOOTP packets.

The **no** form of this command disables the relaying of plain BOOTP packets.

Platforms

7705 SAR Gen 2

relay-plain-bootp

Syntax

[**no**] **relay-plain-bootp**

Context

[Tree] (config>router>if>dhcp relay-plain-bootp)

Full Context

configure router interface dhcp relay-plain-bootp

Description

This command enables the relaying of plain BOOTP packets.

The **no** form of this command disables the relaying of plain BOOTP packets.

Default

no relay-plain-bootp

Platforms

7705 SAR Gen 2

25.65 relay-proxy

relay-proxy

Syntax

relay-proxy [**release-update-src-ip**] [**siaddr-override** *ip-address*]

no relay-proxy

Context

[Tree] (config>service>vprn>if>dhcp relay-proxy)

[Tree] (config>service>ies>if>dhcp relay-proxy)

Full Context

configure service vprn interface dhcp relay-proxy

configure service ies interface dhcp relay-proxy

Description

This command enables the DHCPv4 relay proxy function on the interface. The command has no effect when no dhcp servers are configured (DHCPv4 relay not configured). By default, unicast DHCPv4 release messages are forwarded transparently.

A relay proxy enhances the relay such that it also relays unicast client DHCPv4 REQUEST messages (lease renewals).

- In the upstream direction, update the source IP address and add the gateway IP address (gi-address) field before sending the message to the intended DHCP server (the message is not broadcasted to all configured DHCP servers.
- In the downstream direction, remove the gi-address and update the destination IP address to the address of the `yiaddr` (your IP address) field.

The optional **release-update-src-ip** parameter updates the source IP address of a DHCP RELEASE message with the address used for relayed DHCPv4 messages.

The optional **siaddr-override** *ip-address* parameter enables DHCP server IP address hiding towards the client. This parameter requires that **lease-populate** is enabled on the interface. The DHCP server ip address is required for the address hiding function and is stored in the lease state record. The client interacts with the relay proxy as if it is the DHCP server. In all DHCP messages to the client, the value of following header fields and DHCP options containing the DHCP server IP address is replaced with the configured <*ip-address*>:

- the "source IP address” field in the IP DHCPv4 packet header
- the "siaddr” field in the DHCPv4 header if not equal to zero in the message received from the server
- the Server Identification option (DHCPv4 option 54) if present in the original server message
- the source IP address field in the IP packet header

DHCP OFFER selection during initial binding is done in the relay-proxy. Only the first DHCP OFFER message is forwarded to the client. Subsequent DHCP OFFER messages from different servers are silently dropped.

Parameters

release-update-src-ip

Updates the source IP address of a DHCP RELEASE message with the address used for relayed DHCPv4 messages.

ip-address

Enables DHCPv4 server address hiding towards the DHCPv4 client and activates DHCPv4 OFFER selection in case multiple DHCP servers are configured. The *ip-address* can be any local address in the same routing instance. If DHCP relay lease-split is enabled, **siaddr-override** *ip-address* has priority over the **emulated-server** *ip-address* configured in the proxy-server and is used as the source IP address.

Platforms

7705 SAR Gen 2

25.66 relay-unsolicited-cfg-attribute

relay-unsolicited-cfg-attribute

Syntax

relay-unsolicited-cfg-attribute

Context

[Tree] (config>ipsec>ike-policy relay-unsolicited-cfg-attribute)

Full Context

configure ipsec ike-policy relay-unsolicited-cfg-attribute

Description

This command enters relay unsolicited configuration attributes context. With this configuration, the configured attributes returned from source (such as a RADIUS server) will be returned to IKEv2 remote-access tunnel client regardless if the client has requested it in the CFG_REQUEST payload.

Platforms

7705 SAR Gen 2

25.67 reliable-delivery

reliable-delivery

Syntax

[**no**] **reliable-delivery**

Context

[Tree] (config>router>rsvp>if>refresh-reduction reliable-delivery)

Full Context

configure router rsvp interface refresh-reduction reliable-delivery

Description

This command enables reliable delivery of RSVP messages over the RSVP interface. When refresh-reduction is enabled on an interface and reliable-delivery is disabled, the router will send a message_id and not set ACK desired in the RSVP messages over the interface. The router does not expect an ACK and but will accept it if received. The node will also accept message ID and reply with an ACK when requested. In this case, if the neighbor set the "refresh-reduction-capable” bit in the flags field of the common RSVP header, the node will enter summary refresh for a specific message_id it sent regardless if it received an ACK or not to this message from the neighbor.

Finally, when 'reliable-delivery’ option is enabled on any interface, RSVP message pacing is disabled on all RSVP interfaces of the system, for example, the user cannot enable the **msg-pacing** option in the **config>router>rsvp** context, and error message is returned in CLI. Conversely, when the **msg-pacing** option is enabled, the user cannot enable the reliable delivery option on any interface on this system. An error message is also generated in CLI after such an attempt.

The **no** form of this command reverts to the default value.

Default

no reliable-delivery

Platforms

7705 SAR Gen 2

25.68 reload

reload

Syntax

reload type {**cert** | **key** | **cert-key-pair**} *filename* **protocol** *protocol* [**key-file** *filename*]

Context

[Tree] (admin>certificate reload)

Full Context

admin certificate reload

Description

This command reloads imported certificate or key file or both at the same time. This command is typically used to update certificate or key file without shutting down **ipsec-tunnel/ipsec-gw/cert-profile/ca-profile**. Note that **type cert** and **type key** is deprecated in a future release. Use **type cert-key-pair** instead. Instead of **type cert** use **type key** instead.

- If the new file exists and valid, then for each tunnel using it:
  - If the key matches the certificate, then the new file is downloaded to the MS-ISA to be used the next time. Tunnels currently up are not affected.
  - If the key does not match the certificate:
    - If **cert** and **key** configuration is used instead of **cert-profile** then the tunnel is brought down.
    - If **cert-profile** is used, then **cert-profile** is brought down. The next authentication fails while the established tunnels are not affected.

If the new file does not exists or somehow invalid (bad format, does not contain right extension, and so on), then this command will abort.

In the case of **type cert-key-pair**, if the new file does not exist or is invalid or **cert** and **key** do not match, then this command aborts with an error message.

Parameters

type

Specifies what item will be reloaded.

cert

Specifies that a certificate cache will be reloaded.

key

Specifies that a key cache will be reloaded.

cert-key-pair

Specifies that a paired certificate and key cache will be reloaded.

filename

Up to 95 characters.

protocol

Specifies which protocol the certificate will be reloaded for.

Values ipsec, tls

Platforms

7705 SAR Gen 2

25.69 remarking

remarking

Syntax

remarking [**force**]

no remarking

Context

[Tree] (config>qos>network>egress remarking)

Full Context

configure qos network egress remarking

Description

This command remarks both customer traffic and egress network IP interface traffic; VPRN customer traffic is not remarked. The remarking is based on the forwarding class to DSCP and LSP EXP bit mapping defined under the egress node of the network QoS policy.

Normally, packets that ingress on network ports have either the DSCP or, for MPLS packets, LSP EXP bit set by an upstream router. The packets are placed in the appropriate forwarding class based on the DSCP-to-forwarding class mapping or the LSP EXP-to-forwarding class mapping. The DSCP or LSP EXP bits of such packets are not altered as the packets egress this router, unless **remarking** is enabled.

Remarking can be required if this router is connected to a different DiffServ domain where the DSCP-to-forwarding class mapping is different.

Normally, no remarking is necessary when all router devices are in the same DiffServ domain.

The network QoS policy supports an egress flag that forces remarking of packets that were received on trusted IES and network IP interfaces. This provides the capability of remarking without regard to the ingress state of the IP interface on which a packet was received. The effect of the egress network remark trusted state on each type of ingress IP interface and trust state is listed in Table 82: Ingress IP Interface Type and Trust State Effect on Egress Network Remarking.

The remark trusted state has no effect on packets received on an ingress VPRN IP interface.

Table 82: Ingress IP Interface Type and Trust State Effect on Egress Network Remarking

Ingress IP Interface Type and Trust State	Egress Network IP Interface Trust Remark Disabled (Default)	Egress Network IP Interface Trust Remark Enabled
IES Non-Trusted (Default)	Egress Remarked	Egress Remarked
IES Trusted	Egress Not Remarked	Egress Remarked
VPRN Non-Trusted	Egress Remarked	Egress Remarked

Ingress IP Interface Type and Trust State	Egress Network IP Interface Trust Remark Disabled (Default)	Egress Network IP Interface Trust Remark Enabled
VPRN Trusted (Default)	Egress Not Remarked	Egress Not Remarked
Network Non-Trusted	Egress Remarked	Egress Remarked
Network Trusted (Default)	Egress Not Remarked	Egress Remarked

The **no** form of this command resets the configuration to the default behavior.

Default

no remarking — Remarking disabled in the Network QoS policy.

Parameters

force

Specifies that all IP routed traffic egressing the associated network interface will have its EXP, DSCP, P-bit, and DE bit setting remarked as defined in the associated QoS policy. Only bit fields configured in the QoS policy will be remarked; all others will be left untouched or set based on the default if the fields were not present at ingress.

Platforms

7705 SAR Gen 2

25.70 remote

remote

Syntax

remote

Context

[Tree] (config>ipsec>ts-list remote)

Full Context

configure ipsec ts-list remote

Description

Commands in this context configure remote TS-list parameters. The TS-list is the traffic selector of the local system, such as TSi, when the system acts as an IKEv2 responder.

Platforms

7705 SAR Gen 2

25.71 remote-age

remote-age

Syntax

remote-age *aging-timer*

no remote-age [*aging-timer*]

Context

[Tree] (config>service>template>vpls-template remote-age)

[Tree] (config>service>vpls remote-age)

Full Context

configure service template vpls-template remote-age

configure service vpls remote-age

Description

This command specifies the aging time for remotely learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance.

In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a service destination point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the FDB. The **remote-age** timer specifies the aging time for remote learned MAC addresses. To reduce the amount of signaling required between switches configure this timer larger than the **local-age** timer.

The **no** form of this command returns the remote aging timer to the default value.

Default

remote-age 900

Parameters

seconds

Specifies the aging time for remote MACs expressed in seconds

Values 60 to 86400

Platforms

7705 SAR Gen 2

25.72 remote-attachment-circuit

remote-attachment-circuit

Syntax

remote-attachment-circuit *ac-name* [**endpoint** *endpoint-name*] [**create**]

no remote-attachment-circuit *ac-name*

Context

[Tree] (config>service>epipe>bgp-evpn remote-attachment-circuit)

Full Context

configure service epipe bgp-evpn remote-attachment-circuit

Description

This command configures the remote attachment circuit.

The **no** form of this command disables the context.

Default

no remote-attachment-circuit

Parameters

ac-name

Specifies the name of the remote attachment circuit, up to 32 characters.

endpoint-name

Specifies the name of the endpoint, up to 32 characters.

create

Keyword used to create the remote AC.

Platforms

7705 SAR Gen 2

25.73 remote-gateway-address

remote-gateway-address

Syntax

remote-gateway-address [*ip-address* | *ipv6-address*]

no remote-gateway-address

Context

[Tree] (config>router>if>ipsec>ipsec-tunnel remote-gateway-address)

Full Context

configure router interface ipsec ipsec-tunnel remote-gateway-address

Description

This command configures the remote IPsec tunnel endpoint address.

Parameters

ip-address

Specifies a remote unicast IPv4 address, up to 64 characters.

ipv6-address

Specifies a remote unicast global unicast IPv6 address, up to 64 characters.

Platforms

7705 SAR Gen 2

25.74 remote-id

remote-id

Syntax

remote-id hex *hex-string*

remote-id string *ascii-string*

no remote-id

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>host-ident remote-id)

Full Context

configure subscriber-mgmt local-user-db ipoe host host-identification remote-id

Description

This command specifies the remote ID to match for a host lookup. When the LUDB is accessed using a DHCPv4 server, the SAP-ID is matched against DHCP option 82.

Note:
This command is used only when **remote-id** is configured as one of the **match-list** parameters.

The **no** form of this command removes the remote ID from the configuration.

Parameters

hex-string

Specifies the hexadecimal format for the remote ID.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

ascii-string

Specifies the string format for the remote ID, up to 255 characters.

Platforms

7705 SAR Gen 2

remote-id

Syntax

remote-id

remote-id mac

remote-id string [*string*]

no remote-id

Context

[Tree] (config>service>vprn>if>ipv6>dhcp6>option remote-id)

[Tree] (config>service>ies>if>ipv6>dhcp6>option remote-id)

Full Context

configure service vprn interface ipv6 dhcp6-relay option remote-id

configure service ies interface ipv6 dhcp6-relay option remote-id

Description

This command enables the sending of remote ID option in the DHCPv6 relay packet.

The client DHCP Unique Identifier (DUID) is used as the remote ID.

The **no** form of this command disables the sending of remote ID option in the DHCPv6 relay packet.

Platforms

7705 SAR Gen 2

remote-id

Syntax

remote-id

remote-id hex [*hex-string*]

remote-id {**mac** | **string** *string*}

no remote-id

Context

[Tree] (config>service>ies>if>dhcp>option remote-id)

[Tree] (config>service>vprn>if>dhcp>option remote-id)

[Tree] (config>service>vpls>sap>dhcp>option remote-id)

Full Context

configure service ies interface dhcp option remote-id

configure service vprn interface dhcp option remote-id

configure service vpls sap dhcp option remote-id

Description

This command specifies what information goes into the remote-id sub-option in the DHCP relay packet.

If disabled, the **remote-id** sub-option of the DHCP packet is left empty. When the command is configured without any parameters, it equals to the remote-id mac option.

The **no** form of this command reverts to the default.

Parameters

string

Specifies the remote-id, up to 32 characters.

hex-string

Specifies the hex value of this option.

Values 0x0 to 0xFFFFFFFF...(up to 64 hex nibbles)

mac

Specifies that the MAC address of the remote end is encoded in the sub-option.

Platforms

7705 SAR Gen 2

remote-id

Syntax

remote-id [{**mac** | **string** *string*}]

no remote-id

Context

[Tree] (config>router>if>dhcp>option remote-id)

Full Context

configure router interface dhcp option remote-id

Description

When enabled, the router sends the MAC address of the remote end (typically the DHCP client) in the **remote-id** suboption of the DHCP packet. This command identifies the host at the other end of the circuit. If disabled, the **remote-id** suboption of the DHCP packet will be left empty.

The **no** form of this command returns the system to the default.

Default

no remote-id

Parameters

mac

This keyword specifies the MAC address of the remote end is encoded in the suboption.

string

Specifies the remote ID.

Platforms

7705 SAR Gen 2

remote-id

Syntax

remote-id mac

remote-id string <*string*>

no remote-id

Context

[Tree] (config>service>vpls>sap>dhcp6>ldra>options remote-id)

Full Context

configure service vpls sap dhcp6 ldra options remote-id

Description

This command configures the information for the remote ID suboption in the DHCP6 LDRA.

The **no** form of this command reverts to the default.

Default

no remote-id

Parameters

mac

Sets the enterprise number field of the Relay Agent remote ID to 6527 and configures the DHCPv6 client source MAC address as six hexadecimal numbers.

string

Sets the enterprise number field of the Relay-Agent remote ID to 6527 and configures the ASCII-encoded string using up to 32 characters.

Platforms

7705 SAR Gen 2

25.75 remote-ip

remote-ip

Syntax

remote-ip *ip-address*

no remote-ip

Context

[Tree] (config>service>ies>if>sap>ip-tunnel remote-ip)

Full Context

configure service ies interface sap ip-tunnel remote-ip

Description

This command configures the primary destination IPv4 or IPv6 address to use for an IP tunnel. This configuration applies to the outer IP header of the encapsulated packets. The **source** address, **remote-ip** address and **backup-remote-ip** address of a tunnel must all belong to the same address family (IPv4 or IPv6). When the remote-ip address contains an IPv6 address it must be a global unicast address.

Default

no remote-ip

Parameters

ip-address

An IPv4 address or an IPv6 address.

Platforms

7705 SAR Gen 2

remote-ip

Syntax

remote-ip *ip-address*

no remote-ip

Context

[Tree] (config>service>vprn>if>sap>ip-tunnel remote-ip)

Full Context

configure service vprn interface sap ip-tunnel remote-ip

Description

This command sets the primary destination IPv4 address of GRE encapsulated packets associated with a particular GRE tunnel. If this address is reachable in the delivery service (there is a route) then this is the destination IPv4 address of GRE encapsulated packets sent by the delivery service.

The **no** form of this command deletes the destination address from the GRE tunnel configuration.

Parameters

ip-address

Specifies the destination IPv4 address of the GRE tunnel.

Values 1.0.0.0 to 223.255.255.255

Platforms

7705 SAR Gen 2

remote-ip

Syntax

remote-ip {*ip-prefix/prefix-length* | *ip-prefix netmask* | **any**}

Context

[Tree] (config>router>ipsec>sec-plcy>entry remote-ip)

[Tree] (config>service>vprn>ipsec>sec-plcy>entry remote-ip)

Full Context

configure router ipsec security-policy entry remote-ip

configure service vprn ipsec security-policy entry remote-ip

Description

This command configures the remote (from the tunnel) IP prefix/mask for the policy parameter entry.

Only one entry is necessary to describe a potential flow. The **local-ip** and **remote-ip** commands can be defined only once. The system evaluates:

- the local IP as the source IP when traffic is examined in the direction of the flows from private to public and as the destination IP when traffic flows from public to private
- the remote IP as the source IP when traffic flows public to private and as the destination IP when traffic flows from private to public

Parameters

ip-prefix

Specifies the destination address of the aggregate route in dotted decimal notation.

Values a.b.c.d (host bits must be 0)

prefix-length 1 to 32

netmask

Specifies the subnet mask in dotted decimal notation.

any

keyword to specify that it can be any address.

Platforms

7705 SAR Gen 2

25.76 remote-lfa

remote-lfa

Syntax

remote-lfa [**max-pq-cost** *value*]

no remote-lfa

Context

[Tree] (config>router>isis>loopfree-alternates remote-lfa)

Full Context

configure router isis loopfree-alternates remote-lfa

Description

This command enables the use of the Remote LFA algorithm in the LFA SPF calculation for this ISIS instance.

The **no** form of this command disables the use of the Remote LFA algorithm in the LFA SPF calculation for this ISIS instance.

Default

no remote-lfa

Parameters

value

Specifies the integer used to limit the search of candidate P and Q nodes in the remote LFA by setting the maximum IGP cost from the router performing the remote LFA calculation to the candidate P or Q node.

Values 0 to 4294967295

Default 4261412864

Platforms

7705 SAR Gen 2

remote-lfa

Syntax

remote-lfa [**max-pq-cost** *value*]

no remote-lfa

Context

[Tree] (config>router>ospf>loopfree-alternates remote-lfa)

Full Context

configure router ospf loopfree-alternates remote-lfa

Description

This command enables the use of the Remote LFA algorithm in the LFA SPF calculation in this OSPF or OSPF3 instance.

The **no** form of this command disables the use of the Remote LFA algorithm in the LFA SPF calculation in this OSPF or OSPF3 instance.

Default

no remote-lfa

Parameters

max-pq-cost value

Specifies the integer used to limit the search of candidate P and Q nodes in the remote LFA by setting the maximum IGP cost from the router performing the remote LFA calculation to the candidate P or Q node.

Values 0 to 4294967295

Default 4261412864

Platforms

7705 SAR Gen 2

25.77 remote-management

remote-management

Syntax

remote-management

Context

[Tree] (config>system>management-interface remote-management)

Full Context

configure system management-interface remote-management

Description

Commands in this context configure the SR OS node to use the remote management service. Configuring remote management enables the SR OS node to report itself to a remote manager service running on a remote server, so that it is included in the dynamic list of available nodes. The manager service streamlines the management of multiple SR OS nodes running different SR OS versions using the same client application providing a similar shell to the MD-CLI.

Platforms

7705 SAR Gen 2

remote-management

Syntax

remote-management

no remote-management

remote-management manager [*manager-name*]

no remote-management manager [*manager-name*]

Context

[Tree] (debug>system>management-interface remote-management)

Full Context

debug system management-interface remote-management

Description

This command configures the management interface to debug the **remote-management** managers.

The **no** form of this command removes the configuration.

Parameters

manager manager-name

Specifies the name of the manager, up to 64 characters. If the parameter is not specified, all configured managers are debugged.

Platforms

7705 SAR Gen 2

25.78 remote-max-checkpoints

remote-max-checkpoints

Syntax

remote-max-checkpoints [*number-of-files*]

no remote-max-checkpoints

Context

[Tree] (config>system>rollback remote-max-checkpoints)

Full Context

configure system rollback remote-max-checkpoints

Description

This command configures the maximum number of rollback checkpoint files when the rollback-location is remote (for example, ftp).

Default

no remote-max-checkpoints

Parameters

number of files

Specifies the maximum rollback files saved at a remote location.

Values 1 to 200

Platforms

7705 SAR Gen 2

25.79 remote-proxy-arp

remote-proxy-arp

Syntax

[**no**] **remote-proxy-arp**

Context

[Tree] (config>service>vprn>if remote-proxy-arp)

[Tree] (config>service>ies>if remote-proxy-arp)

Full Context

configure service vprn interface remote-proxy-arp

configure service ies interface remote-proxy-arp

Description

This command enables remote proxy ARP on the interface.

Remote proxy ARP is similar to proxy ARP. It allows the router to answer an ARP request on an interface for a subnet that is not provisioned on that interface. This allows the router to forward to the other subnet on behalf of the requester. To distinguish remote proxy ARP from local proxy ARP, local proxy ARP performs a similar function but only when the requested IP is on the receiving interface.

The **no** form of this command reverts to the default.

Platforms

7705 SAR Gen 2

remote-proxy-arp

Syntax

[**no**] **remote-proxy-arp**

Context

[Tree] (config>router>if remote-proxy-arp)

Full Context

configure router interface remote-proxy-arp

Description

This command enables remote proxy ARP on the interface.

Default

no remote-proxy-arp

Platforms

7705 SAR Gen 2

25.80 remote-servers

remote-servers

Syntax

remote-servers

Context

[Tree] (config>service>vprn>aaa remote-servers)

Full Context

configure service vprn aaa remote-servers

Description

Commands in this context configure AAA remote servers on the VPRN.

Platforms

7705 SAR Gen 2

25.81 remote-source

remote-source

Syntax

[**no**] **remote-source**

Context

[Tree] (config>mirror>mirror-dest remote-source)

Full Context

configure mirror mirror-dest remote-source

Description

This command is used on a destination router in a remote mirroring solution. The mirroring (packet copy) is performed on the source router and sent via an SDP to the destination router. Remote mirroring requires remote source configuration on the destination router.

Remote mirroring allows a destination router to terminate SDPs from multiple remote source routers. This allows consolidation of packet sniffers or analyzers at a single or small set of points in a network (for example, a sniffer or analyze farm, or lawful interception gateway).

A **remote-source** entry must be configured on the destination router for each source router from which mirrored traffic is being sent via SDPs.

A mirror destination service that is configured for a destination router must not be configured as for a source router.

The emote source configuration is not applicable when routable LI encapsulation is being used on the mirror source router. The remote source configuration is only used when a source router is sending mirrored traffic to a destination router via SDPs.

Two types of remote-source entries can be configured:

- far end
- spoke SDP

Certain remote source types are applicable with certain SDP types. For descriptions of the command usage in the **mirror-dest** context, see the **far-end** and **spoke-sdp** commands.

The **no** form of this command removes all remote-source entries.

Platforms

7705 SAR Gen 2

25.82 remote-v6-ip

remote-v6-ip

Syntax

remote-v6-ip any

remote-v6-ip *ipv6-prefix/prefix-length*

no remote-v6-ip

Context

[Tree] (config>service>vprn>ipsec>sec-plcy>entry remote-v6-ip)

[Tree] (config>router>ipsec>sec-plcy>entry remote-v6-ip)

Full Context

configure service vprn ipsec security-policy entry remote-v6-ip

configure router ipsec security-policy entry remote-v6-ip

Description

This command specifies the remote v6 prefix for the security-policy entry.

Parameters

ipv6-prefix/prefix-length

Specifies the local v6 prefix and length.

Values	ipv6-address/prefix: ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x [0 to FFFF]H
		d [0 to 255]D
		host bits must be 0
		:: not allowed
		prefix-length [1 to 28]

any

A keyword to specify that any address can be used.

Platforms

7705 SAR Gen 2

25.83 remote-ve-name

remote-ve-name

Syntax

[**no**] **remote-ve-name** *name*

Context

[Tree] (config>service>epipe>bgp-vpws remote-ve-name)

Full Context

configure service epipe bgp-vpws remote-ve-name

Description

This command creates or edits a remote-ve-name. A single remote-ve-name can be created per BGP VPWS instance if the service is single-homed or uses a single pseudowire to connect to a pair of dual-homed systems. When the service requires active/standby pseudowires to be created to remote dual-homed systems then two remote-ve-names must be configured.

This context defines the remote PE to which a pseudowire will be signaled.

remote-ve-name commands can be added even if bgp-vpws is not shutdown.

The **no** form of this command removes the configured remote-ve-name from the bgp vpws node. It can be used when the BGP VPWS status is either shutdown or "no shutdown”.

Parameters

name

Specifies a site name up to 32 characters in length.

Platforms

7705 SAR Gen 2

25.84 remove

remove

Syntax

[**no**] **remove**

Context

[Tree] (config>service>vprn>bgp>attribute-set remove)

Full Context

configure service vprn bgp attribute-set remove

Description

This command configures BGP to ignore and silently discard ATTR_SETs in BGP routes received from PE-CE peers of the VPRN. The discarded ATTR_SETs do not affect BGP best-path selection in the VPRN, and they do not appear in the VPN-IP routes that result from the VRF export of the BGP routes. Nokia recommends enabling this command in most deployments.

The **no** form of this command configures BGP to ignore ATTR_SETs in BGP routes received from PE-CE peers of the VPRN without discarding them. This allows the ATTR_SETs to propagate between CE devices connected to the VPRN and to other PE devices when the BGP routes are exported as VPN-IP routes.

Note: If the configuration of this command is changed, ROUTE_REFRESH messages are sent to all PE-CE peers of the VPRN.

Default

no remove

Platforms

7705 SAR Gen 2

25.85 remove-private

remove-private

Syntax

remove-private [**limited**] [**skip-peer-as**] [**replace**]

no remove-private

Context

[Tree] (config>service>vprn>bgp>group remove-private)

[Tree] (config>service>vprn>bgp>group>neighbor remove-private)

[Tree] (config>service>vprn>bgp remove-private)

Full Context

configure service vprn bgp group remove-private

configure service vprn bgp group neighbor remove-private

configure service vprn bgp remove-private

Description

When this command is configured private AS numbers are removed or replaced when they are found inside the AS path of BGP routes advertised to peers within the scope of the command.

The set of AS numbers that are defined by IANA as private are in the range of 64512 to 65534, and 4200000000 to 4294967294, inclusive. In SR OS, this command also removes ASN 65535 and ASN 4294967295, which are reserved values.

The **no** form of this command (at the BGP instance level) implements the default behavior, private AS numbers are allowed without restriction or modification in routes advertised to peers.

Default

no remove-private

Parameters

limited

This keyword instructs BGP to process private ASNs only up to the first public ASN encountered. Private ASNs beyond that first public AS will not be stripped or replaced.

skip-peer-as

This keyword instructs BGP to not strip or replace a private ASN from the AS-Path if that ASN is the same as the BGP peer AS number.

replace

When this keyword is configured, private ASNs are not stripped. Each occurrence is replaced by the ASN of the advertising BGP router (the ASN the router advertised to its peer in its OPEN message). When the **replace** keyword is not configured, private ASNs are stripped, subject to influence by the other keyword options. This generally results in a shortening of AS_PATH length.

Platforms

7705 SAR Gen 2

remove-private

Syntax

remove-private [**limited**] [**skip-peer-as**] [**replace**]

no remove-private

Context

[Tree] (config>router>bgp remove-private)

[Tree] (config>router>bgp>group remove-private)

[Tree] (config>router>bgp>group>neighbor remove-private)

Full Context

configure router bgp remove-private

configure router bgp group remove-private

configure router bgp group neighbor remove-private

Description

When this command is configured private AS numbers are removed or replaced when they are found inside the AS path of BGP routes advertised to peers within the scope of the command.

The set of AS numbers that are defined by IANA as private are in the range of 64512 to 65534, and 4200000000 to 4294967294, inclusive. In SR OS, this command also removes ASN 65535 and ASN 4294967295, which are reserved values.

The **no** form of this command (at the BGP instance level) implements the default behavior, private AS numbers are allowed without restriction or modification in routes advertised to peers.

Default

no remove-private

Parameters

limited

This keyword instructs BGP to process private ASNs only up to the first public ASN encountered. Private ASNs beyond that first public AS will not be stripped or replaced.

skip-peer-as

This keyword instructs BGP to not strip or replace a private ASN from the AS-Path if that ASN is the same as the BGP peer AS number.

replace

When this keyword is configured, private ASNs are not stripped. Each occurrence is replaced by the ASN of the advertising BGP router (the ASN the router advertised to its peer in its OPEN message). When the **replace** keyword is not configured, private ASNs are stripped, subject to influence by the other keyword options. This generally results in a shortening of AS_PATH length.

Platforms

7705 SAR Gen 2

25.86 renew

renew

Syntax

renew est-profile *name* **cert** *cert-filename* **key** *key-filename* [**hash-alg** *hash-algorithm*] **output** *output-cert-filename* [**validate-cert-chain**] [**force**]

Context

[Tree] (admin>certificate>est renew)

Full Context

admin certificate est renew

Description

This command renews an imported certificate (specified by the **cert** *cert-filename*) with a Certificate Authority (CA) using the EST protocol specified by the **est-profile** name, with an imported private key specified the key parameter. The key can be either the key of the certificate to be renewed or a new key.

The authentication between system and EST server is specified by the est-profile.

The **hash-alg** *hash-alorithm* parameter is used to generate the CSR (Certificate Signing Request) in the EST request message.

Parameters

name

Specifies EST profile name, up to 32 characters

cert-filename

Specifies the certificate file name, up to 95 characters

key-filename

Specifies the file name of a key, up to 95 characters

hash-algorithm

Specifies the hash algorithm to be used in a certificate request.

Values sha1, sha224, sha256, sha384, sha512

output-cert-filename

Specifies the output cert file name, up to 200 characters

validate-cert-chain

Specifies that the the system validates the certificate chain of the result certificate before importing it

force

Specifies the system to overwrite the existing file with same **output** *output-cert-filename*

Platforms

7705 SAR Gen 2

25.87 renew-timer

renew-timer

Syntax

renew-timer [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no renew-timer

Context

[Tree] (config>router>dhcp6>server>pool>prefix renew-timer)

[Tree] (config>service>vprn>dhcp6>server>pool>prefix renew-timer)

Full Context

configure router dhcp6 local-dhcp-server pool prefix renew-timer

configure service vprn dhcp6 local-dhcp-server pool prefix renew-timer

Description

This command configures the lease renew time (T1) via LUDB.

The T1 is the time at which the client contacts the addressing authority to extend the lifetimes of the DHCPv6 leases (addresses or prefixes). T1 is a time duration relative to the current time expressed in units of seconds.

The IP addressing authority controls the time at which the client contacts the addressing authority to extend the lifetimes on assigned addresses through the T1 and T2 parameters assigned to an IA. At time T1 for an IA, the client initiates a Renew/Reply message exchange to extend the lifetimes on any addresses in the IA. The client includes an IA option with all addresses currently assigned to the IA in its Renew message. Recommended values for T1 and T2 are .5 and .8 times the shortest preferred lifetime of the addresses in the IA that the addressing authority is willing to extend, respectively.

The configured renew timer should always be smaller than or equal to the rebind timer.

The T1 and T2 are carried in the IPv6 address option that is within the IA.

The **no** form of this command reverts to the default.

Default

renew-timer min 30

Parameters

renew-timer

Specifies the preferred lifetime.

Values	**days** *days*	0 to 7
	hrs *hours*	0 to 23
	min *minutes*	0 to 59
	sec *seconds*	0 to 59

Platforms

7705 SAR Gen 2

25.88 renum

renum

Syntax

renum *old-entry-id new-entry-id*

Context

[Tree] (config>qos>sap-ingress>ipv6-criteria renum)

[Tree] (config>qos>sap-ingress>mac-criteria renum)

[Tree] (config>qos>sap-egress>ip-criteria renum)

[Tree] (config>qos>sap-egress>ipv6-criteria renum)

[Tree] (config>qos>sap-ingress>ip-criteria renum)

Full Context

configure qos sap-ingress ipv6-criteria renum

configure qos sap-ingress mac-criteria renum

configure qos sap-egress ip-criteria renum

configure qos sap-egress ipv6-criteria renum

configure qos sap-ingress ip-criteria renum

Description

This command renumbers existing QoS policy criteria entries to properly sequence policy entries.

This can be required in some cases since the router exits when the first match is found and executes the actions in accordance with the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.

Parameters

old-entry-id

Enter the entry number of an existing entry.

Values 1 to 65535

new-entry-id

Enter the new entry number to be assigned to the old entry.

Values 1 to 65535

Platforms

7705 SAR Gen 2

renum

Syntax

renum *old-entry-number new-entry-number*

Context

[Tree] (config>qos>network>egress>ipv6-criteria renum)

[Tree] (config>qos>network>ingress>ip-criteria renum)

[Tree] (config>qos>network>ingress>ipv6-criteria renum)

[Tree] (config>qos>network>egress>ip-criteria renum)

Full Context

configure qos network egress ipv6-criteria renum

configure qos network ingress ip-criteria renum

configure qos network ingress ipv6-criteria renum

configure qos network egress ip-criteria renum

Description

This command renumbers existing QoS policy criteria entries to properly sequence policy entries.

This can be required in some cases since the router exits when the first match is found and executes the actions in accordance with the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.

Parameters

old-entry-number

Enter the entry number of an existing entry.

Values 1 to 65535

new-entry-number

Enter the new entry number to be assigned to the old entry.

Values 1 to 65535

Platforms

7705 SAR Gen 2

renum

Syntax

renum *old-entry-id new-entry-id*

Context

[Tree] (config>filter>ip-filter renum)

[Tree] (config>filter>ipv6-exception renum)

[Tree] (config>filter>ip-exception renum)

[Tree] (config>filter>ipv6-filter renum)

Full Context

configure filter ip-filter renum

configure filter ipv6-exception renum

configure filter ip-exception renum

configure filter ipv6-filter renum

Description

This command renumbers existing MAC, IPv4/IPv6, IP exception filter, or IPv6 exception filter entries to properly sequence filter entries.

This may be required in some cases since the OS exits when the first match is found and executes the actions according to the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.

Parameters

old-entry-id

Specifies the entry number of an existing entry, as a decimal integer.

Values 1 to 2097151

new-entry-id

Specifies the new entry-number to be assigned to the old entry, as a decimal integer.

Values 1 to 2097151

Platforms

7705 SAR Gen 2

renum

Syntax

renum *old-entry-number new-entry-number*

Context

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter renum)

[Tree] (config>system>security>mgmt-access-filter>mac-filter renum)

[Tree] (config>system>security>mgmt-access-filter>ip-filter renum)

Full Context

configure system security management-access-filter ipv6-filter renum

configure system security management-access-filter mac-filter renum

configure system security management-access-filter ip-filter renum

Description

This command renumbers existing management access filter entries for an IP(v4), IPv6, or MAC filter to re-sequence filter entries.

The exits on the first match found and executes the actions in accordance with the accompanying **action** command. This may require some entries to be re-numbered differently from most to least explicit.

Parameters

old-entry-number

Specifies the entry number of the existing entry.

Values 1 to 9999

new-entry-number

Specifies the new entry number that will replace the old entry number.

Values 1 to 9999

Platforms

7705 SAR Gen 2

renum

Syntax

renum *old-entry-number new-entry-number*

Context

[Tree] (config>system>security>profile renum)

Full Context

configure system security profile renum

Description

This command renumbers profile entries to re-sequence the entries.

Since the OS exits when the first match is found and executes the actions according to accompanying action command, re-numbering is useful to rearrange the entries from most explicit to least explicit.

Parameters

old-entry-number

Enter the entry number of an existing entry.

Values 1 **to** 9999

new-entry-number

Enter the new entry number.

Values 1 **to** 9999

Platforms

7705 SAR Gen 2

26 r Commands – Part II

26.1 renumber

renumber

Syntax

renumber from *entry-id* **to** *entry-id*

Context

[Tree] (config>router>policy-options>policy-statement renumber)

Full Context

configure router policy-options policy-statement renumber

Description

This command allows the operator to renumber the existing entry ID to a new entry ID. When performing the renumbering action, the two entry IDs must be different. The existing (**from**) *entry-id* must exist. The new (**to**) *entry-id* must not exist.

Renumbering is not saved in the configuration because it is a performing action.

Parameters

from *entry-id*

Specifies the existing entry ID to be renumbered.

Values 1 to 4294967295

to *entry-id*

Specifies the new entry ID to be assigned.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

26.2 repair

repair

Syntax

repair [*cflash-id*]

Context

[Tree] (file repair)

Full Context

file repair

Description

This command checks a compact flash device for errors and repairs any errors found.

Parameters

cflash-id

Specifies the compact flash slot ID to be repaired. When a specific *cflash-id* is specified, that drive is repaired. If no *flash-id* is specified, the drive referred to by the current working directory is assumed. If a slot number is not specified, the active CPM is assumed.

Values cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Default the current compact flash device

Platforms

7705 SAR Gen 2

26.3 repeated-characters

repeated-characters

Syntax

repeated-characters *count*

no repeated-characters

Context

[Tree] (config>system>security>password>complexity-rules repeated-characters)

Full Context

configure system security password complexity-rules repeated-characters

Description

The number of times a characters can be repeated consecutively.

The **no** form of this command resets to default.

Default

no repeated-characters

Parameters

count

Specifies the minimum count of consecutively repeated characters.

Values 2 to 8

Platforms

7705 SAR Gen 2

26.4 replace

replace

Syntax

replace [*line*]

Context

[Tree] (candidate replace)

Full Context

candidate replace

Description

This command displays the specified line (a single line only) and allows it to be changed.

Parameters

line

Indicates which line to replace starting at the point indicated by the following options.

Values line, offset, **first**, **edit-point**, **last**

line	absolute line number
offset	relative line number to current edit point. Prefixed with '+' or '-'
first	keyword - first line
edit-point	keyword - current edit point
last	keyword - last line that is not 'exit'

Platforms

7705 SAR Gen 2

26.5 replay-protection

replay-protection

Syntax

[**no**] **replay-protection**

Context

[Tree] (config>macsec>connectivity-association replay-protection)

Full Context

configure macsec connectivity-association replay-protection

Description

Specifies the size of the replay protection window.

This command must be configured to force packet discard when it has detected a packet that is not within the replay-window-size.

When replay protection is enabled, the sequence of the ID number of the received packets are checked. If the packet arrives out of sequence and the difference between the packet numbers exceeds the replay window size, the packet is counted by the receiving port and then discarded. For example, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is counted and discarded because it falls outside the parameters of the replay window size.

Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

Default

no replay-protection

Platforms

7705 SAR Gen 2

26.6 replay-window

replay-window

Syntax

replay-window *replay-window-size*

no replay-window

Context

[Tree] (config>ipsec>trans-mode-prof replay-window)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel replay-window)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel replay-window)

[Tree] (config>ipsec>tnl-temp replay-window)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel replay-window)

Full Context

configure ipsec ipsec-transport-mode-profile replay-window

configure service ies interface ipsec ipsec-tunnel replay-window

configure service vprn interface sap ipsec-tunnel replay-window

configure ipsec tunnel-template replay-window

configure service vprn interface ipsec ipsec-tunnel replay-window

Description

This command specifies the size of the anti-replay window. The anti-replay window protocol further secures IPsec against an entity that can inject a recorded message in a message stream from a source to a destination computer on the Internet.

Default

no replay-window

Parameters

replay-window-size

Specifies the size of the SA anti-replay window.

Values 32, 64, 128, 256, 512

Platforms

7705 SAR Gen 2

26.7 replay-window-size

replay-window-size

Syntax

replay-window-size *number-of-packet*s

no replay-window-size

Context

[Tree] (config>macsec>connectivity-association replay-window-size)

Full Context

configure macsec connectivity-association replay-window-size

Description

This command specifies the size of the replay protection window.

This command must be configured to enable replay protection. When replay protection is enabled, the sequence of the ID number of received packets are checked. If the packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving port. For example, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.

Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

When the *number-of-packets* variable is set to 0, all packets that arrive out-of-order are dropped.

The **no** form of this command reverts to the default value.

Default

replay-window-size 0

Parameters

number-of-packets

Specifies the window for which the packets can arrive out of order.

Values 0 to 4294967294

Platforms

7705 SAR Gen 2

26.8 report-alarm

report-alarm

Syntax

[**no**] **report-alarm** [**signal-fail**] [**remote**] [**local**] [**no-frame-lock**] [**high-ber**] [**no-block-lock**] [**no-am-lock**] [**duplicate-lane**]

Context

[Tree] (config>port>ethernet report-alarm)

Full Context

configure port ethernet report-alarm

Description

This command specifies when and if to generate alarms and alarm clear notifications for this port.

Note:
For some DWDM transceivers, if the **configure port dwdm coherent rx-los-reaction squelch** command is disabled the signal-fail and no-am-lock alarm conditions are not reported when the media side of the transceiver has an RX LOS condition.

Parameters

signal-fail

Reports an Ethernet signal lost alarm.

remote

Reports remote faults.

local

Reports local faults.

no-frame-lock

Reports a 'not locked on the Ethernet framing sequence' alarm.

high-ber

Reports High Bit Error Rate.

no-block-lock

Reports 40G/100G PCS Lanes Not Block Locked.

no-am-lock

Reports 40G/100G PCS Alignment Marker Loss of Lock.

duplicate-lane

Reports 40G/100G PCS Duplicate Lane Marker.

Platforms

7705 SAR Gen 2

26.9 report-alarms

report-alarms

Syntax

[**no**] **report-alarms** [**modflt**] [**mod**] [**netrx**] [**nettx**] [**hosttx**]

Context

[Tree] (config>port>dwdm>coherent report-alarms)

Full Context

configure port dwdm coherent report-alarms

Description

This command configures the alarms that will be reported for the coherent module.

Default

modflt mod netrx nettx hosttx

Parameters

modflt

Reports module fault alarm.

mod

Reports module alarm.

netrx

Reports network (optical side) receive alarm.

nettx

Reports network (optical side) transmit alarm.

hosttx

Reports host (electrical side) transmit alarm.

Platforms

7705 SAR Gen 2

26.10 report-path-constraints

report-path-constraints

Syntax

report-path-constraints

no report-path-constraints

Context

[Tree] (config>router>pcep>pcc report-path-constraints)

Full Context

configure router pcep pcc report-path-constraints

Description

This command enables the inclusion of LSP path constraints in the PCE report messages sent from the PCC to a PCE.

In order for the PCE to know about the original constraints for an LSP which is delegated, but for which there is no prior state in its LSP database, such as if no PCReq message was sent for the same PLSP-ID, the following proprietary behavior is observed:

- PCC appends a duplicate of each of the LSPA, METRIC, and BANDWIDTH objects in the PCRpt message. The only difference between two objects of the same type is that the P-flag is set in the common header of the duplicate object to indicate that it is a mandatory object for processing by PCE.
- The value of the metric or bandwidth in the duplicate object contains the original constraint value, while the first object contains the operational value. This is applicable to hop metrics in the METRIC and BANDWIDTH objects only. The SR OS PCC does not support configuring a boundary on the path computation IGP or TE metrics.
- The path computation on the PCE must use the first set of objects when updating a path if the PCRpt contained a single set. If the PCRpt contained a duplicate set, PCE path computation must use the constraints in the duplicate set.

The **no** form of the command disables the above behavior in case of interoperability issues with third-party PCE implementations.

Default

report-path-constraints

Platforms

7705 SAR Gen 2

26.11 report-src-ip

report-src-ip

Syntax

report-src-ip *ip-address*

no report-src-ip

Context

[Tree] (config>service>vpls>igmp-snooping report-src-ip)

Full Context

configure service vpls igmp-snooping report-src-ip

Description

This command configures the source IPv4 address used when generating IGMP reports. According the IGMPv3 standard, a zero source address is allowed in sending IGMP reports. However, for interoperability with some multicast routers, the source IP address of IGMP group reports can be configured using this command.

Default

report-src-ip 0.0.0.0

Parameters

ip-address

Specifies the source IPv4 address in transmitted IGMP reports.

Values a.b.c.d

Platforms

7705 SAR Gen 2

report-src-ip

Syntax

report-src-ip *ipv6-address*

no report-src-ip

Context

[Tree] (config>service>vpls>mld-snooping report-src-ip)

Full Context

configure service vpls mld-snooping report-src-ip

Description

This command configures the source IPv6 address used when generating MLD reports. A zero source address is allowed in sending MLD reports. However, for interoperability with some multicast routers, the source IP address of MLD reports can be configured using this command.

Default

report-src-ip 0:0:0:0:0:0:0:0

Parameters

ipv6-address

Specifies the source IPv6 address in transmitted MLD reports.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)

Platforms

7705 SAR Gen 2

26.12 request-format

request-format

Syntax

request-format

Context

[Tree] (config>service>vprn>aaa>rmt-srv>tacplus request-format)

[Tree] (config>system>security>tacplus request-format)

Full Context

configure service vprn aaa remote-servers tacplus request-format

configure system security tacplus request-format

Description

Commands in this context configure access operations that are sent to the TACACS+ server during authorization.

Platforms

7705 SAR Gen 2

26.13 request-timer

request-timer

Syntax

request-timer *timer1* **retry-timer** *timer2* **timeout-multiplier** *multiplier*

no request-timer

Context

[Tree] (config>service>epipe>spoke-sdp>control-channel-status request-timer)

[Tree] (config>service>vpls>spoke-sdp>control-channel-status request-timer)

Full Context

configure service epipe spoke-sdp control-channel-status request-timer

configure service vpls spoke-sdp control-channel-status request-timer

Description

This command configures the control channel status request mechanism. When it is configured, control channel status request procedures are used. These augment the procedures for control channel status messaging from RFC 6478. This command cannot be used with a non-zero refresh-timer value.

Parameters

timer1

Specifies the interval, in seconds, at which pseudowire status messages, including a reliable delivery TLV with the "request" bit set, are sent.

Values 10 to 65535

timer2

specifies the timeout interval, in seconds, if no response to a pseudowire status request is received. This parameter must be configured. A value of zero (0) disables retries.

Values 0, 3 to 60

multiplier

If a requesting node does not receive a valid response to a pseudowire status request within a number of seconds equal to the retry timer multiplied by this multiplier, then it will assume the pseudowire is down. This parameter is optional.

Values 3 to 20

Platforms

7705 SAR Gen 2

request-timer

Syntax

request-timer *request-timer-secs* **retry-timer** *retry-timer-secs* **timeout-multiplier** *multiplier*

no request-timer

Context

[Tree] (config>service>vpls>spoke-sdp>control-channel-status request-timer)

Full Context

configure service vpls spoke-sdp control-channel-status request-timer

Description

This command configures the control channel status request mechanism. When it is configured, control channel status request procedures are used. These augment the procedures for control channel status messaging from RFC 6478, *Pseudowire Status for Static Pseudowires*. This command cannot be used with a non-zero refresh-timer value.

Parameters

request-timer-secs

Specifies the interval, in seconds, at which pseudowire status messages, including a reliable delivery TLV with the "request" bit set, are sent.

Values 10 to 65535

retry-timer-secs

specifies the timeout interval, in seconds, if no response to a pseudowire status request is received. This parameter must be configured. A value of zero (0) disables retries.

Values 0, 3 to 60

multiplier

If a requesting node does not receive a valid response to a pseudowire status request within a number of seconds equal to the retry timer multiplied by this multiplier, then it will assume the pseudowire is down. This parameter is optional.

Values 3 to 20

Platforms

7705 SAR Gen 2

request-timer

Syntax

request-timer *timer1* **retry-timer** *timer2* **timeout-multiplier** *multiplier*

no request-timer

Context

[Tree] (config>service>ies>if>spoke-sdp>control-channel-status request-timer)

Full Context

configure service ies interface spoke-sdp control-channel-status request-timer

Description

This command configures the control channel status request mechanism. When it is configured, control channel status request procedures are used. These augment the procedures for control channel status messaging from RFC 6478. This command is mutually exclusive with a non-zero refresh-timer value.

Parameters

timer1

Specifies the interval at which pseudowire status messages, including a reliable delivery TLV, with the "request" bit set, are sent.

Values 10 to 65535 seconds

retry-timer *timer2*

Specifies the timeout interval if no response to a pseudowire status request is received. This parameter must be configured. A value of zero (0) disables retries.

Values 0, 3 to 60 seconds

timeout-multiplier *multiplier*

If a requesting node does not receive a valid response to a pseudowire status request within this multiplier times the retry timer, then it will assume the pseudowire is down. This parameter is optional.

Values 3 to 20 seconds

Platforms

7705 SAR Gen 2

request-timer

Syntax

request-timer *request-timer-secs* **retry-timer** *retry-timer-secs* **timeout-multiplier** *multiplier*

no request-timer

Context

[Tree] (config>service>vprn>if>spoke-sdp>control-channel-status request-timer)

Full Context

configure service vprn interface spoke-sdp control-channel-status request-timer

Description

This command configures the control channel status request mechanism. When it is configured, control channel status request procedures are used. These augment the procedures for control channel status messaging from RFC 6478. This command is mutually exclusive with a non-zero refresh-timer value.

Parameters

request-timer-secs

Specifies the interval, in seconds, at which pseudowire status messages, including a reliable delivery TLV, with the "request" bit set, are sent.

Values 10 to 65535

retry-timer *retry-timer-secs*

Specifies the timeout interval, in seconds, if no response to a pseudowire status request is received. This parameter must be configured. A value of zero (0) disables retries.

Values 0, 3 to 60

timeout-multiplier *multiplier*

Specifies the multiplier, in seconds. If a requesting node does not receive a valid response to a pseudowire status request within this multiplier times the retry timer, then it assume the pseudowire is down. This parameter is optional.

Values 3 to 15

Platforms

7705 SAR Gen 2

26.14 requests

requests

Syntax

[**no**] **requests** [**neighbor** *ip-int-name* | *ip-address*]

Context

[Tree] (debug>router>rip requests)

Full Context

debug router rip requests

Description

This command enables debugging for RIP requests.

Parameters

ip-int-name* | *ip-address

Debugs the RIP requests sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

requests

Syntax

[**no**] **requests** [**neighbor** *ip-int-name* | *ipv6-address*]

Context

[Tree] (debug>router>ripng requests)

Full Context

debug router ripng requests

Description

This command enables debugging for RIP requests.

Parameters

ip-int-name*| *ipv6-address

Debugs the RIP requests sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

26.15 required

required

Syntax

required [**lowercase** *count*] [**uppercase** *count*] [**numeric** *count*] [**special-character** *count*]

no required

Context

[Tree] (config>system>security>password>complexity-rules required)

Full Context

configure system security password complexity-rules required

Description

Force the minimum number of different character classes required.

The **no** form of this command resets to default.

Default

required lowercase 0 uppercase 0 numeric 0 special-character 0

Parameters

count

Specifies the minimum count of characters classes.

Values 0 to 10

Platforms

7705 SAR Gen 2

26.16 rescue-location

rescue-location

Syntax

rescue-location *file-url*

no rescue-location

Context

[Tree] (config>system>rollback rescue-location)

Full Context

configure system rollback rescue-location

Description

The location and filename of the rescue configuration is configurable to be local (on compact flash) or remote. The suffix .rc will be automatically appended to the filename when a rescue configuration file is saved. Trivial FTP (TFTP) is not supported for remote locations.

Default

no rescue location

Parameters

file-url

Specifies the URL or filename.

Values

local-url | *remote-url*	
local-url	[*cflash-id/*][*file-path*] up to 200 characters, including cflash-id directory length of up to 99 characters each
remote-url	[{ftp://}*login*:*pswd*@ *remote-locn*/][*file-path*] up to 255 characters, directory length of up to 99 characters each
remote-locn	[*hostname* | *ipv4-address* | *ipv6-address*]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D
	interface - 32 chars max, for link local addresses
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:
rescue filename	suffixed with .rc during the rescue file creation

Platforms

7705 SAR Gen 2

26.17 reserved-label-block

reserved-label-block

Syntax

[**no**] **reserved-label-block** *name*

Context

[Tree] (config>router>mpls-labels reserved-label-block)

Full Context

configure router mpls-labels reserved-label-block

Description

Commands in this context configure a block of labels from the dynamic range to be locally assigned for specific applications, such as Segment Routing adjacency SIDs. The reserved label block is not advertised by the IGP.

The **no** form of this command removes a reserved label block.

Parameters

name

Specifies the name of the reserved label block, up to 64 characters

Platforms

7705 SAR Gen 2

reserved-label-block

Syntax

reserved-label-block *name*

no reserved-label-block

Context

[Tree] (config>router>mpls>fwd-policies reserved-label-block)

Full Context

configure router mpls forwarding-policies reserved-label-block

Description

This command specifies the reserved label block to use for all MPLS forwarding policies. The named reserved label block must already have been configured under **config**>**router**>**mpls-labels**.

The **no** form of the command removes the assignment of the reserved label block.

Parameters

name

Specifies the name of the reserved label block, up to 64 characters.

Platforms

7705 SAR Gen 2

reserved-label-block

Syntax

reserved-label-block *name*

no reserved-label-block

Context

[Tree] (config>router>segment-routing>sr-policies reserved-label-block)

Full Context

configure router segment-routing sr-policies reserved-label-block

Description

This command associates a reserved label block with segment routing policies. The *name* must already exist. Reserved label blocks are configured under the **config**>**router**>**mpls-labels** hierarchy.

A locally-targeted segment routing policy (statically configured or BGP signaled) cannot be activated if its binding SID (BSID) is not an available label between the start-label and end-label of the referenced reserved label block.

The **no** form of this command removes any association of segment routing policies with a reserved label block.

Default

no reserved-label-block

Parameters

name

Specifies the name of a **reserved-label-block** that has already been configured, up to 64 characters.

Platforms

7705 SAR Gen 2

26.18 reset-policy-exclusive

reset-policy-exclusive

Syntax

reset-policy-exclusive

Context

[Tree] (admin reset-policy-exclusive)

Full Context

admin reset-policy-exclusive

Description

This command allows an authorized administrator to reset the exclusive policy editing lock. This will reset the lock flag and end the policy editing session in progress, discarding any policy edits.

Platforms

7705 SAR Gen 2

26.19 reset-query

reset-query

Syntax

[**no**] **reset-query**

Context

[Tree] (debug>router>rpki-session>packet reset-query)

Full Context

debug router rpki-session packet reset-query

Description

This command enables debugging for reset query RPKI packets.

The **no** form of this command disables debugging for reset query RPKI packets.

Platforms

7705 SAR Gen 2

26.20 resignal-on-igp-event

resignal-on-igp-event

Syntax

[**no**] **resignal-on-igp-event**

Context

[Tree] (config>router>mpls>sr-te-resignal resignal-on-igp-event)

Full Context

configure router mpls sr-te-resignal resignal-on-igp-event

Description

This command enables the ad hoc reoptimzation of all CSPF paths in the operational UP state of all SR-TE LSPs at the receipt of an IGP link event. The following link events are supported:

- link down
- link up
- IGP or TE metric change
- SRLG change
- admin group change

The ad hoc reoptimization follows the same behavior as in the timer-based resignal Make-Before-Break (MBB) feature. MPLS reevaluates all the paths in operational UP state of all SR-TE LSPs. The reevaluation consists of updating the total IGP or TE metric of the current path, checking the validity of the hops and labels, and computing a new CSPF path. MPLS programs the new path only if its total metric is different than the updated metric of the current path, or if one or more hops or labels of the current path are invalid. Otherwise, the current path is considered to be the most optimal and retained.

This feature does not require that the timer-based resignal (**configure router mpls sr-te-resignal resignal-timer**) command be enabled. If enabled, the resignal timer is aborted and an ad hoc reoptimization is performed.

The **no** form of this command disables ad hoc reoptimization of SR-TE LSPs.

Default

no resignal-on-igp-event

Platforms

7705 SAR Gen 2

resignal-on-igp-event

Syntax

[**no**] **resignal-on-igp-event**

Context

[Tree] (config>router>mpls resignal-on-igp-event)

Full Context

configure router mpls resignal-on-igp-event

Description

This command enables the ad hoc reoptimization of the active CSPF path of all RSVP-TE LSPs at the receipt of an IGP link event. The following link events are supported:

- link down
- link up
- IGP or TE metric change
- SRLG change
- admin group change

The ad hoc reoptimization follows the same behavior as in the timer-based resignal Make-Before-Break (MBB) feature. MPLS reevaluates the active paths of all RSVP-TE LSPs. The reevaluation consists of updating the total IGP or TE metric of the current path, checking the validity of the hops, and computing a new CSPF path. MPLS signals and programs the new path only if its total metric is different than the updated metric of the current path, or if one or more hops of the current path are invalid. Otherwise, the current path is considered to be the most optimal and retained.

This feature does not require that the timer-based resignal (**configure router mpls resignal-timer**) command be enabled. If enabled, the resignal timer is aborted and an ad hoc reoptimization is performed.

The **no** form of this command disables ad hoc reoptimization of the active RSVP-TE LSPs.

Default

no resignal-on-igp-event

Platforms

7705 SAR Gen 2

26.21 resignal-on-igp-overload

resignal-on-igp-overload

Syntax

[**no**] **resignal-on-igp-overload**

Context

[Tree] (config>router>mpls resignal-on-igp-overload)

Full Context

configure router mpls resignal-on-igp-overload

Description

This command enables the resignaling of all RSVP-TE LSPs at the receipt of the IS-IS overload bit in the TE-DB.

Once the re-optimization is triggered, the behavior is the same as the timer-based resignal or the **delay** option of the manual-based resignal. MPLS forces the expiry of the resignal timer and requests the TE-DB to compute a new CSPF for each RSVP-TE LSP active path.

This re-optimization effectively causes the immediate move of transit RSVP-TE LSP paths away from the IS-IS node in overload.

By default, MPLS re-optimizes, using the MBB procedure, the transit paths away from the node in an IS-IS overload state only at the time a manual or timer-based resignal is performed for the LSP paths. MPLS does not act immediately on the receipt of the IS-IS overload bit.

Note:
This command and the **retry-on-overload** command are mutually exclusive.

The **no** form of this command results in the MPLS not acting immediately to the request of the IS-IS overload bit.

Default

no resignal-on-overload

Platforms

7705 SAR Gen 2

resignal-on-igp-overload

Syntax

[**no**] **resignal-on-igp-overload**

Context

[Tree] (config>router>mpls>sr-te-resignal resignal-on-igp-overload)

Full Context

configure router mpls sr-te-resignal resignal-on-igp-overload

Description

This command enables the ad-hoc re-optimzation of the CSPF paths of all SR-TE LSPs when IS-IS receives an IS-IS overload bit advertisement from a remote router.

When this command is enabled on the router and an IGP overload bit is set in a Layer 1 or Layer 2 IS-IS LSP received from a remote router, MPLS performs an ad-hoc re-optimization of all the paths of all the SR-TE LSPs that have paths computed by the local CSPF. For each SR-TE LSP current path that transits the router in overload, the CSPF looks for a new path that avoids the router. For each SR-TE LSP current path that terminates on the router in overload, the CSPF checks if a better path exists. In both cases, if a new path is not found the system maintains the current path when operationally up.

The ad-hoc re-optimization triggers the timer-based re-optimization by forcing the resignal timer to expire. Therefore, the user must use the following command to configure the resignal timer for the SR-TE application.

```
configure router mpls sr-te-resignal resignal-timer
```

The **no** form of this command configures MPLS to not act immediately on an IS-IS overload bit from a remote router. MPLS will act on it at the next timer-based or manual re-optimization of the SR-TE LSPs.

Default

no resignal-on-igp-overload

Platforms

7705 SAR Gen 2

26.22 resignal-timer

resignal-timer

Syntax

resignal-timer *minutes*

no resignal-timer

Context

[Tree] (config>router>mpls resignal-timer)

Full Context

configure router mpls resignal-timer

Description

This command specifies the value for the LSP resignal timer. The resignal timer is the time, in minutes, the software waits before attempting to resignal the LSPs.

When the resignal timer expires, if the new computed path for an LSP has a better metric than the current recorded hop list, an attempt is made to resignal that LSP using the make-before-break mechanism. If the attempt to resignal an LSP fails, the LSP continues to use the existing path and a resignal will be attempted the next time the timer expires.

The **no** form of this command disables timer-based LSP resignaling.

Default

no resignal-timer

Parameters

minutes

Specifies the time the software waits before attempting to resignal the LSPs.

Values 30 to 10080

Platforms

7705 SAR Gen 2

resignal-timer

Syntax

resignal-timer *minutes*

no resignal-timer

Context

[Tree] (config>router>mpls>sr-te-resignal resignal-timer)

Full Context

configure router mpls sr-te-resignal resignal-timer

Description

This command specifies the value for the SR-TE LSP resignal timer when the path computation method is set to the local CSPF or the PCE.

The resignal timer is the time, in minutes, MPLS waits before attempting to re-optimize all paths of all SR-TE LSPs. The re-optimization is performed by the local CSPF or the PCE, depending on the value of the parameter **path-computation-method**.

When local CSPF is used and the resignal timer expires, MPLS provides the current path of the SR-TE LSP and TE-DB updates the total IGP or TE metric of the current path and checks the validity of the hops and labels. CSPF then computes a new path for each SR-TE LSP. MPLS programs the new path only if the total metric of the new computed path is different than the updated metric of the current path, or if one or

more hops or labels of the current path are invalid. Otherwise, the current path is considered to be one of the most optimal ECMP paths and is not updated in data path.

The **no** form of this command disables timer-based LSP resignaling.

Default

no resignal-timer

Parameters

minutes

Specifies the time, in minutes, the software waits before attempting to resignal the SR-TE LSPs.

Values 30 to 10080

Platforms

7705 SAR Gen 2

26.23 resolution

resolution

Syntax

resolution {**disabled** | **any** | **filter**}

Context

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel resolution)

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel resolution)

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel resolution)

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel resolution)

Full Context

configure service epipe bgp-evpn mpls auto-bind-tunnel resolution

configure service vprn bgp-ipvpn mpls auto-bind-tunnel resolution

configure service vpls bgp-evpn mpls auto-bind-tunnel resolution

configure service vprn bgp-evpn mpls auto-bind-tunnel resolution

Description

This command configures the resolution mode in the automatic binding of a BGP-EVPN or BGP-IPVPN MPLS service to tunnels to MP-BGP peers.

Default

resolution disabled

Parameters

any

Enables the binding to any supported tunnel type in a BGP-EVPN or BGP-IPVPN MPLS context following TTM preference.

disabled

Disables the automatic binding of a BGP-EVPN or BGP-IPVPN MPLS service to tunnels to MP-BGP peers.

filter

Enables the binding to the subset of tunnel types configured the **resolution-filter** context.

Platforms

7705 SAR Gen 2

resolution

Syntax

resolution {**any** | **disabled** | **filter**}

Context

[Tree] (config>router>static-route-entry>indirect>tunnel-next-hop resolution)

Full Context

configure router static-route-entry indirect tunnel-next-hop resolution

Description

This command determines the resolution mode for the associated static route to a tunnel next hop.

Default

resolution any

Parameters

any

Allows the associated static route to be resolved to any active entry in the TTM, following the TTM preference order.

disabled

Disables the resolution of the associated static route to any active entry in the TTM. As a result, the static route can only be resolved via IP RTM resolution of the static route's next hop.

filter

Allows the associated static route to be resolved to active tunnels in the TTM using the resolution-filter restrictions.

Platforms

7705 SAR Gen 2

resolution

Syntax

resolution {**any** | **filter** | **disabled**}

Context

[Tree] (config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family resolution)

Full Context

configure router bgp next-hop-resolution labeled-routes transport-tunnel family resolution

Description

This command configures the resolution mode in the resolution of BGP label routes using tunnels to BGP peers.

Parameters

any

Enables the binding to any supported tunnel type in the BGP label route context following TTM preference.

filter

Enables the binding to the subset of tunnel types configured under **resolution-filter**.

disabled

Disables the resolution of BGP label routes using tunnels to BGP peers.

Platforms

7705 SAR Gen 2

resolution

Syntax

resolution {**any** | **filter** | **disabled**}

Context

[Tree] (config>router>bgp>next-hop-res>shortcut-tunn>family resolution)

Full Context

configure router bgp next-hop-resolution shortcut-tunnel family resolution

Description

This command configures the resolution mode in the resolution of BGP prefixes using tunnels to BGP peers.

Parameters

any

Enables the binding to any supported tunnel type in BGP shortcut context following TTM preference.

filter

Enables the binding to the subset of tunnel types configured under **resolution-filter**.

disabled

Disables the resolution of BGP prefixes using tunnels to BGP peers.

Platforms

7705 SAR Gen 2

resolution

Syntax

resolution {**any** | **disabled** | **filter** | **match-family-ip**}

Context

[Tree] (config>router>isis>igp-shortcut>tunnel-next-hop>family resolution)

Full Context

configure router isis igp-shortcut tunnel-next-hop family resolution

Description

This command configures resolution mode in the resolution of the IP prefix or SR tunnel family using IGP shortcuts.

Parameters

any

Enables the binding to any supported tunnel type following TTM preference.

disabled

Disables the resolution of the IP prefix or SR tunnel family using IGP shortcuts.

filter

Enables the binding to the subset of tunnel types configured under **resolution-filter**.

match-family-ip

Enables the resolution of the SR tunnel family to match that of the corresponding IP prefix family.

Platforms

7705 SAR Gen 2

resolution

Syntax

resolution {**any** | **disabled** | **filter** | **match-family-ip**}

Context

[Tree] (config>router>ospf>igp-shortcut>tunnel-next-hop>family resolution)

Full Context

configure router ospf igp-shortcut tunnel-next-hop family resolution

Description

This command configures resolution mode in the resolution of the IP prefix or SR tunnel family using IGP shortcuts.

Parameters

any

Enables the binding to any supported tunnel type following TTM preference.

disabled

Disables the resolution of the IP prefix or SR tunnel family using IGP shortcuts.

filter

Enables the binding to the subset of tunnel types configured under **resolution-filter**.

match-family-ip

Enables the resolution of the SR tunnel family to match that of the corresponding IP prefix family.

Platforms

7705 SAR Gen 2

resolution

Syntax

resolution {**any** | **disabled** | **filter**}

Context

[Tree] (config>router>ospf3>igp-shortcut>tunnel-next-hop>family resolution)

Full Context

configure router ospf3 igp-shortcut tunnel-next-hop family resolution

Description

This command configures resolution mode in the resolution of the IPv6 prefix using IGP shortcuts.

Parameters

any

Enables the binding to any supported tunnel type following TTM preference.

disabled

Disables the resolution of the IPv6 prefix using IGP shortcuts.

filter

Enables the binding to the subset of tunnel types configured under **resolution-filter**.

Platforms

7705 SAR Gen 2

resolution

Syntax

resolution {**any** | **disabled** | **filter**}

Context

[Tree] (config>service>vprn>auto-bind-tunnel resolution)

Full Context

configure service vprn auto-bind-tunnel resolution

Description

This command configures the resolution method for tunnel selection.

Default

resolution any

Parameters

any

Allows the associated static route to be resolved to any active entry in the TTM, following the TTM preference order.

disabled

Disables the associated static route to be resolved to any active entry in the TTM. As a result, the static route can only be resolved via IP RTM resolution of the static route's nexthop.

filter

Allows the associated static route to be resolved to active tunnels in the TTM using the resolution-filter restrictions.

Platforms

7705 SAR Gen 2

resolution

Syntax

resolution

Context

[Tree] (config>service>vprn>auto-bind-tunnel resolution)

Full Context

configure service vprn auto-bind-tunnel resolution

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

26.24 resolution-filter

resolution-filter

Syntax

resolution-filter

Context

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel resolution-filter)

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel resolution-filter)

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel resolution-filter)

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel resolution-filter)

Full Context

configure service epipe bgp-evpn mpls auto-bind-tunnel resolution-filter

configure service vpls bgp-evpn mpls auto-bind-tunnel resolution-filter

configure service vprn bgp-evpn mpls auto-bind-tunnel resolution-filter

configure service vprn bgp-ipvpn mpls auto-bind-tunnel resolution-filter

Description

Commands in this context configure the subset of tunnel types that can be used in the resolution of BGP-EVPN or BGP-IPVPN routes within the automatic binding of BGP-EVPN or BGP-IPVPN MPLS service to tunnels to MP-BGP peers.

The following tunnel types are supported in a BGP-EVPN or BGP-IPVPN MPLS context: BGP, LDP, RIB-API, RSVP, SR-ISIS, SR-OSPF, SR-policy, SR-TE, UDP, and MPLS forwarding policy.

The user must set **resolution** to **filter** to activate the list of tunnel-types configured under **resolution-filter**.

Note:
UDP tunnels are created through import policies with action **create-udp-tunnel**.

Platforms

7705 SAR Gen 2

resolution-filter

Syntax

resolution-filter

Context

[Tree] (config>router>static-route-entry>indirect>tunnel-next-hop resolution-filter)

Full Context

configure router static-route-entry indirect tunnel-next-hop resolution-filter

Description

This command creates the context to configure the tunnel next-hop resolution options.

If one or more tunnel filter criteria are specified, the static route nexthop is resolved to an available tunnel from one of those LSP types. The tunnel type is selected based on the TTM preference.

Platforms

7705 SAR Gen 2

resolution-filter

Syntax

resolution-filter

Context

[Tree] (config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family resolution-filter)

Full Context

configure router bgp next-hop-resolution labeled-routes transport-tunnel family resolution-filter

Description

Commands in this context set resolution filter types.

Platforms

7705 SAR Gen 2

resolution-filter

Syntax

resolution-filter [**bgp**] [**ldp**] [**rsvp**] [**sr-isis**] [**sr-ospf**] [**sr-policy**] [**sr-te**]

Context

[Tree] (config>router>bgp>next-hop-res>shortcut-tunn>family resolution-filter)

Full Context

configure router bgp next-hop-resolution shortcut-tunnel family resolution-filter

Description

This command configures the subset of tunnel types that can be used to resolve BGP unlabeled routes.

Parameters

bgp

Selects the BGP label route tunnel type.

ldp

Selects the LDP tunnel type.

rsvp

Selects the RSVP-TE tunnel type.

sr-isis

Selects the SR tunnel type programmed by an IS-IS instance in TTM.

sr-ospf

Selects the SR tunnel type programmed by an OSPF instance in TTM.

sr-policy

Selects the SR tunnel type programmed by an SR policy instance in TTM.

sr-te

Selects the SR tunnel type programmed by a TE instance in TTM.

Platforms

7705 SAR Gen 2

resolution-filter

Syntax

resolution-filter

Context

[Tree] (config>router>isis>igp-shortcut>tunnel-next-hop>family resolution-filter)

Full Context

configure router isis igp-shortcut tunnel-next-hop family resolution-filter

Description

Commands in this context configure the subset of tunnel types which can be used in the resolution of the IP prefix or SR tunnel family using IGP shortcuts.

Parameters

rsvp

Selects the RSVP-TE tunnel type.

sr-te

Selects the SR-TE tunnel type.

Platforms

7705 SAR Gen 2

resolution-filter

Syntax

resolution-filter

Context

[Tree] (config>router>ospf3>igp-shortcut>tunnel-next-hop>family resolution-filter)

[Tree] (config>router>ospf>igp-shortcut>tunnel-next-hop>family resolution-filter)

Full Context

configure router ospf3 igp-shortcut tunnel-next-hop family resolution-filter

configure router ospf igp-shortcut tunnel-next-hop family resolution-filter

Description

Commands in this context configure the subset of tunnel types that can be used in the resolution of the IP prefix or SR tunnel family using IGP shortcuts.

Platforms

7705 SAR Gen 2

resolution-filter

Syntax

resolution-filter

Context

[Tree] (config>service>vprn>auto-bind-tunnel resolution-filter)

Full Context

configure service vprn auto-bind-tunnel resolution-filter

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

26.25 resolve

resolve

Syntax

resolve *minutes*

Context

[Tree] (config>service>vpls>proxy-arp>dynamic resolve)

[Tree] (config>service>vpls>proxy-nd>dynamic resolve)

Full Context

configure service vpls proxy-arp dynamic resolve

configure service vpls proxy-nd dynamic resolve

Description

This command configures the frequency at which a resolve message is sent. The resolve message is an ARP-request or NS message flooded to all the non-EVPN endpoints in the service irrespective of the current status of the **unknown-arp-request-flood-evpn** or **unknown-ns-flood-evpn** commands.

Default

resolve 5

Parameters

minutes

Specifies the frequency in minutes at which the **resolve** message is issued.

Values 1 to 60

Default 5

Platforms

7705 SAR Gen 2

26.26 resolve-static

resolve-static

Syntax

[**no**] **resolve-static**

Context

[Tree] (config>router>policy-options>policy-statement>entry>action resolve-static)

[Tree] (config>router>policy-options>policy-statement>default-action resolve-static)

Full Context

configure router policy-options policy-statement entry action resolve-static

configure router policy-options policy-statement default-action resolve-static

Description

This command has an affect only in BGP route-table-import policies and applies only to BGP IPv4 and IPv6 routes created by importing static routes with indirect next-hops. When such a route matches a policy entry with this action, the BGP next-hop is the resolved next-hop of the static route.

The **no** form of this command reverts to the default behavior, which copies the indirect next-hop of the static route into the BGP next-hop without resolving it further.

Default

no resolve-static

Platforms

7705 SAR Gen 2

26.27 resolve-v6-prefix-over-shortcut

resolve-v6-prefix-over-shortcut

Syntax

[**no**] **resolve-v6-prefix-over-shortcut**

Context

[Tree] (config>router>ldp>targ-session resolve-v6-prefix-over-shortcut)

Full Context

configure router ldp targeted-session resolve-v6-prefix-over-shortcut

Description

This command allows an IPv6 prefix FEC to be resolved over an IGP shortcut.

The **no** form of this command disables the resolution.

Platforms

7705 SAR Gen 2

26.28 responder-url

responder-url

Syntax

responder-url *url-string*

no responder-url

Context

[Tree] (config>system>security>pki>ca-profile>ocsp responder-url)

Full Context

configure system security pki ca-profile ocsp responder-url

Description

This command specifies HTTP URL of the OCSP responder for the CA, this URL will only be used if there is no OCSP responder defined in the AIA extension of the certificate to be verified.

Default

no responder-url

Parameters

url-string

Specifies the HTTP URL of the OCSP responder

Platforms

7705 SAR Gen 2

26.29 response-signing-cert

response-signing-cert

Syntax

response-signing-cert *filename*

no response-signing-cert

Context

[Tree] (config>system>security>pki>ca-profile>cmpv2 response-signing-cert)

Full Context

configure system security pki ca-profile cmpv2 response-signing-cert

Description

This command specifies a imported certificate that is used to verify the CMP response message if they are protected by signature. If this command is not configured, CA's certificate is used.

Default

no response-signing-cert

Parameters

filename

Specifies the filename of the imported certificate.

Platforms

7705 SAR Gen 2

26.30 restart-time

restart-time

Syntax

restart-time *seconds*

no restart-time

Context

[Tree] (config>service>vprn>bgp>group>graceful-restart restart-time)

[Tree] (config>service>vprn>bgp>graceful-restart restart-time)

[Tree] (config>service>vprn>bgp>group>neighbor>graceful-restart restart-time)

Full Context

configure service vprn bgp group graceful-restart restart-time

configure service vprn bgp graceful-restart restart-time

configure service vprn bgp group neighbor graceful-restart restart-time

Description

This command sets the value of the restart-time that is advertised in the router's graceful-restart capability. If this command is not configured, the default is 300.

Default

no restart-time

Parameters

seconds

Specifies the restart-time that is advertised in the router's graceful-restart capability.

Values 0 to 4095 seconds

Default 300

Platforms

7705 SAR Gen 2

restart-time

Syntax

restart-time *seconds*

no restart-time

Context

[Tree] (config>router>bgp>group>graceful-restart restart-time)

[Tree] (config>router>bgp>group>neighbor>graceful-restart restart-time)

[Tree] (config>router>bgp>graceful-restart restart-time)

Full Context

configure router bgp group graceful-restart restart-time

configure router bgp group neighbor graceful-restart restart-time

configure router bgp graceful-restart restart-time

Description

This command sets the value of the restart-time that is advertised in the router's graceful-restart capability. If this command is not configured, the default is 300.

Default

no restart time

Parameters

seconds

Specifies the restart-time that is advertised in the router's graceful-restart capability.

Values 0 to 4095 seconds

Default config>router>bgp>graceful-restart: 120 seconds

config>router>bgp>group>graceful-restart: 300 seconds

config>router>bgp>group>neighbor>graceful-restart: 300 seconds

Platforms

7705 SAR Gen 2

26.31 restrict-non-configured-ip-address

restrict-non-configured-ip-address

Syntax

restrict-non-configured-ip-address [**sponge-mac** *mac-address*]

no restrict-non-configured-ip-address

Context

[Tree] (config>service>vpls>proxy-arp restrict-non-configured-ip-address)

[Tree] (config>service>vpls>proxy-nd restrict-non-configured-ip-address)

Full Context

configure service vpls proxy-arp restrict-non-configured-ip-address

configure service vpls proxy-nd restrict-non-configured-ip-address

Description

This command configures whether all the configured dynamic IP address entries are considered the only authorized entries in the proxy ARP or ND table. ARP or ND packets coming from a unauthorized sender IP addresses are dropped. Therefore, unauthorized IP addresses are not learned in the proxy ARP or ND table, and ARP requests or neighbor solicitations (NS) coming from a unauthorized sender IP addresses are not replied to, unless the **sponge-mac** option is configured.

The **no** form of this command does not drop ARP or ND packets coming from a unauthorized sender IP addresses.

Parameters

sponge-mac

Keyword to specify that ARP requests or NSs from an unauthorized IP address are not learned in the proxy ARP or ND table and ARP requests or NSs from an unauthorized IP address are replied with the configured sponge MAC address. Any IP address that is not configured as proxy ARP, ND dynamic ARP, or neighbor IP address is considered unauthorized and dropped.

mac-address

Specifies the MAC address.

The configured sponge MAC address is not installed in the FDB or advertised in EVPN. If needed, the sponge MAC address can be configured as a static MAC in the same service in the node or a remote node.

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

Platforms

7705 SAR Gen 2

26.32 restrict-protected-src

restrict-protected-src

Syntax

restrict-protected-src discard-frame

restrict-protected-src [**alarm-only**]

no restrict-protected-src

Context

[Tree] (config>service>pw-template>split-horizon-group restrict-protected-src)

[Tree] (config>service>pw-template restrict-protected-src)

[Tree] (config>service>vpls>mesh-sdp restrict-protected-src)

[Tree] (config>service>vpls>spoke-sdp restrict-protected-src)

[Tree] (config>service>vpls>split-horizon-group restrict-protected-src)

[Tree] (config>service>vpls>sap restrict-protected-src)

[Tree] (config>service>vpls>endpoint restrict-protected-src)

Full Context

configure service pw-template split-horizon-group restrict-protected-src

configure service pw-template restrict-protected-src

configure service vpls mesh-sdp restrict-protected-src

configure service vpls spoke-sdp restrict-protected-src

configure service vpls split-horizon-group restrict-protected-src

configure service vpls sap restrict-protected-src

configure service vpls endpoint restrict-protected-src

Description

This command indicates how the agent will handle relearn requests for protected MAC addresses, either manually added using the mac-protect command or automatically added using the **auto-learn-mac-protect** command. While enabled all packets entering the configured SAP, spoke SDP, mesh SDP, or any SAP that is part of the configured split horizon group (SHG) is verified not to contain a protected source MAC address. If the packet is found to contain such an address, the action taken depends on the parameter specified on the **restrict-protected-src** command, namely:

- No parameter — The packet is discarded, an alarm is generated and the SAP, spoke SDP or mesh SDP is set operationally down. The SAP, spoke SDP or mesh SDP must be shut down and enabled (**no shutdown**) for this state to be cleared.
- **alarm-only** — The packet is forwarded, an alarm is generated but the source MAC is not learned on the SAP, spoke SDP or mesh SDP.

- **discard-frame** — The packet is discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per MAC address per FP per 10 minutes in a given VPLS service. This parameter is only applicable to automatically protected MAC addresses.

When the **restrict-protected-src** is enabled on an SHG the action only applies to the associated SAPs (no action is taken by default for spoke SDPs in the SHG) and is displayed in the SAP show output as the oper state unless it is overridden by the configuration of **restrict-protected-src** on the SAP itself. To enable this function for spoke SDPs within a SHG, the **restrict-protected-src** must be enabled explicitly under the spoke SDP. If required, **restrict-protected-src** can also be enabled explicitly under specific SAPs within the SHG.

When this command is applied or removed, with either the **alarm-only** or **discard-frame** parameters, the MAC addresses are cleared from the related object.

The use of **restrict-protected-src discard-frame** is mutually exclusive with the configuration of manually protected MAC addresses within a given VPLS.

The **no** form of the command reverts to the default.

Default

no restrict-protected-src

Parameters

alarm-only

Specifies that the packet is forwarded, an alarm is generated but the source MAC is not learned on the SAP, spoke SDP, or mesh SDP.

Default no alarm-only

discard-frame

Specifies that the packet is discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per MAC address per FP per 10 minutes within a given VPLS service.

Default no discard-frame

Platforms

7705 SAR Gen 2

26.33 restrict-unprotected-dst

restrict-unprotected-dst

Syntax

restrict-unprotected-dst

no restrict-unprotected-dst

Context

[Tree] (config>service>vpls>sap restrict-unprotected-dst)

[Tree] (config>service>vpls>split-horizon-group restrict-unprotected-dst)

[Tree] (config>service>pw-template>split-horizon-group restrict-unprotected-dst)

Full Context

configure service vpls sap restrict-unprotected-dst

configure service vpls split-horizon-group restrict-unprotected-dst

configure service pw-template split-horizon-group restrict-unprotected-dst

Description

This command indicates how the system will forward packets destined for an unprotected MAC address, either manually added using the **mac-protect** command or automatically added using the **auto-learn-mac-protect** command. While enabled all packets entering the configured SAP or SAPs within a split horizon group (but not spoke or mesh-SDPs) will be verified to contain a protected destination MAC address. If the packet is found to contain a non-protected destination MAC, it will be discarded. Detecting a non-protected destination MAC on the SAP will not cause the SAP to be placed in the operationally down state. No alarms are generated.

If the destination MAC address is unknown, even if the packet is entering a restricted SAP, with **restrict-unprotected-dst** enabled, it will be flooded.

Default

no restrict-unprotected-dst

Platforms

7705 SAR Gen 2

26.34 restricted-to-home

restricted-to-home

Syntax

[**no**] **restricted-to-home**

Context

[Tree] (config>system>security>user restricted-to-home)

[Tree] (config>system>security>user-template restricted-to-home)

Full Context

configure system security user restricted-to-home

configure system security user-template restricted-to-home

Description

This command denies the user from accessing files outside of their home directory. Files can be accessed locally by CLI file commands and output modifiers such as **>** (file redirect), or remotely via FTP and SCP.

When enabled, the system denies all configuration save operations (such as **admin save**) via any management interface (such as CLI and NETCONF) unless **save-when-restricted** is enabled.

When **restricted-to-home** is configured, file access is denied unless the **home-directory** is configured and the directory is created by an administrator.

The **no** form of this command permits the user to access all files on the system.

Default

restricted-to-home

Platforms

7705 SAR Gen 2

26.35 results

results

Syntax

results *file-url*

no results

Context

[Tree] (config>system>script-control>script-policy results)

Full Context

configure system script-control script-policy results

Description

This command is used to specify the location where the system writes the output of an event script's execution.

The **no** form of the command removes the file location from the configuration. Scripts will not execute if there is no result location defined.

Default

no results

Parameters

file-url

Specifies the location to send CLI output from script runs. The *file-url* is a location, directory, and filename prefix to which a data and timestamp suffix is added when the results files are created during a script run, as follows:

*file-url*_YYYYMMDD-hhmmss.uuuuuu.out

where:

YYYYMMDD — date

hhmmss — hours, minutes, and seconds

uuuuuu — microseconds (padded to 6 characters with leading zeros)

Values *local-url* | *remote-url*

local-url — [*cflash-id/*] [*file-path*] 167 chars max, including *cflash-idfile-path* 166 chars max

remote url — [{ftp:// | tftp://}*login*:*password@remote-location/*][*file-path*] 255 characters max directory length 99 characters max each

remote-location — [*hostname* | *ipv4-address* | *ipv6-address*]

ipv4-address — *a.b.c.d*

ipv6-address — x:x:x:x:x:x:x:x[-*interface*] x:x:x:x:x:x:d.d.d.d[-*interface*] x — [0 to FFFF]H d — [0 to 255]D interface — 32 characters max, for link local addresses

cflash-id — cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

26.36 resv

resv

Syntax

resv [**detail**]

no resv

Context

[Tree] (debug>router>rsvp>event resv)

Full Context

debug router rsvp event resv

Description

This command debugs RSVP reservation events.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about RSVP reservation events.

Platforms

7705 SAR Gen 2

resv

Syntax

resv [**detail**]

no resv

Context

[Tree] (debug>router>rsvp>packet resv)

Full Context

debug router rsvp packet resv

Description

This command enables debugging for RSVP resv packets.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about RSVP Resv events.

Platforms

7705 SAR Gen 2

26.37 resverr

resverr

Syntax

resverr [**detail**]

no resverr

Context

[Tree] (debug>router>rsvp>packet resverr)

Full Context

debug router rsvp packet resverr

Description

This command debugs ResvErr packets.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about ResvErr packets.

Platforms

7705 SAR Gen 2

26.38 resvtear

resvtear

Syntax

resvtear [**detail**]

no resvtear

Context

[Tree] (debug>router>rsvp>packet resvtear)

Full Context

debug router rsvp packet resvtear

Description

This command debugs ResvTear packets.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about ResvTear packets.

Platforms

7705 SAR Gen 2

26.39 retransmit-interval

retransmit-interval

Syntax

retransmit-interval *seconds*

no retransmit-interval

Context

[Tree] (config>service>vprn>isis>if retransmit-interval)

Full Context

configure service vprn isis interface retransmit-interval

Description

This command configures the minimum time between LSP PDU retransmissions on a point-to-point interface.

The **no** form of this command reverts to the default value.

Default

retransmit-interval 5

Parameters

seconds

Specifies the interval in seconds that IS-IS LSPs can be sent on the interface

1 to 65535.

Platforms

7705 SAR Gen 2

retransmit-interval

Syntax

retransmit-interval *seconds*

no retransmit-interval

Context

[Tree] (config>service>vprn>ospf3>area>if retransmit-interval)

[Tree] (config>service>vprn>ospf>area>if retransmit-interval)

[Tree] (config>service>vprn>ospf>area>sham-link retransmit-interval)

[Tree] (config>service>vprn>ospf>area>virtual-link retransmit-interval)

[Tree] (config>service>vprn>ospf3>area>virtual-link retransmit-interval)

Full Context

configure service vprn ospf3 area interface retransmit-interval

configure service vprn ospf area interface retransmit-interval

configure service vprn ospf area sham-link retransmit-interval

configure service vprn ospf area virtual-link retransmit-interval

configure service vprn ospf3 area virtual-link retransmit-interval

Description

This command specifies the length of time, in seconds, that OSPF will wait before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor.

The value should be longer than the expected round trip delay between any two routers on the attached network. Once the retransmit interval expires and no acknowledgment is received, the LSA is retransmitted.

The **no** form of this command reverts to the default interval.

Default

retransmit-interval 5

Parameters

seconds

The retransmit interval in seconds expressed as a decimal integer.

Values 1 to 3600

Platforms

7705 SAR Gen 2

retransmit-interval

Syntax

retransmit-interval *seconds*

no retransmit-interval

Context

[Tree] (config>router>isis>interface retransmit-interval)

Full Context

configure router isis interface retransmit-interval

Description

This command configures the minimum time between LSP PDU retransmissions on a point-to-point interface.

The **no** form of this command reverts to the default value.

Default

retransmit-interval 5

Parameters

seconds

Specifies the interval, in seconds, that IS-IS LSPs can be sent on the interface.

Values 1 to 65535

Platforms

7705 SAR Gen 2

retransmit-interval

Syntax

retransmit-interval *seconds*

no retransmit-interval

Context

[Tree] (config>router>ospf3>area>virtual-link retransmit-interval)

[Tree] (config>router>ospf>area>interface retransmit-interval)

[Tree] (config>router>ospf>area>virtual-link retransmit-interval)

[Tree] (config>router>ospf3>area>interface retransmit-interval)

Full Context

configure router ospf3 area virtual-link retransmit-interval

configure router ospf area interface retransmit-interval

configure router ospf area virtual-link retransmit-interval

configure router ospf3 area interface retransmit-interval

Description

This command specifies the length of time, in seconds, that OSPF will wait before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor.

The value should be longer than the expected round trip delay between any two routers on the attached network. After the retransmit-interval expires and no acknowledgment has been received, the LSA will be retransmitted.

The **no** form of this command reverts to the default interval.

Default

retransmit-interval 5

Parameters

seconds

Specifies the retransmit interval in seconds expressed as a decimal integer.

Values 1 to 1800

Platforms

7705 SAR Gen 2

26.40 retransmit-time

retransmit-time

Syntax

retransmit-time *milli-seconds*

no retransmit-time

Context

[Tree] (config>service>vprn>router-advert>if retransmit-time)

[Tree] (config>router>router-advert>if retransmit-time)

Full Context

configure service vprn router-advertisement interface retransmit-time

configure router router-advertisement interface retransmit-time

Description

This command configures the value to be placed in the retransmit timer field in router advertisements sent from this interface.

The **no** form of this command reverts to the default.

Default

retransmit-time 0

Parameters

milli-seconds

Specifies the retransmit time, in milli-seconds, for advertisement from this group-interface.

Values 0 to 1800000

Platforms

7705 SAR Gen 2

26.41 retries

retries

Syntax

retries *count*

no retries

Context

[Tree] (config>system>grpc>tcp-keepalive retries)

Full Context

configure system grpc tcp-keepalive retries

Description

This command configures the number of TCP keepalive probes sent by the router that must be unacknowledged before the connection is closed.

The **no** form of this command reverts to the default value.

Default

retries 4

Parameters

count

Specifies the number of missed keep-alives before the TCP connection is declared down.

Values 3 to 100

Default 4

Platforms

7705 SAR Gen 2

retries

Syntax

retries *count*

no retries

Context

[Tree] (config>system>grpc-tunnel>destination-group>tcp-keepalive retries)

[Tree] (config>system>telemetry>destination-group>tcp-keepalive retries)

Full Context

configure system grpc-tunnel destination-group tcp-keepalive retries

configure system telemetry destination-group tcp-keepalive retries

Description

This command configures the number of missed TCP keepalive probes before the TCP connection is closed and attempts are made to reach other destinations within the same destination group.

The **no** form of this command reverts to the default value.

Default

retries 4

Parameters

count

Specifies the number of missed keep-alives before the TCP connection is declared down.

Values 3 to 100

Default 4

Platforms

7705 SAR Gen 2

26.42 retry

retry

Syntax

retry *count*

no retry

Context

[Tree] (config>aaa>radius-srv-plcy>servers retry)

Full Context

configure aaa radius-server-policy servers retry

Description

This command configures the number of times the router attempts to contact the RADIUS server, if not successful the first time.

The **no** form of this command reverts to the default.

Default

retry 3

Parameters

count

Specifies the number of times a signaling request message is transmitted towards the same peer.

Values 1 to 256

Platforms

7705 SAR Gen 2

retry

Syntax

retry *minutes*

no retry

Context

[Tree] (config>service>vpls>bgp-evpn>mac-duplication retry)

Full Context

configure service vpls bgp-evpn mac-duplication retry

Description

Specifies the timer after which the MAC in hold-down state is automatically flushed and the mac-duplication process starts again. This value is expected to be equal to two times or more than that of window.

If **no** retry is configured, this implies that, when mac-duplication is detected, MAC updates for that MAC will be held down till the user intervenes or a network event (that flushes the MAC) occurs.

Default

retry 9

Parameters

minutes

Specifies the BGP EVPN MAC duplication retry in minutes.

Values 2 to 60

Platforms

7705 SAR Gen 2

retry

Syntax

retry *count*

no retry

Context

[Tree] (config>service>vprn>aaa>rmt-srv>radius retry)

[Tree] (config>system>security>radius retry)

Full Context

configure service vprn aaa remote-servers radius retry

configure system security radius retry

Description

This command configures the number of times the router attempts to contact the RADIUS server for authentication if there are problems communicating with the server.

The **no** form of this command reverts to the default value.

Default

retry 3

Parameters

count

Specifies the retry count.

Values 1 to 10

Platforms

7705 SAR Gen 2

retry

Syntax

retry *count*

no retry

Context

[Tree] (config>system>file-trans-prof retry)

Full Context

configure system file-transmission-profile retry

Description

This command specifies the number of retries on transport protocol level.

When the virtual router does not receive any data from a server (e.g., FTP or HTTP server) after the configured **timeout** *seconds*, the router may repeat the request to the server. The number of retries specifies the maximum number of repeated requests.

The **no** form of this command disables the retry.

Default

no retry

Parameters

count

Specifies the number of retries.

Values 1 to 256

Platforms

7705 SAR Gen 2

retry

Syntax

retry *count*

no retry

Context

[Tree] (config>system>security>ldap retry)

Full Context

configure system security ldap retry

Description

This command configures the number of retries for the SR OS in its attempt to reach the current LDAP server before attempting the next server.

The **no** form of this command reverts to the default value.

Default

retry 3

Parameters

count

Specifies the number of retransmissions.

Values 1 to 10

Default 3

Platforms

7705 SAR Gen 2

26.43 retry-count

retry-count

Syntax

retry-count *retry-count*

no retry-count

Context

[Tree] (config>service>epipe>spoke-sdp-fec retry-count)

Full Context

configure service epipe spoke-sdp-fec retry-count

Description

This optional command specifies the number of attempts software should make to reestablish the spoke SDP after it has failed. After each successful attempt, the counter is reset to zero.

When the specified number is reached, no more attempts are made and the spoke-sdp is put into the shutdown state.

Use the no shutdown command to bring up the path after the retry limit is exceeded.

The **no** form of this command reverts the parameter to the default value.

Default

retry-count 30

Parameters

retry-count

The maximum number of retries before putting the spoke-sdp into the shutdown state.

Values 10 to 10000

Platforms

7705 SAR Gen 2

retry-count

Syntax

retry-count [*count*]

no retry-count

Context

[Tree] (config>service>pw-routing retry-count)

Full Context

configure service pw-routing retry-count

Description

This optional command specifies the number of attempts software should make to re-establish the spoke SDP after it has failed. After each successful attempt, the counter is reset to zero.

When the specified number is reached, no more attempts are made and the spoke SDP is put into the shutdown state.

Use the **no shutdown** command to bring up the path after the retry limit is exceeded.

The **no** form of this command reverts the parameter to the default value.

Default

no retry-count

Parameters

count

Specifies the maximum number of retries before putting the spoke SDP into the shutdown state.

Values 10 to 10000

Platforms

7705 SAR Gen 2

26.44 retry-interval

retry-interval

Syntax

retry-interval *seconds*

no retry-interval

Context

[Tree] (config>system>security>pki>ca-prof>auto-crl-update retry-interval)

Full Context

configure system security pki ca-profile auto-crl-update retry-interval

Description

This command specifies the interval, in seconds, that the system waits before retrying the configured **url-entry** list when **schedule-type** is **next-update-based** and none of the URLs return a qualified CRL.

The **no** form of this command causes the system to retry immediately without waiting.

Default

retry-interval 3600

Parameters

seconds

Specifies an interval, in seconds, before retrying to update the CRL.

Values 1 to 31622400

Platforms

7705 SAR Gen 2

retry-interval

Syntax

retry-interval *seconds*

Context

[Tree] (config>system>security>pki>cert-upd-prof retry-interval)

Full Context

configure system security pki certificate-update-profile retry-interval

Description

This command configures the retry interval after the update fails.

Default

retry-interval 3600

Parameters

seconds

Specifies a retry interval, in seconds, after a failed update.

Values 60 to 36000

Platforms

7705 SAR Gen 2

26.45 retry-limit

retry-limit

Syntax

retry-limit *number*

no retry-limit

Context

[Tree] (config>router>mpls>lsp-template retry-limit)

[Tree] (config>router>mpls>lsp retry-limit)

Full Context

configure router mpls lsp-template retry-limit

configure router mpls lsp retry-limit

Description

This optional command specifies the number of attempts software should make to re-establish the LSP after it has failed LSP. After each successful attempt, the counter is reset to zero.

When the specified number is reached, no more attempts are made and the LSP path is put into the **shutdown** state.

Use the config router **mpls lsp** *lsp-name* **no shutdown** command to bring up the path after the retry-limit is exceeded.

For P2MP LSP that are created based on the LSP template, all S2Ls must attempt to retry-limit before the client application is informed of failure.

The **no** form of this command reverts to the default value.

Default

retry-limit 0 (no limit, retries forever)

Parameters

number

Specifies the number of times software will attempt to re-establish the LSP after it has failed. Allowed values are integers in the range of 0 to 10000.

Values 0 to 10000

Platforms

7705 SAR Gen 2

26.46 retry-on-igp-overload

retry-on-igp-overload

Syntax

[**no**] **retry-on-igp-overload**

Context

[Tree] (config>router>mpls retry-on-igp-overload)

Full Context

configure router mpls retry-on-igp-overload

Description

This command allows for the global configuration of the handling in the ingress LER of the LSP paths which transit an LSR that advertised the IS-IS overload bit.

By default, MPLS re-optimizes using make-before-break (MBB) the transit paths away from the node in an IS-IS overload state only at the time a manual or timer-based re-signal is performed for the LSP paths. MPLS will not act immediately on the receipt of the IS-IS overload bit.

When this command is enabled, MPLS in the ingress LER immediately tears down and re-signals all LSP paths away from a transit LSR node which advertised the IS-IS overload bit.

LSP paths that terminate on the node that advertised the IS-IS overload bit are not acted on whether this command is enabled or disabled.

The **no** form of this command returns to the default behavior.

Platforms

7705 SAR Gen 2

26.47 retry-timeout

retry-timeout

Syntax

retry-timeout *timeout*

no retry-timeout

Context

[Tree] (config>service>template>vpls-template>mac-move retry-timeout)

[Tree] (config>service>vpls>mac-move retry-timeout)

Full Context

configure service template vpls-template mac-move retry-timeout

configure service vpls mac-move retry-timeout

Description

This indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled.

It is recommended that the retry-timeout value is larger or equal to 5s * cumulative factor of the highest priority port so that the sequential order of port blocking will not be disturbed by re-initializing lower priority ports.

A zero value indicates that the SAP will not be automatically re-enabled after being disabled. If, after the SAP is re-enabled it is disabled again, the retry timeout is increased with the provisioned retry timeout in order to avoid thrashing. For example, when retry-timeout is set to 15, it increments (15,30,45,60...).

The **no** form of this command reverts to the default value.

Default

retry-timeout 10 (when mac-move is enabled)

Parameters

timeout

Specifies the time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled.

Values 0 to 120

Platforms

7705 SAR Gen 2

26.48 retry-timer

retry-timer

Syntax

retry-timer *retry-timer*

no retry-timer

Context

[Tree] (config>service>epipe>spoke-sdp-fec retry-timer)

Full Context

configure service epipe spoke-sdp-fec retry-timer

Description

This command specifies a retry-timer for the spoke SDP. This is a configurable exponential back-off timer that determines the interval between retries to reestablish a spoke SDP if it fails and a label withdraw message is received with the status code "AII unreachable”.

The **no** form of this command reverts the timer to its default value.

Default

retry-timer 30

Parameters

retry-timer

The initial retry-timer value in seconds.

Values 10 to 480

Platforms

7705 SAR Gen 2

retry-timer

Syntax

retry-timer *seconds*

no retry-timer

Context

[Tree] (config>router>mpls>lsp retry-timer)

[Tree] (config>router>mpls>lsp-template retry-timer)

Full Context

configure router mpls lsp retry-timer

configure router mpls lsp-template retry-timer

Description

This command configures the time (in s), for LSP re-establishment attempts after it has failed. The retry time is jittered to +/- 25% of its nominal value.

For P2MP LSP created based on LSP template, all S2Ls must attempt to retry-limit before client application is informed of failure.

The **no** form of this command reverts to the default value.

Default

retry-timer 30

Parameters

seconds

Specifies the amount of time (in s), between attempts to re-establish the LSP after it has failed. Allowed values are integers in the range of 1 to 600.

Values 1 to 600

Platforms

7705 SAR Gen 2

retry-timer

Syntax

retry-timer *secs*

no retry-timer

Context

[Tree] (config>service>pw-routing retry-timer)

Full Context

configure service pw-routing retry-timer

Description

This command configures a retry-timer for the spoke-SDP. This is a configurable exponential back-off timer that determines the interval between retries to re-establish a spoke-SDP if it fails and a label withdraw message is received with the status code "AII unreachable”.

The **no** form of this command reverts the timer to its default value.

Default

no retry-timer

Parameters

secs

Specifies initial retry-timer value in seconds.

Values 10 to 480

Platforms

7705 SAR Gen 2

26.49 return-path-label

return-path-label

Syntax

return-path-label *label-value*

no return-path-label

Context

[Tree] (config>router>segment-routing>main-plcy return-path-label)

Full Context

configure router segment-routing maintenance-policy return-path-label

Description

This command configures the Seamless Bidirectional Forwarding Detection (S-BFD) session to echo mode and adds an additional MPLS label, referring to an MPLS-labeled reply path for the S-BFD packet, to the bottom of the label stack for the S-BFD packet.

The command applies to the initiator of the S-BFD sessions. The return-path label may be a binding SID for an SR policy or other MPLS path configured on the reflector router. Instead of being routed through the IGP path, the S-BFD packet returns to the initiator through this MPLS return path.

The **no** form of this command disables the controlled return-path label and echo mode for S-BFD. S-BFD returns to asynchronous mode and the initiator node does not push a return-path label. Any S-BFD packets for this LSP or path that the reflector receives are sent back using a routed return path.

Default

no return-path-label

Parameters

label-value

Specifies the label value.

Values 32 to 1048575

Platforms

7705 SAR Gen 2

26.50 reuse

reuse

Syntax

reuse *integer*

no reuse

Context

[Tree] (config>router>policy-options>damping reuse)

Full Context

configure router policy-options damping reuse

Description

This command configures the reuse parameter for the route damping profile.

When the Figure of Merit (FoM) value falls below the **reuse** threshold, the route is once again considered valid and can be reused or included in route advertisements.

The **no** form of this command removes the reuse parameter from the damping profile.

Default

no reuse

Parameters

integer

Specifies the reuse value expressed as a decimal integer.

Values 1 to 20000

Platforms

7705 SAR Gen 2

26.51 reverse-route

reverse-route

Syntax

reverse-route

Context

[Tree] (config>ipsec>tnl-temp reverse-route)

Full Context

configure ipsec tunnel-template reverse-route

Description

Commands in this context configure the dynamic LAN-to-LAN (DL2L) tunnel reverse-route options for the tunnel template.

Platforms

7705 SAR Gen 2

26.52 revert

revert

Syntax

revert {**latest-rb**| *checkpoint-id* | **rescue**} [**now**]

Context

[Tree] (admin>rollback revert)

Full Context

admin rollback revert

Description

This command initiates a configuration rollback revert operation that will return the configuration state of the node to a previously saved checkpoint. The rollback revert minimizes impacts to running services. There are no impacts in areas of configuration that did not change since the checkpoint. Configuration parameters that changed (or items on which changed configuration have dependencies) are first removed (revert to default) and the previous values are then restored (can be briefly service impacting in changed areas).

Parameters

latest-rb

Specifies the most recently created rollback checkpoint (corresponds to the file-url.rb rollback checkpoint file).

checkpoint-id

Specifies the configuration to return to (which rollback checkpoint file to use). Checkpoint-id of 1 corresponds to the file-url.rb.1 rollback checkpoint file. The higher the id, the older the checkpoint. Max is the highest rollback checkpoint supported or configured.

Values 1 to 9

rescue

Specifies to revert to the rescue checkpoint.

now

Forces a rollback revert without any interactive confirmations (assumes 'y' for any confirmations that would have occurred).

Platforms

7705 SAR Gen 2

26.53 revert-members

revert-members

Syntax

revert-members [*1..8*]

no revert-members

Context

[Tree] (config>service>vprn>isis>link-group>level revert-members)

Full Context

configure service vprn isis link-group level revert-members

Description

This command sets the threshold for the minimum number of operational links to return the associated link group to its normal operating state and remove the associated offsets to the IS-IS metrics. If the number of operational links is equal to or greater than the configured **revert-members** threshold, the configured offsets are removed.

The **no** form of this command reverts the threshold back to the default, which is equal to the **oper-members** threshold value.

Default

no revert-members *oper-members*

Parameters

1..8

Specifies the number of revert members.

Values 1 to 8

Platforms

7705 SAR Gen 2

revert-members

Syntax

revert-members [*1..8*]

no revert-members

Context

[Tree] (config>router>isis>link-group>level revert-members)

Full Context

configure router isis link-group level revert-members

Description

This command sets the threshold for the minimum number of operational links to return the associated link group to its normal operating state and remove the associated offsets to the IS-IS metrics. If the number of operational links is equal to or greater than the configured revert-member threshold then the configured offsets are removed.

The **no** form of this command reverts the threshold back to the default which is equal to the oper-member threshold value.

Default

no revert-members oper-members

Parameters

1..8

Specifies the threshold for revertive members.

Values 1 to 8

Platforms

7705 SAR Gen 2

26.54 revert-time

revert-time

Syntax

revert-time [*revert-time* | **infinite**]

no revert-time

Context

[Tree] (config>service>epipe>endpoint revert-time)

Full Context

configure service epipe endpoint revert-time

Description

This command configures the time to wait before reverting back to the primary spoke SDP defined on this service endpoint, after having failed over to a backup spoke SDP.

Parameters

revert-time

Specifies the time, in seconds, to wait before reverting to the primary SDP.

Values 0 to 600

Default 0

infinite

Causes the endpoint to be non-revertive.

Platforms

7705 SAR Gen 2

revert-time

Syntax

revert-time *revert-time* | **infinite**

no revert-time

Context

[Tree] (config>service>vpls>endpoint revert-time)

Full Context

configure service vpls endpoint revert-time

Description

This command configures the time to wait before reverting to primary spoke-SDP.

In a regular endpoint the revert-time setting affects just the pseudowire defined as primary (precedence 0). For a failure of the primary pseudowire followed by restoration the revert-timer is started. After it expires the primary pseudowire takes the active role in the endpoint. This behavior does not apply for the case when both pseudowires are defined as secondary. For example, if the active secondary pseudowire fails and is restored it will stay in standby until a configuration change or a force command occurs.

Parameters

revert-time

Specifies the time to wait, in seconds, before reverting back to the primary spoke-SDP defined on this service endpoint, after having failed over to a backup spoke-SDP

Values 0 to 600

infinite

Specifying this keyword makes endpoint non-revertive

Platforms

7705 SAR Gen 2

revert-time

Syntax

revert-time {*revert-time* | **infinite**}

no revert-time

Context

[Tree] (config>mirror>mirror-dest>endpoint revert-time)

Full Context

configure mirror mirror-dest endpoint revert-time

Description

This command configures the time to wait before reverting to the primary spoke SDP. This command has an effect only when used in conjunction with an endpoint which contains a SDP of type 'primary'. It is ignored and has no effect in all other cases. The revert-timer is the delay in seconds the system waits before it switches the path of the mirror service from an active secondary SDP in the endpoint into the endpoint primary SDP after the latter comes back up.

The **no** form of this command resets the timer to the default value of 0. This means that the mirror-service path is switched back to the endpoint primary sdp immediately after it comes back up.

Parameters

revert-time

Specifies a delay, in seconds, the system waits before it switches the path of the mirror service from an active secondary SDP in the endpoint into the endpoint primary SDP after the latter comes back up.

Values 0 to 600

infinite

Forces the mirror or LI service path to never revert to the primary SDP as long as the currently active secondary SDP is UP.

Platforms

7705 SAR Gen 2

revert-time

Syntax

revert-time {*revert-time* | **infinite**}

no revert-time

Context

[Tree] (config>service>sdp>mixed-lsp-mode revert-time)

Full Context

configure service sdp mixed-lsp-mode revert-time

Description

This command configures the delay period the SDP must wait before it reverts to a higher priority LSP type when one becomes available.

The **no** form of the command resets the timer to the default value of 0. This means the SDP reverts immediately to a higher priority LSP type when one becomes available.

Default

no revert-time

Parameters

revert-time

Specifies the delay period, in seconds, that the SDP must wait before it reverts to a higher priority LSP type when one becomes available. A value of zero means the SDP reverts immediately to a higher priority LSP type when one becomes available.

Values 0 to 600

infinite

This keyword forces the SDP to never revert to another higher priority LSP type unless the currently active LSP type is down.

Platforms

7705 SAR Gen 2

26.55 revert-timer

revert-timer

Syntax

revert-timer *timer-value*

no revert-timer

Context

[Tree] (config>router>mpls>lsp revert-timer)

Full Context

configure router mpls lsp revert-timer

Description

This command configures a revert timer on an LSP. The timer starts when the LSP primary path recovers from a failure. The LSP reverts from a secondary path to the primary path when the timer expires, or when the secondary path fails.

The **no** form of this command cancels any currently outstanding revert timer. If the LSP is up when a no revert-timer is issued, the LSP will revert to the primary path. Otherwise the LSP reverts when the primary path is restored.

Default

no revert-timer

Parameters

timer-value

Specifies the amount of time, in one minute increments, between attempts to re-establish the LSP after it has failed.

Values 1 to 4320

Platforms

7705 SAR Gen 2

revert-timer

Syntax

revert-timer *seconds*

no revert-timer

Context

[Tree] (config>router>mpls>fwd-policies>fwd-policy revert-timer)

Full Context

configure router mpls forwarding-policies forwarding-policy revert-timer

Description

This command configures the revert timer in an MPLS forwarding policy.

When the primary direct or indirect next hop is restored and is added back into the routing table, CPM waits for an amount of time equal to the user-programmed revert timer before activating it and updating the data path. However, if the backup direct or indirect next hop fails while the timer is running, CPM activates it and updates the data path immediately.

A value of 0 disables the revert timer; meaning the policy reverts immediately.

The **no** form of this command removes the revert timer from the MPLS forwarding policy.

Default

revert-timer 0

Parameters

seconds

Specifies the revert-timer value, in number of seconds.

Values 1 to 600

Platforms

7705 SAR Gen 2

revert-timer

Syntax

revert-timer *revert-timer*

no revert-timer

Context

[Tree] (config>router>segment-routing>maintenance-policy revert-timer)

Full Context

configure router segment-routing maintenance-policy revert-timer

Description

This command configures the revert timer for SR Policy candidate paths.

The revert timer is started when the primary path (for example, the best preference programmed candidate path) recovers (for example, after the number of S-BFD sessions that are up is ≥ **threshold** and the **hold-down-timer** has expired) and switches back when the timer expires.

The **no** form of this command removes the revert timer from the SR policy.

Default

no revert-timer

Parameters

revert-timer

Specifies the revert timer, in minutes.

Values 1 to 4320

Platforms

7705 SAR Gen 2

26.56 revertive

revertive

Syntax

[**no**] **revertive**

Context

[Tree] (config>redundancy>multi-chassis>ipsec-domain revertive)

Full Context

configure redundancy multi-chassis ipsec-domain revertive

Description

This command configures whether to allow a revertive activity state after a designated active state recovers from an ineligibility event. The revertive function allows a router in an N:M domain to automatically take over as the active router in the domain, when it becomes eligible to do so.

The **no** form of this command reverts to the default value.

Default

no revertive

Platforms

7705 SAR Gen 2

26.57 revocation-check

revocation-check

Syntax

revocation-check {**crl** | **crl-optional**}

Context

[Tree] (config>system>security>pki>ca-profile revocation-check)

Full Context

configure system security pki ca-profile revocation-check

Description

This command specifies the revocation method the system uses to check the revocation status of certificate issued by the CA. If the **crl-optional** option is configured, when the user disables the **ca-profile**, the system tries to load the configured CRL (specified by the **crl-file** command). However, if the system fails to load the configured CRL for the following reasons, the system still brings the **ca-profile** operationally up, but leaves the CRL configured as non-existent:

- CRL file does not exist
- CRL is not properly encoded - maybe due to interrupted file transfer
- CRL does not match cert
- Wrong CRL version
- CRL expired

Note:
The **crl-optional** command option makes configuration of a valid CRL in a **ca-profile** optional. However, from a security point of view, it is important to always verify the revocation status of a certificate.

If the system needs to use the CRL of a specific CA profile to check the revocation status of an end-entity certificate, and the CRL is non-existent due to the preceding reasons, the system treats a case like this as being unable to get an answer from CRL and falls back to the next status verify method or default result.

If the system needs to check the revocation of a CA certificate in a certificate chain, and if the CRL is non-existent due to the preceding reasons, the system skips checking the revocation status of the CA certificate. For example, if CA1 is issued by CA2, if the revocation-check for CA2 is **crl-optional** and the CRL for CA2 is non-existent, the system does not check the certificate revocation status of CA1 and it is considered as "good”.

Note:
Users must shut down the **ca-profile** to change the **revocation-check** configuration.

Default

revocation-check crl

Parameters

crl

Specifies to use the configured CRL.

crl-optional

Specifies that the CRL is optional.

Platforms

7705 SAR Gen 2

26.58 revoke-key

revoke-key

Syntax

revoke-key card *cpm-slot* **serial-number** *cpm-serial-number* **confirmation-code** *code*

Context

[Tree] (admin>system>security>secure-boot revoke-key)

Full Context

admin system security secure-boot revoke-key

Description

This command revokes secure boot keys.

Parameters

cpm-slot

Specifies the CPM slot.

Values A,B

cpm-serial-number

Specifies the CPM serial number, up to 256 characters.

code

Specifies the signed software confirmation code, up to 32 characters.

Platforms

7705 SAR Gen 2

26.59 rib-management

rib-management

Syntax

rib-management

Context

[Tree] (config>service>vprn>bgp rib-management)

Full Context

configure service vprn bgp rib-management

Description

Commands in this context configure RIB management parameters.

Platforms

7705 SAR Gen 2

rib-management

Syntax

rib-management

Context

[Tree] (config>router>bgp rib-management)

Full Context

configure router bgp rib-management

Description

Commands in this context configure RIB management parameters.

Platforms

7705 SAR Gen 2

26.60 rib-priority

rib-priority

Syntax

rib-priority high {*prefix-list-name* | **tag** *tag*}

no rib-priority

Context

[Tree] (config>service>vprn>isis rib-priority)

Full Context

configure service vprn isis rib-priority

Description

This command enabled RIB prioritization for the IS-IS protocol and specifies the prefix list or IS-IS tag value that will be used to select the specific routes that should be processed through the IS-IS route calculation process at a higher priority.

The **no** form of this command disables RIB prioritization.

Default

no rib-priority

Parameters

prefix-list-name

Specifies the prefix list which is used to select the routes that are processed at a higher priority through the route calculation process.

tag tag-value

Specifies the tag value that is used to match IS-IS routes that are to be processed at a higher priority through the route calculation process.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

rib-priority

Syntax

rib-priority high

no rib-priority

Context

[Tree] (config>service>vprn>ospf>area>if rib-priority)

[Tree] (config>service>vprn>ospf3>area>if rib-priority)

Full Context

configure service vprn ospf area interface rib-priority

configure service vprn ospf3 area interface rib-priority

Description

This command enables RIB prioritization for the OSPF/OSPFv3 protocol. When enabled at the OSPF interface level, all routes learned through the associated OSPF interface will be processed through the OSPF route calculation process at a higher priority.

The **no** form of **rib-priority** command disables RIB prioritization at the associated level.

Default

no rib-priority

Platforms

7705 SAR Gen 2

rib-priority

Syntax

rib-priority {**high**} *prefix-list-name*

no rib-priority

Context

[Tree] (config>service>vprn>ospf3 rib-priority)

[Tree] (config>service>vprn>ospf rib-priority)

Full Context

configure service vprn ospf3 rib-priority

configure service vprn ospf rib-priority

Description

This command enabled RIB prioritization for the OSPF protocol and specifies the prefix list that will be used to select the specific routes that should be processed through the OSPF route calculation process at a higher priority.

The **no** form of **rib-priority** command disables RIB prioritization at the associated level.

Default

no rib-priority

Parameters

prefix-list-name

Specifies the prefix list which is used to select the routes that are processed at a higher priority through the route calculation process.

Platforms

7705 SAR Gen 2

rib-priority

Syntax

rib-priority high {*prefix-list-name* | **tag** *tag-value*}

no rib-priority

Context

[Tree] (config>router>isis rib-priority)

Full Context

configure router isis rib-priority

Description

This command enabled RIB prioritization for the IS-IS protocol and specifies the prefix list or IS-IS tag value that will be used to select the specific routes that should be processed through the IS-IS route calculation process at a higher priority.

The no rib-priority form of command disables RIB prioritization.

Default

no rib-priority high

Parameters

prefix-list-name

Specifies the prefix list which is used to select the routes that are processed at a higher priority through the route calculation process.

tag tag-value

Specifies the tag value that is used to match IS-IS routes that are to be processed at a higher priority through the route calculation process.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

rib-priority

Syntax

rib-priority {**high**} *prefix-list-name*

no rib-priority {**high**}

Context

[Tree] (config>router>ospf3 rib-priority)

[Tree] (config>router>ospf rib-priority)

Full Context

configure router ospf3 rib-priority

configure router ospf rib-priority

Description

This command enables RIB prioritization for the OSPF protocol and specifies the prefix list used to select the specific routes that should be processed through the OSPF route calculation process at a higher priority.

The **no** form of this command disables RIB prioritization at the associated level.

Default

no rib-priority high

Parameters

prefix-list-name

Specifies the prefix list, up to 32 characters, which is used to select the routes that are processed at a higher priority through the route calculation process.

Platforms

7705 SAR Gen 2

rib-priority

Syntax

rib-priority {**high**}

no rib-priority

Context

[Tree] (config>router>ospf>area>interface rib-priority)

[Tree] (config>router>ospf3>area>interface rib-priority)

Full Context

configure router ospf area interface rib-priority

configure router ospf3 area interface rib-priority

Description

This command enables RIB prioritization for the OSPF/OSPFv3 protocol. When enabled at the OSPF interface level, all routes learned through the associated OSPF interface are processed through the OSPF route calculation process at a higher priority.

The **no** form of this command disables RIB prioritization at the associated level.

Default

no rib-priority

Parameters

high

Specifies that the name of the prefix list which contains prefixes get high priority for RIB-download. The high priority prefixes are downloaded first to the RIB. In doing so, the convergence time for these prefixes is better.

Platforms

7705 SAR Gen 2

26.61 ring-node

ring-node

Syntax

ring-node *ring-node-name*

no ring-node

Context

[Tree] (config>service>epipe>sap ring-node)

Full Context

configure service epipe sap ring-node

Description

This command configures a multi-chassis ring-node for this SAP.

The **no** form of this command removes the name from the configuration.

Platforms

7705 SAR Gen 2

26.62 rip

rip

Syntax

[**no**] **rip**

Context

[Tree] (config>service>vprn rip)

Full Context

configure service vprn rip

Description

This command enables the RIP protocol on the given VPRN IP interface.

The **no** form of this command disables the RIP protocol from the given VPRN IP interface.

Default

no rip

Platforms

7705 SAR Gen 2

rip

Syntax

[**no**] **rip**

Context

[Tree] (config>router rip)

Full Context

configure router rip

Description

This command creates the context to configure the RIP protocol instance.

When a RIP instance is created, the protocol is enabled by default. To start or suspend execution of the RIP protocol without affecting the configuration, use the [**no**] **shutdown** command.

The **no** form of the command deletes the RIP protocol instance removing all associated configuration parameters.

Default

no rip

Platforms

7705 SAR Gen 2

26.63 rip-policy

rip-policy

Syntax

rip-policy *policy-name*

no rip-policy

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host rip-policy)

Full Context

configure subscriber-mgmt local-user-db ipoe host rip-policy

Description

This command configures the RIP policy name. This policy is applied to a subscriber IPv4 host to enable the BNG to learn RIP routes from the host. RIP routes are never sent to the hosts.

The **no** form of this command removes the RIP policy name from the configuration.

Parameters

policy-name

Specifies the RIP policy name, up to 32 characters.

Platforms

7705 SAR Gen 2

rip-policy

Syntax

rip-policy *policy-name* [**create**]

no rip-policy *policy-name*

Context

[Tree] (config>subscr-mgmt rip-policy)

Full Context

configure subscriber-mgmt rip-policy

Description

This command creates a RIP policy. This policy is applied to a subscriber IPv4 host to enable the BNG to learn RIP routes from the host. RIP routes are never sent to the hosts.

Parameters

policy-name

Specifies the RIP policy name up to 32 characters.

create

Keyword required to create the configuration context.

Platforms

7705 SAR Gen 2

26.64 ripng

ripng

Syntax

[**no**] **ripng**

Context

[Tree] (config>router ripng)

Full Context

configure router ripng

Description

This command creates the context to configure the RIPng protocol instance.

When a RIPng instance is created, the protocol is enabled by default. To start or suspend execution of the RIP protocol without affecting the configuration, use the [**no**] **shutdown** command.

The **no** form of this command deletes the RIP protocol instance removing all associated configuration parameters.

Default

no ripng

Platforms

7705 SAR Gen 2

26.65 rmon

rmon

Syntax

rmon

Context

[Tree] (config>system>thresholds rmon)

Full Context

configure system thresholds rmon

Description

This command creates the context to configure generic RMON alarms and events.

Generic RMON alarms can be created on any SNMP object-ID that is valid for RMON monitoring (for example, an integer-based datatype).

The configuration of an event controls the generation and notification of threshold crossing events configured with the alarm command.

Platforms

7705 SAR Gen 2

26.66 robust-count

robust-count

Syntax

robust-count *robust-count*

no robust-count

Context

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping robust-count)

[Tree] (config>service>vpls>mesh-sdp>mld-snooping robust-count)

[Tree] (config>service>vpls>sap>igmp-snooping robust-count)

[Tree] (config>service>vpls>igmp-snooping robust-count)

[Tree] (config>service>vpls>sap>mld-snooping robust-count)

[Tree] (config>service>vpls>spoke-sdp>mld-snooping robust-count)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping robust-count)

Full Context

configure service vpls mesh-sdp igmp-snooping robust-count

configure service vpls mesh-sdp mld-snooping robust-count

configure service vpls sap igmp-snooping robust-count

configure service vpls igmp-snooping robust-count

configure service vpls sap mld-snooping robust-count

configure service vpls spoke-sdp mld-snooping robust-count

configure service vpls spoke-sdp igmp-snooping robust-count

Description

If the **send-queries** command is enabled, this parameter allows tuning for the expected packet loss on a SAP or SDP. The robust-count variable allows tuning for the expected packet loss on a subnet and is comparable to a retry count. If this SAP or SDP is expected to be 'lossy', this parameter may be increased. IGMP snooping on this SAP or SDP is robust to (robust-count-1) packet losses.

If **send-queries** is not enabled, this parameter will be ignored.

Default

robust-count 2

Parameters

robust-count

Specifies the robust count for the SAP or SDP

Values 2 to 7 (for **config>service>vpls>sap>igmp-snooping**) 1 to 255 (for **config>service>vpls>igmp-snooping**)

Platforms

7705 SAR Gen 2

robust-count

Syntax

robust-count *robust-count*

no robust-count

Context

[Tree] (config>service>vprn>mld robust-count)

[Tree] (config>service>vprn>igmp robust-count)

Full Context

configure service vprn mld robust-count

configure service vprn igmp robust-count

Description

This command configures the robust count. The robust-count variable allows tuning for the expected packet loss on a subnet. If a subnet anticipates losses, the robust-count variable can be increased.

Default

robust-count 2

Parameters

robust-count

Specifies the robust count value.

Values 2 to 10

Platforms

7705 SAR Gen 2

robust-count

Syntax

robust-count *robust-count*

no robust-count

Context

[Tree] (config>router>igmp robust-count)

Full Context

configure router igmp robust-count

Description

This command configures the robust count. The robust-count variable allows tuning for the expected packet loss on a subnet. If a subnet anticipates losses, the robust-count variable can be increased.

Default

robust-count 2

Parameters

robust-count

Specify the robust count value.

Values 2 to 10

Platforms

7705 SAR Gen 2

robust-count

Syntax

robust-count *robust-count*

no robust-count

Context

[Tree] (config>router>mld robust-count)

Full Context

configure router mld robust-count

Description

This command configures the robust count. The robust-count variable allows tuning for the expected packet loss on a subnet. If a subnet anticipates losses, the robust-count variable can be increased.

Default

robust-count 2

Parameters

robust-count

Specify the robust count value.

Values 2 to 10

Platforms

7705 SAR Gen 2

robust-count

Syntax

robust-count *robust-count*

no robust-count

Context

[Tree] (config>service>pw-template>igmp-snooping robust-count)

Full Context

configure service pw-template igmp-snooping robust-count

Description

If the **send-queries** command is enabled, this parameter allows tuning for the expected packet loss. The robust-count variable allows tuning for the expected packet loss on a subnet and is comparable to a retry count.

If send-queries is not enabled, this parameter will be ignored.

Default

robust-count 2

Parameters

robust-count

Specifies the robust count for the SAP or SDP.

Values 2 to 7

Platforms

7705 SAR Gen 2

26.67 rollback

rollback

Syntax

rollback

Context

[Tree] (config>system rollback)

Full Context

configure system rollback

Description

Configure parameters of the classic CLI configuration rollback functionality. Configuration rollback provides the ability to undo configuration and revert back to previous router configuration states.

Platforms

7705 SAR Gen 2

rollback

Syntax

rollback

Context

[Tree] (admin rollback)

Full Context

admin rollback

Description

Commands in this context configure rollback operations.

Platforms

7705 SAR Gen 2

26.68 rollback-location

rollback-location

Syntax

rollback-location *file-url* /*rollback filename*

no rollback-location

Context

[Tree] (config>system>rollback rollback-location)

Full Context

configure system rollback rollback-location

Description

The location and name of the rollback checkpoint files is configurable to be local (on compact flash) or remote. The *file-url* must not contain a suffix (just a path/directory + filename). The suffixes for rollback checkpoint files are ".rb", ".rb.1", ..., ".rb.9" and are automatically appended to rollback checkpoint files.

Default

no rollback-location

Parameters

file-url

Specifies the URL or rollback filename.

Values

local-url | *remote-url*	
local-url	[*cflash-id/*][*file-path*] up to 200 characters, including cflash-id directory length of up to 99 characters each
remote-url	[{ftp://}*login*:*pswd*@ *remote-locn*/][*file-path*] up to 255 characters, directory length of up to 99 characters each
remote-locn	[*hostname* | *ipv4-address* | *ipv6-address*]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D

	interface - up to 32 characters each, for link local addresses
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

rollback-filename

Specifies the rollback file name.

Values suffixed with .rb, .rb.1 up to .9 during rollback checkpoint creation

Platforms

7705 SAR Gen 2

26.69 rollback-sync

rollback-sync

Syntax

rollback-sync

Context

[Tree] (admin>redundancy rollback-sync)

Full Context

admin redundancy rollback-sync

Description

This command copies the entire set of rollback checkpoint files from the active CPM CF to the standby CPM CF.

Platforms

7705 SAR Gen 2

rollback-sync

Syntax

[**no**] **rollback-sync**

Context

[Tree] (config>redundancy rollback-sync)

Full Context

configure redundancy rollback-sync

Description

The operator can enable automatic synchronization of classic CLI rollback checkpoint files between the active CPM and standby CPM. When this automatic synchronization is enabled, a classic CLI **rollback save** causes the new classic CLI checkpoint file to be saved on both the active and standby CPMs. The suffixes of the old checkpoint files on both active and standby CPMs are incremented. Note that automatic sync only causes the one new checkpoint file to be copied to both CFs (the other checkpoint files are not automatically copied from active to standby but that can be done manually with **admin redundancy rollback-sync**).

Automatic synchronization of classic CLI rollback checkpoint files across CPMs is only performed if the rollback-location is configured as a local file-url (for example, "cf3:/rollback-files/rollback). Synchronization is not done if the rollback-location is remote.

The **config redundancy synchronize {boot-env | config}** and **admin redundancy synchronize {boot-env | config}** do not apply to classic CLI rollback checkpoint files. These commands do not manually or automatically sync classic CLI rollback checkpoint files. The dedicated **rollback-sync** command must be used to sync classic CLI rollback checkpoint files.

Default

no rollback-sync

Platforms

7705 SAR Gen 2

26.70 rollover

rollover

Syntax

rollover *minutes* [**retention** *hours*]

no rollover

Context

[Tree] (config>log>file-id rollover)

Full Context

configure log file-id rollover

Description

This command configures how often an event or accounting log is rolled over or partitioned into a new file.

An event or accounting log is actually composed of multiple, individual files. The system creates a new file for the log based on the **rollover** time, expressed in minutes.

The *retention* option, expressed in hours, allows you to modify the default time to keep the file in the system. The retention time is based on the rollover time of the file.

If logs are needed to be retained for more than 16 days, use a CRON job to move the logs to a different location, either on a local drive or a remote server. For more information, contact Nokia support.

When multiple **rollover** commands for a *file-id* are entered, the last command overwrites the previous command.

The **no** form of this command reverts to the default values.

Default

rollover 1440 retention 12

Parameters

minutes

Specifies the rollover time, in minutes.

Values 5 to 10080

retention hours

Specifies the retention period in hours, expressed as a decimal integer. The retention time is based on the time creation time of the file. The file becomes a candidate for removal once the creation datestamp + rollover time + retention time is less than the current timestamp.

Default 12

Values 1 to 500

Platforms

7705 SAR Gen 2

26.71 root

root

Syntax

root

Context

[Tree] (config>qos>policer-control-policy root)

Full Context

configure qos policer-control-policy root

Description

The **root** node contains the policer control policies configuration parameters for the root arbiter. Within the node, the parent policer's maximum rate limit can be set, the strict priority level, and fair threshold portions may be defined per priority level.

The root node always exists and does not need to be created.

Platforms

7705 SAR Gen 2

26.72 root-guard

root-guard

Syntax

[**no**] **root-guard**

Context

[Tree] (config>service>vpls>sap>stp root-guard)

[Tree] (config>service>vpls>spoke-sdp>stp root-guard)

[Tree] (config>service>template>vpls-sap-template>stp root-guard)

Full Context

configure service vpls sap stp root-guard

configure service vpls spoke-sdp stp root-guard

configure service template vpls-sap-template stp root-guard

Description

This command specifies whether this port is allowed to become an STP root port. It corresponds to the restrictedRole parameter in 802.1Q. If set, it can cause lack of spanning tree connectivity.

Default

no root-guard

Platforms

7705 SAR Gen 2

root-guard

Syntax

[**no**] **root-guard**

Context

[Tree] (config>service>pw-template>stp root-guard)

Full Context

configure service pw-template stp root-guard

Description

This command specifies whether this port is allowed to become an STP root port. It corresponds to the restrictedRole parameter in 802.1Q. If set, it can cause lack of spanning tree connectivity.

Default

no root-guard

Platforms

7705 SAR Gen 2

26.73 route-admin-tag-policy

route-admin-tag-policy

Syntax

[**no**] **route-admin-tag-policy** *policy-name*

Context

[Tree] (config>router>admin-tags route-admin-tag-policy)

Full Context

configure router admin-tags route-admin-tag-policy

Description

This command configures a route admin tag policy.

Up to 2,000 policies can be configured per system.

The **no** form of this command removes the route admin tag policy.

Parameters

policy-name

The name of the route admin tag policy, up to 32 characters.

Platforms

7705 SAR Gen 2

26.74 route-distinguisher

route-distinguisher

Syntax

route-distinguisher auto-rd

no route-distinguisher

route-distinguisher *rd*

Context

[Tree] (config>service>vpls>bgp route-distinguisher)

[Tree] (config>service>epipe>bgp route-distinguisher)

Full Context

configure service vpls bgp route-distinguisher

configure service epipe bgp route-distinguisher

Description

This command configures the Route Distinguisher (RD) component that will be signaled in the MP-BGP NLRI for L2VPN and EVPN families. This value will be used for BGP-AD, BGP VPLS and BGP multi-homing NLRI if these features are configured.

If this command is not configured, the RD is automatically built using the BGP-AD VPLS ID. The following rules apply:

- if BGP AD VPLS-id is configured and no RD is configured under BGP node - RD=VPLS-ID
- if BGP AD VPLS-id is not configured then an RD value must be configured under BGP node (this is the case when only BGP VPLS is configured)
- if BGP AD VPLS-id is configured and an RD value is also configured under BGP node, the configured RD value prevails

Values and format (6 bytes, other 2 bytes of type will be automatically generated)

Alternatively, the **auto-rd** option allows the system to automatically generate an RD based on the **bgp-auto-rd-range** command configured at the service level. For **BGP-EVPN** enabled VPLS and Epipe services, the **route-distinguisher** value can also be auto-derived from the **evi** value (**config>service>vpls>bgp-evpn>evi** or **config>service>epipe>bgp-evpn>evi**) if this command is

not configured. See the **config>service>system>bgp-evpn>eth-seg>service-carving>manual evi** command description for more information.

Parameters

ip-addr:comm-val

Specifies the IP address.

Values *ip-addr*: a.b.c.d

comm-val: 0 to 65535

as-number:ext-comm-val

Specifies the AS number.

Values *as-number*: 1 to 65535

ext-comm-val: 0 to 4294967295

auto-rd

The system will generate an RD for the service according to the IP address and range configured in the **bgp-auto-rd-range** command.

Platforms

7705 SAR Gen 2

route-distinguisher

Syntax

route-distinguisher [*ip-addr:comm-val* | *as-number:ext-comm-val*]

no route-distinguisher

Context

[Tree] (config>service>system>bgp-evpn route-distinguisher)

Full Context

configure service system bgp-evpn route-distinguisher

Description

This command configures the Route Distinguisher (RD) component that will be signaled in the MP-BGP NLRI for EVPN corresponding to the base EVPN instance (Ethernet Segment routes). If the route-distinguisher component is not configured, the system will use system:ip-address as the default route-distinguisher

Default

no route-distinguisher

Parameters

ip-addr:comm-val

Specifies the IP address.

Values *ip-addr*: a.b.c.d

comm-val: 0 to 65535

as-number:ext-comm-val

Specifies the AS number.

Values *as-number*: 1 to 65535

ext-comm-val: 0 to 4294967295

Platforms

7705 SAR Gen 2

route-distinguisher

Syntax

route-distinguisher *rd*

route-distinguisher auto-rd

no route-distinguisher

Context

[Tree] (config>service>vprn>bgp-ipvpn>mpls route-distinguisher)

[Tree] (config>service>vprn>bgp-evpn>mpls route-distinguisher)

Full Context

configure service vprn bgp-ipvpn mpls route-distinguisher

configure service vprn bgp-evpn mpls route-distinguisher

Description

This command specifies an identifier attached to a route, which enables the user to identify the VPN to which the route belongs. Each routing instance must have a unique (within the carrier's domain) route distinguisher (RD) associated with it.

Alternatively, the **auto-rd** option allows the system to automatically generate an RD based on the **configure service system bgp-auto-rd-range** command.

The **no** form of this command removes the RD configuration.

Default

no route-distinguisher

Parameters

auto-rd

Keyword that allows the system to generate an RD for the service according to the IP address and range configured in the **bgp-auto-rd-range** command.

rd

Specifies the route distinguisher.

Values rd: *ip-addr:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*

ip-addr: a.b.c.d

comm-val: [0 to 65535]

2byte-asnumber: [1 to 65535]

ext-comm-val: [0 to 4294967295]

4byte-asnumber: [1 to 4294967295]

Platforms

7705 SAR Gen 2

route-distinguisher

Syntax

route-distinguisher

Context

[Tree] (config>service>vprn route-distinguisher)

Full Context

configure service vprn route-distinguisher

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

26.75 route-distinguisher-list

route-distinguisher-list

Syntax

route-distinguisher-list *name*

no route-distinguisher-list *name*

Context

[Tree] (config>router>policy-options route-distinguisher-list)

Full Context

configure router policy-options route-distinguisher-list

Description

This command creates a list of entries used to match the RD in BGP routes of specific address families.

Parameters

name

Specifies the name of the RD list, up to 64 characters.

Platforms

7705 SAR Gen 2

route-distinguisher-list

Syntax

route-distinguisher-list *name*

no route-distinguisher-list *name*

Context

[Tree] (config>router>policy-options>policy-statement>entry>from route-distinguisher-list)

Full Context

configure router policy-options policy-statement entry from route-distinguisher-list

Description

This command configures a route distinguisher (RD) list as a match criterion for the policy statement entry.

This match condition is supported by policies applied as VRF import or BGP peer import policies. A BGP route can match a policy entry with this match criterion if the NLRI field contains an RD that is matched by at least one of the entries in the **route-distinguisher-list**.

BGP routes belonging to address families other than VPN-IPv4, VPN-IPv6, MCAST-VPN-IPv4, MCAST-VPN-IPv6, EVPN, FlowSpec-VPN IPv4, FlowSpec-VPN IPv6, MVPN-IPv4 or MVPN-IPv6 routes do not match policy entries with this match criterion.

Parameters

name

Specifies the (possibly parameterized) name of an RD list.

Platforms

7705 SAR Gen 2

26.76 route-exists

route-exists

Syntax

route-exists *expression*

no route-exists

Context

[Tree] (config>router>policy-options>policy-statement>entry>cond-expr route-exists)

Full Context

configure router policy-options policy-statement entry conditional-expression route-exists

Description

This command is used to specify a route existence expression to control evaluation of the policy entry. If the route existence expression evaluates to 'true' the matching and action commands of the policy entry are applied as normal. If the route existence expression evaluates to 'false' the entire policy entry is skipped and processing continues with the next entry; however, conditional expressions are only parsed when the route policy is used as a BGP export policy or VRF export policy.

Default

no route-exists

Parameters

expression

"["<pfx-list-name>"]" [all | none]

If neither the **all** nor the **none** keyword are used the match logic is 'any' – that is, the route expression evaluates as 'true' if any exact match entry in the referenced prefix-list has an active route in the route table associated with the policy.

all – the route expression evaluates as 'true' only if all the exact match entries in the referenced prefix-list have an active route in the route table associated with the policy.

none – the route expression evaluates as 'true' only if none of the exact match entries in the referenced prefix-list have an active route in the route table associated with the policy.

Platforms

7705 SAR Gen 2

26.77 route-next-hop

route-next-hop

Syntax

route-next-hop {**system-ipv4** | **system-ipv6** | *ip-address*}

Context

[Tree] (config>service>vpls>bgp-evpn>mpls route-next-hop)

[Tree] (config>service>epipe>bgp-evpn>mpls route-next-hop)

Full Context

configure service vpls bgp-evpn mpls route-next-hop

configure service epipe bgp-evpn mpls route-next-hop

Description

This command configures the next hop of the EVPN routes.

Default

route-next-hop system-ipv4

Parameters

system-ipv4

Specifies the system IPv4 address as the next hop for the service EVPN routes.

system-ipv6

Specifies the system IPv6 address as the next hop for the service EVPN routes.

ip-address

Specifies the IPv4 address value as the next hop for the service EVPN.

Values a.b.c.d

Platforms
7705 SAR Gen 2

26.78 route-next-hop-policy

route-next-hop-policy

Syntax
route-next-hop-policy

Context
[Tree] (config>router route-next-hop-policy)

Full Context
configure router route-next-hop-policy

Description
This command creates the context to configure route next-hop policies.

Platforms
7705 SAR Gen 2

26.79 route-preference

route-preference

Syntax
route-preference primary {**inband** | **outband**} **secondary** {**inband** | **outband** | **none**}
no route-preference

Context
[Tree] (config>log route-preference)

Full Context
configure log route-preference

Description

This command specifies the primary and secondary routing preference for traffic generated for SNMP notifications and syslog messages. If the remote destination is not reachable through the routing context specified by primary route preference then the secondary routing preference will be attempted.

The **no** form of this command reverts to the default values.

Default

no route-preference

Parameters

primary

Specifies the primary routing preference for traffic generated for SNMP notifications and syslog messages.

Default outband

secondary

Specifies the secondary routing preference for traffic generated for SNMP notifications and syslog messages. The routing context specified by the secondary route preference will be attempted if the remote destination was not reachable by the primary routing preference, specified by primary route preference. The value specified for the secondary routing preference must be distinct from the value for primary route preference.

Default inband

inband

Specifies that the logging utility will attempt to use the base routing context to send SNMP notifications and syslog messages to remote destinations.

outband

Specifies that the logging utility will attempt to use the management routing context to send SNMP notifications and syslog messages to remote destinations.

none

Specifies that no attempt will be made to send SNMP notifications and syslog messages to remote destinations.

Platforms

7705 SAR Gen 2

route-preference

Syntax

route-preference {**both** | **inband** | **outband**}

no route-preference

Context

[Tree] (config>system>security>ldap route-preference)

[Tree] (config>system>security>radius route-preference)

[Tree] (config>system>security>tacplus route-preference)

Full Context

configure system security ldap route-preference

configure system security radius route-preference

configure system security tacplus route-preference

Description

This command specifies the routing preference to reach the AAA server. If the configured option is to use both in-band and out-of-band routes, the out-of-band routes in the management routing instance are used to reach the server before the in-band routes in the Base routing instance.

The **no** form of this command reverts to the default value.

Default

route-preference both

Parameters

both

Specifies the use of out-of-band routes before in-band routes.

inband

Specifies the use of in-band routes only.

outband

Specifies the use of out-of-band routes only.

Platforms

7705 SAR Gen 2

route-preference

Syntax

route-preference {**both** | **inband** | **outband**}

no route-preference

Context

[Tree] (config>router>pcep>pcc>peer route-preference)

Full Context

configure router pcep pcc peer route-preference

Description

This command specifies the routing preference to reach the PCE server. If the configured option is to use both in-band and out-of-band routes, the out-of-band routes in the management routing instance are used to reach the server before the in-band routes in the Base routing instance.

The **no** form of this command reverts to the default value.

Default

route-preference both

Parameters

both

Specifies the use of out-of-band routes before in-band routes.

inband

Specifies the use of in-band routes only.

outband

Specifies the use of out-of-band routes only.

Platforms

7705 SAR Gen 2

26.80 route-recovery-wait

route-recovery-wait

Syntax

route-recovery-wait *seconds*

no route-recovery-wait

Context

[Tree] (config>log>app-route-notifications route-recovery-wait)

Full Context

configure log app-route-notifications route-recovery-wait

Description

The time delay that must pass before notifying specific CPM applications after the recovery or change of a route during normal operation.

The **no** form of this command disables the time-delay configuration.

Default

no route-recovery-wait

Parameters

seconds

Time delay in seconds.

Values 1 to 100

Platforms

7705 SAR Gen 2

26.81 route-refresh

route-refresh

Syntax

route-refresh [**neighbor** *ip-address* | **group** *name*]

no route-refresh

Context

[Tree] (debug>router>bgp route-refresh)

Full Context

debug router bgp route-refresh

Description

This command enables debugging for BGP route-refresh.

The **no** form of this command disables debugging.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x [-interface] (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H

- d: [0 to 255]D
- interface: up to 32 characters for link local addresses

group *name*

Debugs only events affecting the specified peer group name, up to 64 characters, and associated neighbors.

Platforms

7705 SAR Gen 2

26.82 route-table

route-table

Syntax

route-table [*ip-prefix*/*prefix-length*]

route-table *ip-prefix*/*prefix-length* **longer**

no route-table

Context

[Tree] (debug>router>ip route-table)

Full Context

debug router ip route-table

Description

This command configures route table debugging.

Parameters

ip-prefix/prefix-length

The IP prefix for prefix list entry in dotted decimal notation.

Values			
	ipv4-prefix	a.b.c.d (host bits must be 0)	
	ipv4-prefix-length	0 to 32	
	ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)	
		x:x:x:x:x:x:d.d.d.d	
		x:	[0 to FFFF]H
		d:	[0 to 255]D

ipv6-prefix-length 0 to 128

longer

Specifies the prefix list entry matches any route that matches the specified *ip-prefix* and prefix *mask* length values greater than the specified *mask*.

Platforms

7705 SAR Gen 2

26.83 route-table-import

route-table-import

Syntax

route-table-import *policy-name*

no route-table-import

Context

[Tree] (config>service>vprn>bgp>rib-management>ipv4 route-table-import)

[Tree] (config>service>vprn>bgp>rib-management>label-ipv4 route-table-import)

[Tree] (config>service>vprn>bgp>rib-management>ipv6 route-table-import)

Full Context

configure service vprn bgp rib-management ipv4 route-table-import

configure service vprn bgp rib-management label-ipv4 route-table-import

configure service vprn bgp rib-management ipv6 route-table-import

Description

This command specifies the name of a route policy to control the importation of active routes from the IP route table into one of the BGP RIBs.

If the **route-table-import** command is not configured, or if the command refers to an empty policy, all non-BGP routes from the IP route table are imported into the applicable RIB.

If the **route-table-import** command is configured, then routes dropped or rejected by the configured policy are not installed in the associated RIB. Rejected routes cannot be advertised to BGP peers associated with the RIB, but they can still be used to resolve BGP next-hops of routes in that RIB. If the active route for a prefix is rejected by the **route-table-import** policy, then the best BGP route for that prefix in the BGP RIB can be advertised to peers as though it is used.

Aggregate routes are always imported into each RIB, independent of the **route-table-import** policy.

Route modifications specified in the actions of a **route-table-import** policy are ignored and have no effect on the imported routes.

Default

no route-table-import

Parameters

policy-name

Specifies the name of a policy-statement (up to 64 characters).

Platforms

7705 SAR Gen 2

route-table-import

Syntax

route-table-import *policy-name*

no route-table-import

Context

[Tree] (config>router>bgp>rib-management>label-ipv6 route-table-import)

[Tree] (config>router>bgp>rib-management>ipv4 route-table-import)

[Tree] (config>router>bgp>rib-management>ipv6 route-table-import)

[Tree] (config>router>bgp>rib-management>label-ipv4 route-table-import)

Full Context

configure router bgp rib-management label-ipv6 route-table-import

configure router bgp rib-management ipv4 route-table-import

configure router bgp rib-management ipv6 route-table-import

configure router bgp rib-management label-ipv4 route-table-import

Description

This command specifies the name of a policy to control the importation of active routes from the IP route table into one of the BGP RIBs.

If the **route-table-import** command is not configured, or if the command refers to an empty policy, all non-BGP routes from the IP route table are imported into the applicable RIB.

If the **route-table-import** command is configured, then routes dropped or rejected by the configured policy are not installed in the associated RIB. Rejected routes cannot be advertised to BGP peers associated with the RIB, but they can still be used to resolve BGP next-hops of routes in that RIB. If the active route for a prefix is rejected by the **route-table-import** policy, then the best BGP route for that prefix in the BGP RIB can be advertised to peers as though it is used.

Aggregate routes are always imported into each RIB, independent of the **route-table-import** policy.

Route modifications specified in the actions of a **route-table-import** policy are ignored and have no effect on the imported routes.

Default

no route-table-import

Parameters

policy-name

Specifies the name of a policy-statement (up to 64 characters).

Platforms

7705 SAR Gen 2

26.84 route-target

route-target

Syntax

route-target {*ext-community* **| {[export** *ext-community***][import** *ext-community***]}}**

no route-target

Context

[Tree] (config>service>epipe>bgp route-target)

[Tree] (config>service>vpls>bgp route-target)

Full Context

configure service epipe bgp route-target

configure service vpls bgp route-target

Description

This command configures the route target (RT) component that will be signaled in the related MP- BGP attribute to be used for BGP auto-discovery, BGP VPLS, BGP multi-homing and EVPN if these features are configured in this VPLS service, or for BGP multi-homing, BGP-VPWS and EVPN in case of Epipe services.

If this command is not used in VPLS services, the RT is built automatically using the VPLS ID. The extended community can have the same two formats as the VPLS ID, a two-octet AS-specific extended community, IPv4 specific extended community. For BGP EVPN enabled VPLS and Epipe services, the route target can also be auto-derived from the **evi** value (**config>service>vpls>bgp-evpn>evi** or **config>service>epipe>bgp-evpn>evi**) if this command is not configured.

Parameters

export *ext-community*

Specifies communities allowed to be sent to remote PE neighbors.

import *ext-community*

Specifies communities allowed to be accepted from remote PE neighbors.

Platforms

7705 SAR Gen 2

26.85 route-target-list

route-target-list

Syntax

route-target-list *comm-id* [*comm-id*]]

no route-target-list [*comm-id*]

Context

[Tree] (config>router>bgp route-target-list)

Full Context

configure router bgp route-target-list

Description

This command specifies the route target(s) to be accepted from or advertised to peers. If the **route-target-list** is a non-null list, only routes with one or more of the given route targets are accepted from or advertised to peers.

The **route-target-list** is assigned at the global level and applies to all peers connected to the system.

This command is only applicable if the router is a route-reflector server.

The **no** form of this command with a specified route target community removes the specified community from the **route-target-list**. The **no** form of this command entered without a route target community removes all communities from the list.

Default

no route-target-list

Parameters

comm-id

Specifies up to 15 route target communities.

Values [**target:** {*ip-address*:*comm-val* | *2byte-asnumber*:*ext-comm-val* | *4byte-asnumber*:*comm-val*}

where:

- *ip-address* — a.b.c.d

- *comm-val* — 0 to 65535
- *2byte-asnumber* — 0 to 65535
- *ext-comm-val* — 0 to 4294967295
- *4byte-asnumber* — 0 to 4294967295

Platforms

7705 SAR Gen 2

26.86 route-unknown

route-unknown

Syntax

[**no**] **route-unknown** [{*ip-prefix/mask* | *ipv6-address/prefix-length*}]

Context

[Tree] (config>vrrp>policy>priority-event route-unknown)

Full Context

configure vrrp policy priority-event route-unknown

Description

This command creates a context to configure a route unknown priority control event that monitors the existence of a specific active IP route prefix within the routing table.

The **route-unknown** command configures a priority control event that defines a link between the VRRP priority control policy and the Route Table Manager (RTM). The RTM registers the specified route prefix as monitored by the policy. If any change (add, delete, new next hop) occurs relative to the prefix, the policy is notified and takes correct action according to the priority event definition. If the route prefix exists and is active in the routing table according to the conditions defined, the event is in the cleared state. If the route prefix is removed, becomes inactive or fails to meet the event criteria, the event is in the set state.

The command creates a **route-unknown** node identified by *prefix/mask-length* and containing event control commands.

Multiple unique (different *prefix/mask-length*) **route-unknown** event nodes can be configured within the **priority-event** node up to the maximum limit of 32 events.

The **route-unknown** command can reference any valid IP address mask-length pair. The IP address and associated mask length define a unique IP router prefix. The dynamic monitoring of the route prefix results in one of the event operational states listed in Table 83: Route-unknown Operational States.

Table 83: Route-unknown Operational States

route-unknown Operational State	Description
Set – non-existent	The route does not exist in the route table
Set – inactive	The route exists in the route table but is not being used
Set – wrong next hop	The route exists in the route table but does not meet the **next-hop** requirements
Set – wrong protocol	The route exists in the route table but does not meet the **protocol** requirements
Set – less specific found	The route exists in the route table but does is not an exact match and does not meet any **less-specific** requirements
Set – default best match	The route exists in the route table as the default route but the default route is not allowed for route matching
Cleared – less specific found	A less specific route exists in the route table and meets all criteria including the **less-specific** requirements
Cleared – found	The route exists in the route table manager and meets all criteria

An existing route prefix in the RTM must be active (used by the IP forwarding engine) to clear the event operational state. It may be less specific (the defined prefix may be contained in a larger prefix according to Classless Inter-Domain Routing (CIDR) techniques) if the event has the **less-specific** statement defined. The less specific route that incorporates the router prefix may be the default route (0.0.0.0) if the **less-specific allow-default** statement is defined. The matching prefix may be required to have a specific next hop IP address if defined by the event **next-hop** command. Finally, the source of the RTM prefix may be required to be one of the dynamic routing protocols or be statically defined if defined by the event **protocol** command. If an RTM prefix is not found that matches all the above criteria (if defined in the event control commands), the event is considered to be set. If a matching prefix is found in the RTM, the event is considered to be cleared.

When an event transitions from clear to set, the set is processed immediately and must be reflected in the associated virtual router instances in-use priority value. As the event transitions from clear to set, a hold-set timer is loaded with the value configured by the events **hold-set** command. This timer prevents the event from clearing until it expires, damping the effect of event flapping. If the event clears and becomes set again before the hold-set timer expires, the timer is reset to the **hold-set** value, extending the time before another clear can take effect.

The **no** form of the command is used to remove the specific *prefix/mask-length* monitoring event. The event can be removed at anytime. When the event is removed, the in-use priority of all associated virtual router instances must be reevaluated. The events **hold-set** timer has no effect on the removal procedure.

Default

no route-unknown — No route unknown priority control events are defined for the priority control event policy.

Parameters

ip-prefix/mask

The IP prefix address in dotted decimal notation and the subnet mask length expressed as a decimal integer associated with the IP *prefix* defining the route prefix to be monitored by the route unknown priority control event.

Values

ip-prefix/mask:	ip-prefix	a.b.c.d (host bits must be 0)
	mask	0 to 32

ipv6-address/prefix-length

The IPv6 address of the host for which the specific event will monitor connectivity. The *ipv6-address* can only be monitored by a single event in this policy. The IPv6 address can be monitored by multiple VRRP priority control policies. The IPv6 address can be used in one or multiple **ping** requests. Each VRRP priority control **host-unreachable** and **ping** destined to the same *ipv6-address* is uniquely identified on a per message basis. Each session originates a unique identifier value for the ICMP echo request messages it generates. This allows received ICMP echo reply messages to be directed to the appropriate sending application.

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0..FFFF]H
prefix-length	0 to 128	

Platforms

7705 SAR Gen 2

26.87 router

router

Syntax

router *router-instance*

router service-name *service-name*

no router

Context

[Tree] (config>aaa>radius-srv-plcy>servers router)

Full Context

configure aaa radius-server-policy servers router

Description

This command specifies the virtual router instance applicable for the set of configured RADIUS servers. This value cannot be changed once a RADIUS server is configured for this policy.

The **no** form of this command reverts to the default.

Parameters

router-instance

Specifies the router instance.

Values

service-name	Service name, up to 64 characters.	
	router-instance:	router-name, service-id
	router-name:	Base, management
	service-id:	1 to 2147483647

service-name

Specifies the router name service-id up to 64 characters.

Platforms

7705 SAR Gen 2

router

Syntax

router [*router-instance*] [**create**]

no router [*router-instance*]

Context

[Tree] (config router)

Full Context

configure router

Description

Commands in this context configure router parameters including interfaces, route policies and protocols. This command is also used to create CPM router instances.

For CPM router instances, this command enters or creates a user-created CPM router instance. A CPM router instance is a not a VPRN router instance. VPRN router instances are configured under **configure service vprn**. CPM router instances are the only type of non-VPRN router instances that can be created

by a user, and they have a user-defined name. CPM router instances only use CPM/CCM ethernet ports as interfaces.

Parameters

router-instance

Specifies the router name or CPM router instance.

Values

router-instance : *router name*

router-name	Base | management | *cpm-vr-name*
cpm-vr-name	[32 characters maximum]

Default Base

create

Mandatory keyword when creating a router instance. The create keyword requirement can be enabled or disabled in the **environment**>**create** context.

Platforms

7705 SAR Gen 2

router

Syntax

router [*router-instance*]

router service-name *service-name*

Context

[Tree] (debug router)

Full Context

debug router

Description

Commands in this context enable debugging of various protocols and areas of a *router-instance*.

Parameters

router-instance

Specifies the router name, CPM router instance, or service ID.

Values *router-name* or *service-id*

router-instance : *router-name*

router-name Base | management | *cpm-vr-name*

cpm-vr-name [32 characters maximum]

service-id: 1 to 2147483647

Default Base

service-name

Specifies the service name, up to 64 characters.

Platforms

7705 SAR Gen 2

router

Syntax

router *router-instance*

router service *vprn-service-name*

Context

[Tree] (config>system>file-trans-prof router)

Full Context

configure system file-transmission-profile router

Description

This command specifies the routing instance that the transport protocol uses.

Default

router Base

Parameters

router-instance

Specifies the router instance on which the file transmission connection will be established.

This variant of this command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **router service** *vprn-service-name* variant can be used in all configuration modes.

Values {router-name | *vprn-svc-id*}

router-name: Base, management

router-name is an alias for input only. The *router-name* gets replaced with

an id automatically by SR OS in the configuration).

vprn-svc-id: 1 to 2147483647

Default Base

service vprn-service-name

Identifies the service, up to 64 characters.

Platforms

7705 SAR Gen 2

router

Syntax

router *router-instance*

router service *vprn-service-instance*

no router

Context

[Tree] (config>system>management-interface>remote-management router)

Full Context

configure system management-interface remote-management router

Description

This command defines the router instance in which all remote managers are reachable.

If this command is also configured for a specific manager in the **config**>**system**> **management-interface**>**remote-management**>**manager** context, that configuration takes precedence.

The **no** form of this command configures management as the router (default).

Default

router management

Parameters

router-instance

Specifies a router instance on which the remote management connection is established, up to 32 characters.

service *vprn-service-instance*

Specifies a VPRN service instance, up to 64 characters.

Platforms

7705 SAR Gen 2

router

Syntax

router *router-instance*

router service *vprn-service-instance*

no router

Context

[Tree] (config>system>management-interface>remote-management>manager router)

Full Context

configure system management-interface remote-management manager router

Description

This command defines the router instance in which this manager is reachable.

This command takes precedence over the same command configured in the global context (**config**>**system**>**management-interface**>**remote-management**).

The **no** form of this command causes the router to be inherited from the global context (**config**>**system**>**management-interface**>**remote-management**).

Default

management

Parameters

router-instance

Specifies the router instance on which the remote management connection is established for this manager, up to 32 characters.

service *vprn-service-instance*

Specifies a VPRN service instance, up to 64 characters.

Platforms

7705 SAR Gen 2

router

Syntax

router *router-or-service*

router service-name *service-name*

no router

Context

[Tree] (config>oam-pm>session>ip router)

Full Context

configure oam-pm session ip router

Description

This command numerically references the source context from which the TWAMP Light packet is launched. The **router-instance** *router-instance* configuration, under the same context as the **router** command, is the preferred method for referencing. This method references the launch context by name, and not number, or alias that converts **service-name** to a number.

The **no** form of this command restores the default value.

Parameters

router-or-service

Specifies the numerical reference to the router instance or service. Well known router-name "Base" is allowed for convenience, but mapped numerically.

Values {*router-name* | *vprn-svc-id*}

router-name: Base

vprn-svc-id: 1 to 2147483647

The parameter *router-instance* is preferred for specifying the router or service.

service-name

Specifies the alias function that allows the service-name to be used converted and stored as service ID, up to 64 characters. The parameter *router-instance* is preferred for specifying the router or service.

Platforms

7705 SAR Gen 2

router

Syntax

router *router-instance*

router service-name *service-name*

no router

Context

[Tree] (config>filter>redirect-policy router)

Full Context

configure filter redirect-policy router

Description

This command enhances VRF support in redirect policies. When a router instance is specified, the configured destination tests are run in the specified router instance, and the PBR action is executed in the specified router instance. If no destination is active or if the hardware does not support PBR action "next-hop router", action forward will be executed (i.e. routing will be performed in the context of the incoming interface routing instance).

The **no** form of the command preserves backward-compatibility. Tests always run in the "Base" routing instance context, and the PBR action executes in the routing context of the ingress interface that the filter using this redirect policy is deployed on.

Default

no router

Parameters

router-instance

Specifies a router instance in the form of **router-name** or **service-id**.

Values **router-name** — Base

service-id — Specifies an existing Layer 3 service [1 to 2147483647]

service-name

Specifies the name of a configured Layer 3 service.

Platforms

7705 SAR Gen 2

router

Syntax

router {**eq** | **neq**} *router-instance* [**regexp**]

no router

Context

[Tree] (config>log>filter>entry>match router)

Full Context

configure log filter entry match router

Description

This command specifies the log event matches for the router instance using a special vrtr-name format used by the logging system.

The **no** form of this command removes the log event matches.

Parameters

eq

Determines if the matching criteria should be equal to the specified value.

neq

Determines if the matching criteria should not be equal to the specified value.

router-instance

Specifies a router name, up to 32 characters, to be used in the match criteria. The router-instance in this command is a name for a router instance in a special format used in the logging system (called the vrtr-name). Examples of vrtr-names include **Base** and **vprn101** (where 101 is the service-id of the VPRN service). It represents the router instance that generated the log event.

regexp

Specifies the type of string comparison to use to determine if the log event matches the value of the specified router instance. When the **regexp** keyword is specified, the string in the **router** command is a regular expression string that is matched against the vrtr-name string in the log event being filtered.

Platforms

7705 SAR Gen 2

router

Syntax

router service-name *service-name*

router *router-instance*

no router

Context

[Tree] (config>system>security>mgmt-access-filter>ip-filter>entry router)

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter>entry router)

Full Context

configure system security management-access-filter ip-filter entry router

configure system security management-access-filter ipv6-filter entry router

Description

This command configures a router name or service ID to be used as a management access filter match criterion.

The **no** form the command removes the router name or service ID from the match criteria.

Parameters

router-instance

Specifies one of the following parameters for the router instance:

router-name — Specifies a router name or CPM router instance, up to 32 characters to be used in the match criteria.

Values "Base" | "management" | "vpls-management"

Default Base

vprn-svc-id — Specifies a CPM router instance to be used in the match criteria.

Values 1 to 2147483647

service name

Specifies an existing service name, up to 64 characters.

Platforms

7705 SAR Gen 2

router

Syntax

router *router-name*

no router

Context

[Tree] (config>mirror>mirror-dest>pcap router)

Full Context

configure mirror mirror-dest pcap router

Description

This command configures the destination router name for the FTP transmission of the PCAP file.

The **no** form of this command configures the router name to **management**, which is the default.

Default

router management

Parameters

router-name

Specifies the router name.

Values Base, management

Platforms

7705 SAR Gen 2

26.88 router-advertisement

router-advertisement

Syntax

[**no**] **router-advertisement**

Context

[Tree] (config>service>vprn router-advertisement)

Full Context

configure service vprn router-advertisement

Description

This command configures router advertisement properties. By default, it is disabled for all IPv6 enabled interfaces.

The **no** form of this command disables all IPv6 interface. However, the **no interface** *interface-name* command disables a specific interface.

Default

no router-advertisement

Platforms

7705 SAR Gen 2

router-advertisement

Syntax

[**no**] **router-advertisement**

Context

[Tree] (config>router router-advertisement)

Full Context

configure router router-advertisement

Description

This command configures router advertisement properties. By default, it is disabled for all IPv6 enabled interfaces.

The **no** form of this command disables all IPv6 interface. However, the **no interface** *interface-name* command disables a specific interface.

Default

disabled

Platforms

7705 SAR Gen 2

26.89 router-id

router-id

Syntax

router-id *ip-address*

no router-id

Context

[Tree] (config>service>vprn router-id)

[Tree] (config>service>vprn>ospf router-id)

[Tree] (config>service>vprn>bgp router-id)

Full Context

configure service vprn router-id

configure service vprn ospf router-id

configure service vprn bgp router-id

Description

This command sets the router ID for a specific VPRN context.

When configuring the router ID in the base instance of OSPF it overrides the router ID configured in the **config>router** context. The default value for the base instance is inherited from the configuration in the **config>router** context. If the router ID in the **config>router** context is not configured, the following applies:

- The system uses the system interface address (which is also the loopback address).
- If a system interface address is not configured, use the last 32 bits of the chassis MAC address.

If neither the router ID nor system interface are defined, the router ID from the base router context is inherited.

This is a **required** command when configuring multiple instances and the instance being configured is not the base instance.

When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for the instance, or reboot the entire router.

It is possible to configure an SR OS to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address.

The **no** form of this command removes the router ID definition from the given VPRN context.

Default

no router-id

Parameters

ip-address

The IP address must be given in dotted decimal notation.

Platforms

7705 SAR Gen 2

router-id

Syntax

router-id *ip-address*

no router-id

Context

[Tree] (config>service>vprn>isis router-id)

Full Context

configure service vprn isis router-id

Description

This command sets the router ID for a specific VPRN context.

If neither the router ID nor system interface are defined, the router ID from the base router context is inherited.

The **no** form of this command removes the router ID definition from the given VPRN context.

Default

no router-id

Parameters

ip-address

The IP address must be given in dotted decimal notation.

Platforms

7705 SAR Gen 2

router-id

Syntax

[**no**] **router-id** *ip*

Context

[Tree] (config>router>mpls>srlg-database router-id)

Full Context

configure router mpls srlg-database router-id

Description

Commands in this context configure the link members of SRLG groups for a specific router in the network. The user must also use this command to enter the local interface SRLG membership into the user SRLG database. Use by CSPF of all interface SRLG membership information of a specific router ID may be temporarily disabled by shutting down the node. If this occurs, CSPF assumes these interfaces have no SRLG membership association.

The **no** form of this command will delete all interface entries under the router ID.

Parameters

ip-address

Specifies the router ID for this system. This must be the router ID configured under the base router instance, the base OSPF instance or the base IS-IS instance.

Platforms

7705 SAR Gen 2

router-id

Syntax

router-id *ip-address*

no router-id

Context

[Tree] (config>router router-id)

Full Context

configure router router-id

Description

This command configures the router ID for the router instance.

The router ID is used by both OSPF and BGP routing protocols in this instance of the routing table manager. IS-IS uses the router ID value as its system ID.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized, the new router ID is used. This can result in an interim period of time when different protocols use different router IDs.

It is possible to configure SR OS to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for each protocol that uses the router ID, or restart the entire router.

The system uses the system interface address which is also the loopback address. If a system interface address is not configured, use the last 32 bits of the chassis MAC address.

The **no** form of this command removes the configured value and the last 32 bits of the chassis MAC address are used.

Default

no router-id

Parameters

ip-address

Specifies the 32 bit router ID expressed in dotted decimal notation or as a decimal value.

Platforms

7705 SAR Gen 2

router-id

Syntax

router-id *ip-address*

no router-id

Context

[Tree] (config>router>bgp router-id)

Full Context

configure router bgp router-id

Description

This command specifies the router ID to be used with this BGP instance.

Changing the BGP router ID on an active BGP instance causes the BGP instance to restart with the new router ID.

It is possible to configure an SR OS to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address.

When no **router-id** is configured for BGP, the system interface IP address is used.

Default

no router-id

Parameters

ip-address

Specifies the router ID, expressed as any non-zero value in the range 0.0.0.1 to 255.255.255.255 (or when converted to decimal it can have any value in the range 1-4294967295). It is recommended to use the system IPv4 address.

Platforms

7705 SAR Gen 2

router-id

Syntax

router-id *router-id*

no router-id

Context

[Tree] (config>router>isis router-id)

Full Context

configure router isis router-id

Description

This command configures the router ID.

The **no** form of this command deletes the router ID.

Parameters

router-id

The IP address of the router.

Platforms

7705 SAR Gen 2

router-id

Syntax

router-id *ip-address*

no router-id

Context

[Tree] (config>router>ospf router-id)

[Tree] (config>router>ospf3 router-id)

Full Context

configure router ospf router-id

configure router ospf3 router-id

Description

This command configures the router ID for the OSPF instance. This command configures the router ID for the OSPF instance.

When configuring the router ID in the base instance of OSPF it overrides the router ID configured in the **config**>**router** context.

The default value for the base instance is inherited from the configuration in the **config**>**router** context. If the router ID in the **config**>**router** context is not configured, the following applies:

- the system uses the system interface address (which is also the loopback address)
- if a system interface address is not configured, it uses the last 32 bits of the chassis MAC address

This is a **required** command when configuring multiple instances and the instance being configured is not the base instance.

When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for the instance, or reboot the entire router.

It is possible to configure an SR OS to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address.

The **no** form of this command to reverts to the default value.

Platforms

7705 SAR Gen 2

26.90 router-instance

router-instance

Syntax

router-instance *router-instance*

no router-instance

Context

[Tree] (config>oam-pm>session>ip router-instance)

Full Context

configure oam-pm session ip router-instance

Description

This command references the source context from which the TWAMP Light packet is launched by name. The **router-instance** *router-instance* configuration is the preferred method for referencing and references the launch context by name, not number or alias that converts **service-name** to a number.

The **no** form of this command restores the default value.

Parameters

router-instance

Specifies the preferred method for entering a service name. Stored as the service name. Only the service linking function is allowed for both mixed-mode and model-driven configuration modes, up to 64 characters.

Platforms

7705 SAR Gen 2

router-instance

Syntax

router-instance *router-instance*

router-instance service *vprn-service-instance*

no router-instance

Context

[Tree] (config>system>telemetry>destination-group>destination router-instance)

[Tree] (config>system>grpc-tunnel>destination-group>destination router-instance)

Full Context

configure system telemetry destination-group destination router-instance

configure system grpc-tunnel destination-group destination router-instance

Description

This command configures the router instance for the destination group.

The **no** form of this command reverts to the default value.

Default

router-instance management

Parameters

router-instance

Specifies the router instance type, up to 32 characters.

Values management, base

vprn-service-instance

Specifies the VPRN service instance, up to 64 characters.

Platforms

7705 SAR Gen 2

router-instance

Syntax

[**no**] **router-instance** *service-id*

Context

[Tree] (config>router>static-route-entry>leak-dest router-instance)

Full Context

configure router static-route-entry leak-destination router-instance

Description

This command configures the static route leak destination router instance.

When a VPRN service is added to the list of VPRNs that receive a leaked copy of the static route, the static route is leaked into that VPRN if the following conditions are met:

- all configured next hops of the static route are direct next hops

- the static route is an active route, or it is capable of immediately becoming an active route when a more-preferred route for the same prefix is removed

Static routes leaked using this method appear as "VPN Leak" protocol routes in the route table of the VPRN.

When a VPRN receives a packet that matches a "VPN Leak" route that leaked using this method, the packet is forwarded according to the configuration of the static route in the GRT, even if the static route is currently non-best in the GRT.

The **no** form of this command removes the configuration.

Parameters

service-id

Specifies the service ID.

Values 1 to 2147483647 | svc-name: up to 64 characters

Platforms

7705 SAR Gen 2

26.91 router-lifetime

router-lifetime

Syntax

router-lifetime *seconds*

no router-lifetime

Context

[Tree] (config>router>router-advert>if router-lifetime)

[Tree] (config>service>vprn>router-advert>if router-lifetime)

Full Context

configure router router-advertisement interface router-lifetime

configure service vprn router-advertisement interface router-lifetime

Description

This command sets the router lifetime.

Default

router life-time 1800

Parameters

seconds

The length of time, in seconds, (relative to the time the packet is sent) that the prefix is valid for route determination.

Values 0, 4 to 9000 seconds. 0 means that the router is not a default router on this link.

Platforms

7705 SAR Gen 2

26.92 router-unsolicited-na-flood-evpn

router-unsolicited-na-flood-evpn

Syntax

[**no**] **router-unsolicited-na-flood-evpn**

Context

[Tree] (config>service>vpls>proxy-nd router-unsolicited-na-flood-evpn)

Full Context

configure service vpls proxy-nd router-unsolicited-na-flood-evpn

Description

This command controls whether the system floods router unsolicited Neighbor Advertisements to EVPN. The NA messages impacted by this command are NA messages with the following flags: S=0 and R=1.

The **no** form of the command will only flood to local SAPs/binds but not to EVPN destinations. This is only recommended in networks where CEs are routers directly connected to the PEs. Networks using aggregation switches between the host/routers and the PEs should flood unsolicited NA messages in EVPN to ensure that the remote caches are updated and BGP does not miss the advertisement of these entries.

Default

router-unsolicited-na-flood-evpn

Platforms

7705 SAR Gen 2

26.93 routing-type0

routing-type0

Syntax

routing-type0 {**true** | **false**}

no routing-type0

Context

[Tree] (config>filter>ipv6-filter>entry>match routing-type0)

Full Context

configure filter ipv6-filter entry match routing-type0

Description

This command enables match on existence of Routing Type Extension Header type 0 in the IPv6 filter policy.

The **no** form of this command ignores Routing Type Extension Header type 0 presence/absence in a packet when evaluating match criteria of a given filter policy entry.

Default

no routing-type0

Parameters

true

Specifies whether a packet contains Routing Type Extension Header type 0.

false

Specifies whether a packet does not contain Routing Type Extension Header type 0.

Platforms

7705 SAR Gen 2

26.94 rp

rp

Syntax

rp

Context

[Tree] (config>service>vprn>pim rp)

Full Context

configure service vprn pim rp

Description

This command enables access to the context to configure the rendezvous point (RP) of a PIM protocol instance.

A Nokia PIM router acting as an RP must respond to a PIM register message specifying an SSM multicast group address by sending stop register message(s) to the first hop router. It does not build an (S, G) shortest path tree toward the first hop router. An SSM multicast group address can be either from the SSM default range of 232/8 or from a multicast group address range that was explicitly configured for SSM.

Default

rp enabled when PIM is enabled.

Platforms

7705 SAR Gen 2

rp

Syntax

rp

Context

[Tree] (config>router>pim rp)

Full Context

configure router pim rp

Description

Commands in this context configure rendezvous point (RP) parameters. The address of the root of the group's shared multicast distribution tree is known as its RP. Packets received from a source upstream and join messages from downstream routers rendezvous at this router.

If this command is not enabled, then the router can never become the RP.

Platforms

7705 SAR Gen 2

26.95 rp-candidate

rp-candidate

Syntax

rp-candidate

Context

[Tree] (config>service>vprn>pim>rp rp-candidate)

[Tree] (config>service>vprn>pim>rp>ipv6 rp-candidate)

Full Context

configure service vprn pim rp rp-candidate

configure service vprn pim rp ipv6 rp-candidate

Description

Commands in this context configure the candidate rendezvous point (RP) parameters.

Default

enabled when PIM is enabled

Platforms

7705 SAR Gen 2

rp-candidate

Syntax

rp-candidate

Context

[Tree] (config>router>pim>rp rp-candidate)

[Tree] (config>router>pim>rp>ipv6 rp-candidate)

Full Context

configure router pim rp rp-candidate

configure router pim rp ipv6 rp-candidate

Description

Commands in this context configure the Candidate RP parameters.

Routers use a set of available rendezvous points distributed in Bootstrap messages to get the proper group-to-RP mapping. A set of routers within a domain are also configured as candidate RPs (C-RPs); typically, these will be the same routers that are configured as candidate BSRs.

Every multicast group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The rendezvous point (RP) is the root of this shared tree.

Default

rp-candidate shutdown

Platforms

7705 SAR Gen 2

26.96 rp-set-peer

rp-set-peer

Syntax

[**no**] **rp-set-peer** *ip-address*

Context

[Tree] (config>service>vprn>pim>rp>anycast rp-set-peer)

Full Context

configure service vprn pim rp anycast rp-set-peer

Description

This command configures a peer in the anycast RP-set. The address identifies the address used by the other node as the RP candidate address for the same multicast group address range as configured on this node.

This is a manual procedure. Caution should be taken to produce a consistent configuration of an RP-set for a given multicast group address range. The priority should be identical on each node and be a higher value than any other configured RP candidate that is not a member of this RP-set.

Although there is no set maximum of addresses that can be configured in an RP-set, up to 15 multicast addresses is recommended.

The **no** form of this command removes an entry from the list.

Parameters

ip-address

Specifies the address used by the other node as the RP candidate address for the same multicast group address range as configured on this node.

Platforms

7705 SAR Gen 2

rp-set-peer

Syntax

[**no**] **rp-set-peer** *ipv6-address*

Context

[Tree] (config>service>vprn>pim>rp>ipv6>anycast rp-set-peer)

Full Context

configure service vprn pim rp ipv6 anycast rp-set-peer

Description

This command configures an IPv6 peer in the anycast rp-set. The address identifies the address used by the other node as the RP candidacy address for the same multicast group address range as configured on this node.

This is a manual procedure. Caution should be taken to produce a consistent configuration of an RP- set for a given multicast group address range. The priority should be identical on each node and be a higher value than any other configured RP candidate that is not a member of this rp-set.

Although there is no set maximum of addresses that can be configured in an rp-set, up to 15 multicast addresses is recommended.

The **no** form of this command removes an entry from the list.

Parameters

ipv6-address

Specifies the address used by the other node as the RP candidacy address for the same multicast group address range as configured on this node.

Values

ipv6-address	: x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x [0 to FFFF]H
	d [0 to 255]D

Platforms

7705 SAR Gen 2

rp-set-peer

Syntax

[**no**] **rp-set-peer** *ip-address*

Context

[Tree] (config>router>pim>rp>anycast rp-set-peer)

Full Context

configure router pim rp anycast rp-set-peer

Description

This command configures an IP peer in the anycast RP-set. The address identifies the address used by the other node as the RP candidate address for the same multicast group address range as configured on this node.

This is a manual procedure. Caution should be taken to produce a consistent configuration of an RP-set for a given multicast group address range. The priority should be identical on each node and be a higher value than any other configured RP candidate that is not a member of this RP-set.

Although there is no set maximum number of addresses that can be configured in an RP-set, up to 15 IP addresses is recommended.

The **no** form of this command removes an entry from the list.

Parameters

ip-address

Specifies an IP peer in the anycast RP-set.

Platforms

7705 SAR Gen 2

rp-set-peer

Syntax

[**no**] **rp-set-peer** *ipv6-address*

Context

[Tree] (config>router>pim>rp>ipv6>anycast rp-set-peer)

Full Context

configure router pim rp ipv6 anycast rp-set-peer

Description

This command configures a peer in the anycast RP-set. The address identifies the address used by the other node as the RP candidate address for the same multicast group address range as configured on this node.

This is a manual procedure. Caution should be taken to produce a consistent configuration of an RP-set for a given multicast group address range. The priority should be identical on each node and be a higher value than any other configured RP candidate that is not a member of this RP-set.

Although there is no set maximum number of addresses that can be configured in an RP-set, up to 15 IP addresses is recommended.

The **no** form of this command removes the IPv6 address from the anycast RP set.

Parameters

ipv6-address

Specifies an IPv6 peer in the anycast RP-set.

Platforms

7705 SAR Gen 2

26.97 rpc-authorization

rpc-authorization

Syntax

rpc-authorization

Context

[Tree] (config>system>security>profile>grpc rpc-authorization)

Full Context

configure system security profile grpc rpc-authorization

Description

This command opens a configuration context for configuring user privileges related to RPCs.

Platforms

7705 SAR Gen 2

26.98 rpf-table

rpf-table

Syntax

rpf-table {**rtable-m** | **rtable-u** | **both**}

no rpf-table

Context

[Tree] (config>service>vprn>pim rpf-table)

Full Context

configure service vprn pim rpf-table

Description

This command configures the sequence of route tables used to find an RPF interface for a multicast route.

By default, only the unicast route table is looked up to calculate RPF interface towards the source or rendezvous point. The user can specify the following options:

- use the unicast route table only
- use the multicast route table only
- use both route tables

The **no** form of this command configures the router to only use the unicast route table.

Default

no rpf-table

Parameters

rtable-m

Keyword to specify that only the multicast route table is used by the multicast protocol (PIM) for IPv4 RPF checks. This route table contains routes submitted by static routes, IS-IS and OSPF.

rtable-u

Keyword to specify that only the unicast route table is used by the PIM for IPv4 RPF checks. This route table contains routes submitted by all the unicast routing protocols.

both

Keyword to specify that the multicast route table is used first by the PIM for checks. If the multicast route table lookup fails, the unicast route table is used.

Platforms

7705 SAR Gen 2

rpf-table

Syntax

rpf-table {**rtable-m** | **rtable-u** | **both**}

no rpf-table

Context

[Tree] (config>router>pim rpf-table)

Full Context

configure router pim rpf-table

Description

This command configures the sequence of route tables used to find a Reverse Path Forwarding (RPF) interface for a particular multicast route.

By default, only the unicast route table is looked up to calculate RPF interface towards the source or rendezvous point. However, the operator can specify one of the following:

- use the unicast route table only
- use the multicast route table only
- use both the route tables

The **no** form of this command reverts to the default value.

Default

rpf-table rtable-u

Parameters

rtable-m

Specifies that only the multicast route table will be used by the multicast protocol (PIM) for IPv4 RPF checks. This route table will contain routes submitted by static routes, ISIS and OSPF.

rtable-u

Specifies only that the unicast route table will be used by the multicast protocol (PIM) for IPv4 RPF checks. This route table will contain routes submitted by all the unicast routing protocols.

both

Specifies to always lookup first in the multicast route table and if there is a route, it will use it. If PIM does not find a route in the first lookup, it will try to find it in the unicast route table. Rtable-m is checked before rtable-u.

Platforms

7705 SAR Gen 2

26.99 rpf6-table

rpf6-table

Syntax

rpf6-table {**rtable6-m** | **rtable6-u** | **both**}

no rpf6-table

Context

[Tree] (config>service>vprn>pim rpf6-table)

Full Context

configure service vprn pim rpf6-table

Description

This command configures the sequence of route tables used to find a Reverse Path Forwarding (RPF) interface for a specific multicast route.

By default, only the unicast route table is looked up to calculate the RPF interface toward the source/rendezvous point. However, the operator can specify to use the following:

- unicast route table only
- multicast route table only
- both route tables

Default

rpf6-table rtable6-u

Parameters

rtable6-m

Specifies that only the multicast route table will be used by the multicast protocol (PIM) for IPv6 RPF checks. This route table will contain routes submitted by static routes, ISIS and OSPF.

rtable6-u

Specifies that only the unicast route table will be used by the multicast protocol (PIM) for IPv6 RPF checks. This route table will contain routes submitted by all unicast routing protocols.

both

Specifies that the multicast route table will be used first by the multicast protocol (PIM) for IPv6 RPF checks, then the unicast route table will be used if the multicast route table lookup fails.

Platforms

7705 SAR Gen 2

rpf6-table

Syntax

rpf6-table {**rtable6-m** | **rtable6-u** | **both**}

no rpf6-table

Context

[Tree] (config>router>pim rpf6-table)

Full Context

configure router pim rpf6-table

Description

This command configures the sequence of route tables used to find a Reverse Path Forwarding (RPF) interface for a particular multicast route.

By default, only the unicast route table is looked up to calculate RPF interface towards the source/rendezvous point. However, the operator can specify the following:

- use unicast route table only
- use multicast route table only or
- use both the route tables

The **no** form of this command reverts to the default value.

Default

rpf6-table rtable6-u

Parameters

rtable6-m

Specifies that only the multicast route table will be used by the multicast protocol (PIM) for IPv6 RPF checks. This route table will contain routes submitted by static routes, ISIS and OSPF.

rtable6-u

Specifies that only the unicast route table will be used by the multicast protocol (PIM) for IPv6 RPF checks. This route table will contain routes submitted by all the unicast routing protocols.

both

Specifies that the multicast route table will be used first by the multicast protocol (PIM) for IPv6 RPF checks, and then the unicast route table will be used if the multicast route table lookup fails.

Platforms

7705 SAR Gen 2

26.100 rpfv

rpfv

Syntax

rpfv [**detail**]

no rpfv

Context

[Tree] (debug>router>pim rpfv)

Full Context

debug router pim rpfv

Description

This command enables debugging for PIM RPF vector.

The **no** form of this command disables debugging for PIM RPF vector.

Parameters

detail

Debugs detailed RPF vector information.

Platforms

7705 SAR Gen 2

rpfv

Syntax

rpfv core

rpfv mvpn

rpfv core mvpn

no rpfv [**core**] [**mvpn**]

Context

[Tree] (config>router>pim rpfv)

Full Context

configure router pim rpfv

Description

This command enables RPF Vector processing for Inter-AS Rosen MVPN Option-B and Option-C. The **rpfv** must be enabled on every node for Inter-AS Option B/C MVPN support.

If **rpfv** is configured, MLDP inter-AS resolution cannot be used. These two features are mutually exclusive.

The **no** form of this command reverts to the default.

Default

no rpfv

Parameters

mvpn

Enables MVPN RPF vector processing for Inter-AS Option B/C MVPN based on RFC 5496 and RFC 6513. If a core RPF vector is received, it will be dropped before a message is processed.

core

Enables core RPF vector (no RD) processing for Inter-AS Option B/C MVPN, which allows SR OS interoperability as P-router with third-party vendors that do not encode RD in the RPF vector for Inter-AS MVPN.

core mvpn

Enables core RPF vector (no RD) processing for Inter-AS Option B/C MVPN, which allows SR OS interoperability as P-router with third-party vendors that do not encode RD in the RPF vector for Inter-AS MVPN.

The **no** version of this command disables RPF Vector processing. If RPF vector is received in a PIM join message, the vector will be removed before local processing of PIM message starts.

Platforms

7705 SAR Gen 2

26.101 rpki-session

rpki-session

Syntax

[**no**] **rpki-session** *ip-address*

Context

[Tree] (config>router>origin-validation rpki-session)

Full Context

configure router origin-validation rpki-session

Description

This command configures a session with an RPKI local cache server by using the RPKI-Router protocol. It is over these sessions that the router learns dynamic VRP entries expressing valid origin AS and prefix associations. SR OS supports the RPKI-Router protocol over TCP/IPv4 or TCP/IPv6 transport. The router can set up an RPKI-Router session using the base routing table (in-band) or the management router (out-of-band). Configure the command in the **config**>**router management** instance to configure a session using the management port.

Default

no rpki-session

Parameters

ip-address

Specifies the IPv4 address or an IPv6 address. If the IPv6 address is link-local then the interface name must be appended to the IPv6 address after a hyphen (-).

Platforms

7705 SAR Gen 2

rpki-session

Syntax

[**no**] **rpki-session** *ip-address*

Context

[Tree] (debug>router rpki-session)

Full Context

debug router rpki-session

Description

This command enables and configures debugging for RPKI session.

The **no** form of this command disables debugging for RPKI session.

Parameters

ip-address

Debugs the RPKI session associated with the specified IP address.

Values

ipv4-address:	a.b.c.d	
ipv6-address	x:x:x:x:x:x:x:x	[-interface]

x:x:x:x:x:x:d.d.d.d	[-interface]
x:	[0 to FFFF]H
d:	[0 to 255]D
interface	up to 32 characters, mandatory for link local addresses

Platforms

7705 SAR Gen 2

26.102 rr

rr

Syntax

[**no**] **rr**

Context

[Tree] (debug>router>rsvp>event rr)

Full Context

debug router rsvp event rr

Description

This command debugs refresh reduction events.

The **no** form of the command disables the debugging.

Platforms

7705 SAR Gen 2

26.103 rr-use-route-table

rr-use-route-table

Syntax

rr-use-route-table

no rr-use-route-table

Context

[Tree] (config>router>bgp>next-hop-res>labeled-routes rr-use-route-table)

Full Context

configure router bgp next-hop-resolution labeled-routes rr-use-route-table

Description

This command enables BGP to perform a lookup of IGP routes in the route table to resolve the BGP next-hop of label-IPv4 and label-IPv6 routes. This is useful for a Route Reflector (RR) that does not participate in tunnel signaling protocols such as LDP and RSVP and therefore, does not have tunnels to resolve the BGP next-hops of label-unicast routes.

Configure the **disable-route-table-install** command before you configure the **rr-use-route-table** command because forwarding would otherwise be incorrect for cases where label routes are resolved this way.

Default

no rr-use-route-table

Platforms

7705 SAR Gen 2

26.104 rs-fec-mode

rs-fec-mode

Syntax

rs-fec-mode *rs-fec-mode*

no rs-fec-mode

Context

[Tree] (config>port>connector rs-fec-mode)

Full Context

configure port connector rs-fec-mode

Description

This command is used for breakout connectors when all connector ports must use the same **rs-fec-mode** setting.

In all other cases, the **rs-fec-mode** is set using the **configure port ethernet rs-fec-mode** command for each individual connector port.

See "Forward Error Correction" in the *Interface Configuration Guide* for more information about **rs-fec-mode** settings.

Default

no rs-fec-mode

Parameters

rs-fec-mode

Specifies the RS-FEC mode to support.

Values cl91-514-528, cl91-514-544

Platforms

7705 SAR Gen 2

rs-fec-mode

Syntax

rs-fec-mode *rs-fec-mode*

no rs-fec-mode

Context

[Tree] (config>port>ethernet rs-fec-mode)

Full Context

configure port ethernet rs-fec-mode

Description

This command enables RS-FEC on the Ethernet port. RS-FEC Clause 91 is required for QSFP28, CFP4, 100GBase-SR4, 100GBase-ER4 lite, and CWDM4 for the QSFP28 package optics for short-reach optics.

See "Forward Error Correction" in the *Interface Configuration Guide* for more information about **rs-fec-mode** settings.

Default

no rs-fec-mode

Parameters

rs-fec-mode

Specifies the RS-FEC mode to support.

Values cl91-514-528, cl74, cl108

Platforms
7705 SAR Gen 2

26.105 rsa

rsa

Syntax
rsa

Context
[Tree] (config>system>security>user>public-keys rsa)

Full Context
configure system security user public-keys rsa

Description
This command allows the user to enter the context to configure RSA public keys.

Platforms
7705 SAR Gen 2

26.106 rsa-key

rsa-key

Syntax
rsa-key *key-id* [**create**]
no rsa-key *key-id*

Context
[Tree] (config>system>security>user>public-keys>rsa rsa-key)

Full Context
configure system security user public-keys rsa rsa-key

Description
This command creates an RSA public key and associates it with the username. Multiple public keys can be associated with the user. The key ID is used to identify these keys for the user.

Parameters

create

Keyword used to create the RSA key. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

key-id

Specifies the key identifier.

Values 1 to 32

Platforms

7705 SAR Gen 2

26.107 rsa-signature

rsa-signature

Syntax

rsa-signature {**pkcs1** | **pss**}

Context

[Tree] (config>ipsec>cert-profile>entry rsa-signature)

Full Context

configure ipsec cert-profile entry rsa-signature

Description

This command specifies the signature scheme for RSA key.

Default

rsa-signature pkcs1

Parameters

pkcs1

Specifies the RSA pkcs#1 v1.5 signature scheme.

pss

Specifies the RSA probabilistic signature scheme.

Platforms

7705 SAR Gen 2

26.108 rsvp

rsvp

Syntax

[**no**] **rsvp**

Context

[Tree] (config>service>vprn>auto-bind-tunnel>res-filter rsvp)

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel>res-filter rsvp)

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>res-filter rsvp)

Full Context

configure service vprn auto-bind-tunnel resolution-filter rsvp

configure service vprn bgp-evpn mpls auto-bind-tunnel resolution-filter rsvp

configure service epipe bgp-evpn mpls auto-bind-tunnel resolution-filter rsvp

Description

This command selects the RSVP-TE tunnel type.

The **rsvp** value instructs BGP to search for the best metric RSVP LSP to the address of the BGP next hop. This address can correspond to the system interface or to another loopback interface used by the BGP instance on the remote node. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest **tunnel-id**.

The **no** form of this command removes the RSVP-TE tunnel type.

Default

no rsvp

Platforms

7705 SAR Gen 2

rsvp

Syntax

rsvp

Context

[Tree] (config>service>vprn>auto-bind-tunnel>res-filter rsvp)

Full Context

configure service vprn auto-bind-tunnel resolution-filter rsvp

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

rsvp

Syntax

[**no**] **rsvp**

Context

[Tree] (config>router rsvp)

Full Context

configure router rsvp

Description

Commands in this context configure RSVP protocol parameters. RSVP is not enabled by default and must be explicitly enabled (**no shutdown**).

RSVP is used to set up LSPs. RSVP should be enabled on all router interfaces that participate in signaled LSPs.

The **no** form of this command deletes this RSVP protocol instance and removes all configuration parameters for this RSVP instance. To suspend the execution and maintain the existing configuration, use the **shutdown** command. RSVP must be shutdown before the RSVP instance can be deleted. If RSVP is not shutdown, the **no rsvp** command does nothing except issue a warning message on the console indicating that RSVP is still administratively enabled.

Default

no shutdown

Platforms

7705 SAR Gen 2

rsvp

Syntax

rsvp [**lsp** *lsp-name*] [**sender** *source-address*] [**endpoint** *endpoint-address*] [**tunnel-id** *tunnel-id*] [**lsp-id** *lsp-id*] [**interface** *ip-int-name*]

no rsvp

Context

[Tree] (debug>router rsvp)

Full Context

debug router rsvp

Description

This command enables and configures debugging for RSVP.

Parameters

lsp *lsp-name*

Specifies the LSP name up to 64 characters in length.

sender *source-address*

Specifies the IP address of the sender.

endpoint *endpoint-address*

Specifies the far-end IP address.

tunnel-id *tunnel-id*

Specifies the RSVP tunnel ID.

Values 0 to 4294967295

lsp-id *lsp-id*

Specifies the LSP ID.

Values 1 to 65535

interface *ip-int-name*

Specifies the interface name. The interface name can be up to 32 characters long and must be unique. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

rsvp

Syntax

[**no**] **rsvp**

Context

[Tree] (config>router>bgp>next-hop-res>shortcut-tunn>family>res-filter rsvp)

[Tree] (config>router>bgp>next-hop-res>lbl-routes>transport-tunn>family>res-filter rsvp)

Full Context

configure router bgp next-hop-resolution shortcut-tunnel family resolution-filter rsvp

configure router bgp next-hop-resolution labeled-routes transport-tunnel family resolution-filter rsvp

Description

This command selects RSVP tunneling for next-hop resolution and specifies RSVP tunnels in a tunnel table to IPv4 destinations. This option allows BGP to use the best metric RSVP LSP to the address of the BGP next-hop. This address can correspond to the system interface or to another loopback interface of the remote BGP router. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel ID.

Platforms

7705 SAR Gen 2

rsvp

Syntax

[**no**] **rsvp**

Context

[Tree] (conf>router>isis>igp-sc>tunn-nh>family>res-filter rsvp)

Full Context

configure router isis igp-shortcut tunnel-next-hop family resolution-filter rsvp

Description

This command selects the RSVP-TE tunnel type in the resolution of the IP prefix or SR tunnel family using IGP shortcuts.

Platforms

7705 SAR Gen 2

rsvp

Syntax

[**no**] **rsvp**

Context

[Tree] (config>router>ospf3>igp-sc>tunnel-nh>family>res-filter rsvp)

[Tree] (config>router>ospf>igp-sc>tunnel-nh>family>res-filter rsvp)

Full Context

configure router ospf3 igp-shortcut tunnel-next-hop family resolution-filter rsvp

configure router ospf igp-shortcut tunnel-next-hop family resolution-filter rsvp

Description

This command selects the RSVP-TE tunnel type in the resolution of the IP prefix or SR tunnel family using IGP shortcuts.

Platforms

7705 SAR Gen 2

26.109 rsvp-resv-style

rsvp-resv-style

Syntax

rsvp-resv-style [**se** | **ff**]

Context

[Tree] (config>router>mpls>lsp rsvp-resv-style)

Full Context

configure router mpls lsp rsvp-resv-style

Description

This command specifies the RSVP reservation style, shared explicit (se) or fixed filter (ff). A reservation style is a set of control options that specify a number of supported parameters. The style information is part of the LSP configuration.

Default

rsvp-resv-style se

Parameters

ff

Fixed filter is single reservation with an explicit scope. This reservation style specifies an explicit list of senders and a distinct reservation for each of them. A specific reservation request is created for data packets from a particular sender. The reservation scope is determined by an explicit list of senders.

se

Shared explicit is shared reservation with a limited scope. This reservation style specifies a shared reservation environment with an explicit reservation scope. This reservation style creates a single reservation over a link that is shared by an explicit list of senders. Because each sender is explicitly listed in the RESV message, different labels can be assigned to different sender-receiver pairs, thereby creating separate LSPs.

Platforms

7705 SAR Gen 2

26.110 rsvp-shortcut

rsvp-shortcut

Syntax

rsvp-shortcut [*ip-address*]

no rsvp-shortcut

Context

[Tree] (debug>router>ospf rsvp-shortcut)

Full Context

debug router ospf rsvp-shortcut

Description

This command debugs the OSPFv2 RSVP shortcut.

Parameters

ip-address

Specifies the IP address to debug.

Platforms

7705 SAR Gen 2

26.111 rsvp-te

rsvp-te

Syntax

rsvp-te *value*

no rsvp-te

Context

[Tree] (config>router>mpls>tunnel-table-pref rsvp-te)

Full Context

configure router mpls tunnel-table-pref rsvp-te

Description

This command configures the tunnel table preference for RSVP-TE LSP tunnel type away from its default value.

The tunnel table preference applies to the next-hop resolution of BGP routes of the following families: EVPN, IPv4, IPv6, VPN-IPv4, VPN-IPv6, label-IPv4, and label-IPv6 in the tunnel table.

This feature does not apply to a VPRN, VPLS, or VLL service with explicit binding to an SDP that enabled the **mixed-lsp-mode** option. The tunnel preference in such an SDP is fixed and is controlled by the service manager. The configuration of the tunnel table preference parameter does not modify the behavior of such an SDP and the services that bind to it.

It is recommended to not set two or more tunnel types to the same preference value. In such a situation, the tunnel table prefers the tunnel type which was first introduced in SR OS implementation historically.

The **no** form of this command reverts to the default.

Default

rsvp-te 7

Parameters

value

Specifies the tunnel table preference value for RSVP-TE LSP.

Values 1 to 255

Default 7

Platforms

7705 SAR Gen 2

rsvp-te

Syntax

[**no**] **rsvp-te**

Context

[Tree] (config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter rsvp-te)

Full Context

configure router static-route-entry indirect tunnel-next-hop resolution-filter rsvp-te

Description

This command enables the use of RSVP-TE sourced tunnel entries in the TTM to resolve the associated static route next-hop.

The rsvp-te value instructs the code to search for the set of lowest metric RSVP-TE LSPs to the address of the indirect next-hop. The LSP metric is provided by MPLS in the tunnel table. The static route treats a set of RSVP-TE LSPs with the same lowest metric as an ECMP set. The user has the option of configuring a list of RSVP-TE LSP names to be used exclusively instead of searching in the tunnel table. In that case, all LSPs must have the same LSP metric in order for the static route to use them as an ECMP set. Otherwise, only the LSPs with the lowest common metric value will be selected.

A P2P auto-lsp that is instantiated via an LSP template can be selected in TTM when resolution is set to any. However, Nokia does not recommend configuring an auto-lsp name explicitly under the rsvp-te node as the auto-generated name can change if the node reboots, which will blackhole the traffic of the static route.

Default

no rsvp-te

Platforms

7705 SAR Gen 2

rsvp-te

Syntax

[**no**] **rsvp-te**

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls rsvp-te)

Full Context

configure oam-pm session ip tunnel mpls rsvp-te

Description

This command configures the specification of RSVP-TE specific tunnel information that is used to transport the test packets. Entering this context removes all other tunnel type options configured under the **configure oam-pm session ip tunnel mpls** context. Only a single **mpls** type can be configured for an OAM-PM session.

The **no** form of this command deletes the context and all configurations under it.

Platforms

7705 SAR Gen 2

26.112 rsvp-te-auto

rsvp-te-auto

Syntax

rsvp-te-auto

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls rsvp-te-auto)

Full Context

configure oam-pm session ip tunnel mpls rsvp-te-auto

Description

This command configures the specification of the RSVP-TE Auto (RSVP-TE with dynamically-created LSPs) tunnel information that is used to transport the test packets. Entering this context removes all other tunnel type options configured under the configure oam-pm session ip tunnel mpls context. Only a single mpls type can be configured for an OAM-PM session.

The **no** form of this command deletes the context and all configurations under it.

Platforms

7705 SAR Gen 2

26.113 rtm

rtm

Syntax

rtm [**detail**]

no rtm

Context

[Tree] (debug>router>pim rtm)

Full Context

debug router pim rtm

Description

This command enables debugging for PIM RTM.

The **no** form of this command disables debugging for PIM RTM.

Parameters

detail

Displays detailed RTM information.

Platforms

7705 SAR Gen 2

rtm

Syntax

rtm [**neighbor** *ip-address* | **group** *name*]

no rtm

Context

[Tree] (debug>router>bgp rtm)

Full Context

debug router bgp rtm

Description

This command logs RTM changes in the debug log.

The **no** form of this command disables debugging.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x [-interface] (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H
- d: [0 to 255]D
- interface: up to 32 characters for link local addresses

group *name*

Debugs only events affecting the specified peer group name, up to 64 characters, and associated neighbors.

Platforms

7705 SAR Gen 2

rtm

Syntax

rtm [*ip-address*]

no rtm

Context

[Tree] (debug>router>isis rtm)

Full Context

debug router isis rtm

Description

This command enables debugging for IS-IS route table manager (RTM).

The **no** form of the command disables debugging.

Parameters

ip-address

The specified IP address.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

rtm

Syntax

rtm [*ip-address*]

no rtm

Context

[Tree] (debug>router>ospf rtm)

[Tree] (debug>router>ospf3 rtm)

Full Context

debug router ospf rtm

debug router ospf3 rtm

Description

This command enables debugging for OSPF RTM.

Parameters

ip-address

Specifies the IP address to debug.

Values ipv4-address:

- a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

26.114 rtr-adv-lsa-limit

rtr-adv-lsa-limit

Syntax

rtr-adv-lsa-limit [*1..4294967295*] [**log-only**] [**threshold** *percent*]

rtr-adv-lsa-limit [*1..4294967295*] [**log-only**] [**threshold** *percent*] **overload-timeout forever**

rtr-adv-lsa-limit [*1..4294967295*] [**log-only**] [**threshold** *percent*] **overload-timeout** *seconds*

no rtr-adv-lsa-limit

Context

[Tree] (config>service>vprn>ospf rtr-adv-lsa-limit)

Full Context

configure service vprn ospf rtr-adv-lsa-limit

Description

This command configures the maximum number of LSAs OSPF can learn from another router, in order to protect the system from a router that accidentally advertises a large number of LSAs. When the number of advertised LSAs reaches the configured percentage of this limit, an SNMP trap is sent. If the limit is exceeded, OSPF goes into overload.

The **overload-timeout** option allows the administrator to control how long OSPF is in overload as a result of the advertised LSA limit being reached. At the end of this duration of time the system automatically attempts to restart OSPF. One possible value for the **overload-timeout** is **forever**, which means OSPF is never restarted automatically and this corresponds to the default behavior when the **overload-timeout** option is not configured.

The **no** form of this command removes the **rtr-adv-lsa-limit**.

Default

rtr-adv-lsa-limit forever

Parameters

log-only

Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, overload is not set.

percent

The threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

seconds

Specifies duration in seconds before restarting OSPF.

Values 1 to 1800

Platforms

7705 SAR Gen 2

rtr-adv-lsa-limit

Syntax

rtr-adv-lsa-limit *limit* [**log-only**] [**threshold** *percent*]

rtr-adv-lsa-limit *limit* [**log-only**] [**threshold** *percent*] [**overload-timeout** {*seconds* | **forever**}]

no rtr-adv-lsa-limit

Context

[Tree] (config>router>ospf rtr-adv-lsa-limit)

[Tree] (config>router>ospf3 rtr-adv-lsa-limit)

Full Context

configure router ospf rtr-adv-lsa-limit

configure router ospf3 rtr-adv-lsa-limit

Description

This command configures the maximum number of LSAs OSPF can learn from another router, in order to protect the system from a router that accidentally advertises a large number of LSAs. When the number of advertised LSAs reaches the configured percentage of this limit, an SNMP trap is sent. If the limit is exceeded, OSPF goes into overload.

The **overload-timeout** option allows the administrator to control how long OSPF is in overload as a result of the advertised LSA limit being reached. At the end of this duration of time, the system automatically exits overload. One possible value for the **overload-timeout** is **forever**, which means OSPF is never exiting overload.

The **no** form of this command removes the **rtr-adv-lsa-limit**.

Default

no rtr-adv-lsa-limit

Parameters

log-only

Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, overload is not set.

percent

Specifies the threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

limit

Specifies the number of LSAs, expressed as a decimal integer, that can be learned.

Values 1 to 4294967295

second

Specifies duration in minutes before restarting OSPF.

Values Values 1 to 1800

forever

Specifies that OSPF is restarted only after the **clear router ospf | ospf3 overload rtr-adv-lsa-limit** command is executed.

Platforms

7705 SAR Gen 2

26.115 rx-los-reaction

rx-los-reaction

Syntax

rx-los-reaction {**squelch**}

no rx-los-reaction

Context

[Tree] (config>port>dwdm>coherent rx-los-reaction)

Full Context

configure port dwdm coherent rx-los-reaction

Description

This command configures the reaction to an RX LOS.

Note:
If **rx-los-reaction squelch** is disabled for some coherent DWDM transceivers, the transceiver only reports local fault alarms when an RX LOS condition occurs; however, the port returns to service faster after the LOS condition is cleared. For these transceivers, if **rx-los-reaction squelch** is enabled, there is better visibility of individual alarms (for example, signal-fail, local fault, and no-am-lock), but the port takes longer to return to service after the LOS condition is cleared.

Parameters

squelch

Specifies to squelch (turn off) the transmit signal on RX LOS.

Platforms

7705 SAR Gen 2

26.116 rx-los-thresh

rx-los-thresh

Syntax

rx-los-thresh *threshold*

Context

[Tree] (config>port>dwdm>coherent rx-los-thresh)

Full Context

configure port dwdm coherent rx-los-thresh

Description

This command configures the average input power LOS threshold.

Default

-23.00

Parameters

threshold

Specifies the RX LOS threshold.

Values -30.00 to -13.00

Platforms

7705 SAR Gen 2

26.117 rx-must-be-encrypted

rx-must-be-encrypted

Syntax

[**no**] **rx-must-be-encrypted**

Context

[Tree] (config>port>ethernet>dot1x>macsec rx-must-be-encrypted)

Full Context

configure port ethernet dot1x macsec rx-must-be-encrypted

Description

When the **rx-must-be-encrypted** option is enabled, all traffic that is not MACsec-secured that is received on the port is dropped.

When the **rx-must-be-encrypted** option is disabled, all arriving traffic, whether MACsec secured or not, will be accepted.

Note:
This command is only available on the NULL port level and does not have per-VLAN granularity.

The **no** form of this command disables the **rx-must-be encrypted** option.

Default

rx-must-be-encrypted

Platforms

7705 SAR Gen 2

27 s Commands – Part I

27.1 s-pmsi

s-pmsi

Syntax

s-pmsi [{*vpnSrcAddr* [*vpnGrpAddr*]} [*mdSrcAddr*]]

no s-pmsi

Context

[Tree] (debug>router>pim s-pmsi)

Full Context

debug router pim s-pmsi

Description

This command enables debugging for PIM selective provider multicast service interface.

The **no** form of this command disables the debugging.

Parameters

vpnSrcAddr

Specifies the VPN source address.

vpnGrpAddr

Specifies the VPN group address.

mdSrcAddr

Specifies the source address of the multicast domain.

Platforms

7705 SAR Gen 2

27.2 sa-mac

sa-mac

Syntax

sa-mac *ieee-address* **da-mac** *ieee-address*

no sa-mac

Context

[Tree] (config>mirror>mirror-dest>sap>egress>ip-mirror sa-mac)

Full Context

configure mirror mirror-dest sap egress ip-mirror sa-mac

Description

This command configures the source and destination MAC addresses for IP mirroring.

The **no** form of this command reverts to the default.

Parameters

sa-mac *ieee-address*

Specifies the source MAC address. Multicast, Broadcast and zeros are not allowed.

da-mac *ieee-address*

Specifies the destination MAC address. Zeros are not allowed.

Platforms

7705 SAR Gen 2

27.3 saa

saa

Syntax

saa

Context

[Tree] (config saa)

Full Context

configure saa

Description

Commands in this context configure the Service Assurance Agent (SAA) tests.

Platforms

7705 SAR Gen 2

saa

Syntax

saa *test-name* [**owner** *test-owner*] {**start** | **stop**} [**no-accounting**]

Context

[Tree] (oam saa)

Full Context

oam saa

Description

This command starts or stops an SAA test that is not configured as continuous.

Parameters

test-name

Specifies the name of the SAA test, up to 32 characters. The test name must already be configured in the **config>saa>test** context.

test-owner

Specifies the owner of an SAA operation, up to 32 characters. If a *test-owner* value is not specified, the default owner is used.

Default "TiMOS CLI"

start

Starts the test. A test cannot be started if the same test is still running.

A test cannot be started if it is in a shut-down state. An error message and log event is generated to indicate a failed attempt to start an SAA test run. A test cannot be started if it is in a continuous state.

stop

Stops a test in progress. A test cannot be stopped if it is not in progress. A log message is generated to indicate that an SAA test run has been aborted. A test cannot be stopped if it is in a continuous state.

no-accounting

Disables the recording results in the accounting policy. When specifying **no-accounting** the MIB record produced at the end of the test is not added to the accounting file. It uses one of the three MIB rows available for the accounting module for collection.

Platforms

7705 SAR Gen 2

27.4 saii-type2

saii-type2

Syntax

saii-type2 *global-id:prefix:ac-id*

no saii-type2

Context

[Tree] (config>service>epipe>spoke-sdp-fec saii-type2)

Full Context

configure service epipe spoke-sdp-fec saii-type2

Description

This command configures the source attachment individual identifier for the spoke-sdp. This is only applicable to FEC129 AII type 2.

Parameters

global-id

A Global ID of this router T-PE. This value must correspond to one of the global_id values configured for a local-prefix under **config>service>pw-routing>local-prefix** context.

Values 1 to 4294967295

prefix

The prefix on this router T-PE that the spoke-sdp SDP is associated with. This value must correspond to one of the prefixes configured under **config>service>pw-routing>local-prefix** context.

Values an IPv4-formatted address a.b.c.d or 1 to 4294967295

ac-id

An unsigned integer representing a locally unique identifier for the spoke SDP.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

27.5 same-recipnonce-for-pollreq

same-recipnonce-for-pollreq

Syntax

[**no**] **same-recipnonce-for-pollreq**

Context

[Tree] (config>system>security>pki>ca-profile>cmpv2 same-recipnonce-for-pollreq)

Full Context

configure system security pki ca-profile cmpv2 same-recipnonce-for-pollreq

Description

This command enables the system to use same recipNonce as the last CMPv2 response for poll request.

The **no** form of this command disables the use of the same recipNonce as the last CMPv2 response for poll request.

Default

no same-recipnonce-for-pollreq

Platforms

7705 SAR Gen 2

27.6 sample-interval

sample-interval

Syntax

sample-interval *interval*

Context

[Tree] (config>system>telemetry>persistent-subscriptions>subscription sample-interval)

Full Context

configure system telemetry persistent-subscriptions subscription sample-interval

Description

This command configures the sample interval for persistent subscription.

This sampling interval only applies when the **mode** command is set to either **target-defined** or **sample**.

Default

sample-interval 10000

Parameters

interval

Specifies the sample interval, in milliseconds.

Values 1000 to 4294967295

Platforms

7705 SAR Gen 2

27.7 sap

sap

Syntax

sap *sap-id* [**split-horizon-group** *group-name*] [**create**] [**capture-sap**] [**eth-ring** *ring-index*]

sap *sap-id* [**split-horizon-group** *group-name*] [**create**] [**capture-sap**] [**eth-ring** *ring-index*] **leaf-ac**

sap *sap-id* [**split-horizon-group** *group-name*] [**create**] [**capture- sap**] [**eth-ring** *ring-index*] **root-leaf-tag leaf-tag** *leaf-tag*

no sap *sap-id*

Context

[Tree] (config>service>vpls sap)

Full Context

configure service vpls sap

Description

This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the 7705 SAR Gen 2. Each SAP must be unique. All SAPs must be explicitly created within a service or on an IP interface.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **configure port** *port-id* **ethernet mode access** command. Channelized TDM ports are always access ports.

If a port is shut down, all SAPs on that port become operationally down. When a service is shut down, SAPs for the service are not displayed as operationally down although all traffic traversing the service is discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP are also deleted. For Internet Ethernet Service (IES), the IP interface must be shut down before the SAP on that interface may be removed.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition.

port-id

Specifies the physical port ID in the *slot/mda/port* format.

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the slot_number/MDA_number/port_number [.channel] format. For example 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

group-name

Specifies the name of the split horizon group to which the SAP belongs.

capture-sap

Specifies a capturing SAP in which triggering packets are sent to the CPM. Non-triggering packets captured by the capture SAP are dropped.

create

Keyword used to create a SAP instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

root-leaf-tag

Specifies a SAP as a root leaf tag SAP. Only SAPs of the form dot1q (for example, 1/1/1:X) or qinq (for example, 1/1/1:X.Y, 1/1/1:X.*) are supported. The default E-Tree SAP type is a root AC, if *root-leaf-tag* **(**or *leaf-ac***)** is not specified at SAP creation. This option is only available when the VPLS is designated as an E-Tree VPLS.

leaf-tag-vid

Specifies to replace the outer SAP-ID for leaf traffic. The leaf tag VID is only significant between peering VPLS but the values must be consistent on each end.

leaf-ac

Specifies a SAP as a leaf access (AC) SAP. The default E-Tree SAP type is root AC if **leaf-ac** (or **root-leaf-tag**) is not specified at SAP creation. This option is only available when the VPLS is designated as an E-Tree VPLS.

Platforms

7705 SAR Gen 2

sap

Syntax

sap *sap-id* [**create**] [**no-endpoint**]

sap *sap-id* [**create**] **endpoint** *endpoint-name*

sap *sap-id* [**create**] [**qtag-normalization**] [[*tag*] | [*s-tag*.*c-tag*]]

no sap *sap-id*

Context

[Tree] (config>service>epipe sap)

Full Context

configure service epipe sap

Description

This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the device. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port. Channelized TDM ports are always access ports.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded.

The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

Ethernet SAPs support null, dot1q, and qinq is supported for all routers.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP will also be deleted. For Internet Enhanced Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.

By default, no SAPs are defined.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP.

port-id

Specifies the physical port ID.

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the slot_number/MDA_number/port_number format. For example 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

port-id	*slot*/*mda*/*port* [.*channel*]		
	eth-sat-id	esat-*id*/*slot*/*port*	
		esat	keyword
		id	1 to 20
	pxc-id	pxc-*id*.*sub-port*	
		pxc	keyword
		id	1 to 64
		sub-port	a, b

endpoint

Adds a SAP endpoint association.

no endpoint

Removes the association of a SAP or a spoke SDP with an explicit endpoint name.

create

Keyword to create a SAP instance. The **create** keyword requirement can be enabled or disabled in the **environment create** context.

qtag-normalization

Keyword to enable Q-tag normalization.

tag

Specifies the value for tag normalization. The tag value is pushed as the S-tag (outer tag) into the frames coming from this SAP and sent to EVPN. On network ingress, the inner and outer VLAN tags are looked up and the frames matching this value and the normalized C-tag value are sent to the associated SAP.

Values 0 to 4094

s-tag

Specifies the value for tag normalization. The tag value is pushed as the S-tag (outer tag) into the frames coming from this SAP and sent to EVPN. On network ingress, the inner and outer VLAN tags are looked up and the frames matching this value and the normalized C-tag value are sent to the associated SAP.

Values 0 to 4094

c-tag

Specifies the value for tag normalization. The tag value is pushed as the C-tag (inner tag) into the frames coming from this SAP and sent to EVPN. On network ingress, the inner and outer VLAN tags are looked up and the frames matching this value and the normalized S-tag value are sent to the associated SAP.

Values 0 to 4094

Platforms

7705 SAR Gen 2

Output

The following output is an example of VLL SAP information.

Output Example

```
*A:test>config>service>epipe 200 name "200" customer 1 info detail
------------------------------------------------------------------
            sap 1/1/c5/1:200.200 create
                no shutdown
            exit
            sap pw-21:200.200 create
                no shutdown
            exit
            no shutdown
        exit
    exit
------------------------------------------------------------------
```

sap

Syntax

sap *sap-id* [**create**]

no sap *sap-id*

Context

[Tree] (config>service>ies>if sap)

[Tree] (config>service>vprn>if sap)

Full Context

configure service ies interface sap

configure service vprn interface sap

Description

This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **configure port** *port-id* **ethernet mode access** command. Channelized TDM ports are always access ports.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

Note:
Configure an IES interface as a loopback interface by issuing the **loopback** command instead of the **sap** *sap-id* command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP are also deleted. For Internet Enhanced Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed. The no form of this command causes the ptp-hw-assist to be disabled.

Default

No SAPs are defined.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition.

port-id

Specifies the physical port ID in the *slot/mda/port* format.

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the slot_number/MDA_number/port_number format. For example 61/2/3 specifies port 3 on MDA 2 in slot 61.

Table 84: Port ID Syntax

null	*port-id* | lag-*id*
dot1q	{port-id | lag-id}:{qtag1 | cp-conn-prof-id
qinq	{*port-id* | lag-*id*}:{*qtag1* | cp-*conn-prof-id*}.{*qtag2* | cp-*conn-prof-id*} cp: keyword

	conn-prof-id: 1 to 8000	
port-id	slot/mda/port [.channel]	
	eth-sat-id	esat-id/slot/port
		esat: keyword
		id: 1 to20
	pxc-id	psc-id.sub-port
		pxc psc-id.sub-port
		pxc: keyword
		id: 1 to 64
		sub-port: a, b
lag-id	lag-id	lag: keyword
		id: 1 to 800
qtag1	0 to 4094	
qtag2	* | null | 0 to 4094	

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

create

Keyword used to create a SAP instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

sap

Syntax

sap *sap-id*

no sap

Context

[Tree] (config>service>vpls>site sap)

Full Context

configure service vpls site sap

Description

This command configures a SAP for the site.

The **no** form of this command removes the SAP ID from the configuration.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition

Platforms

7705 SAR Gen 2

sap

Syntax

sap *sap-id*

no sap

Context

[Tree] (config>service>epipe>site sap)

Full Context

configure service epipe site sap

Description

This command configures a SAP for the site.

The **no** form of this command removes the SAP ID from the configuration.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition.

Platforms

7705 SAR Gen 2

sap

Syntax

[**no**] **sap** *sap-id*

Context

[Tree] (debug>service>id>stp sap)

[Tree] (debug>service>id sap)

[Tree] (debug>service>id>dhcp sap)

Full Context

debug service id stp sap

debug service id sap

debug service id dhcp sap

Description

This command enables STP debugging for a specific SAP.

The **no** form of the command disables debugging.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition.

Platforms

7705 SAR Gen 2

sap

Syntax

sap [**split-horizon-group** *group-name*] [**create**] [**capture-sap**]

no sap *sap-id*

Context

[Tree] (config>service>vpls>mac-move>primary-ports sap)

[Tree] (config>service>vpls>mac-move>secondary-ports sap)

Full Context

configure service vpls mac-move primary-ports sap

configure service vpls mac-move secondary-ports sap

Description

This command declares a specified SAP as a primary (or secondary) VPLS port.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition

Platforms

7705 SAR Gen 2

sap

Syntax

[**no**] **sap** *sap-id*

Context

[Tree] (debug>service>id>igmp-snooping sap)

Full Context

debug service id igmp-snooping sap

Description

This command shows IGMP packets for a specific SAP.

The **no** form of this command disables the debugging for the SAP.

Platforms

7705 SAR Gen 2

sap

Syntax

[**no**] **sap** *sap-id*

Context

[Tree] (debug>service>id>mld sap)

Full Context

debug service id mld-snooping sap

Description

This command shows MLD packets for a specific SAP.

The **no** form of this command disables the debugging for the SAP.

Platforms

7705 SAR Gen 2

sap

Syntax

sap *sap-id* [**create**] [**no-endpoint**]

sap *sap-id* [**create**] **endpoint** *name*

no sap

Context

[Tree] (config>mirror>mirror-dest sap)

Full Context

configure mirror mirror-dest sap

Description

This command creates a service access point (SAP) within a mirror destination service. The SAP is owned by the mirror destination service ID.

The SAP is defined with port and encapsulation parameters to uniquely identify the (mirror) SAP on the interface and within the box. The specified SAP may be defined on an Ethernet access port with a dot1q, null, or q-in-q encapsulation type.

Only one SAP can be created within a **mirror-dest** service ID. If the defined SAP has not been created on any service within the system, the SAP is created and the context of the CLI will change to the newly created SAP. In addition, the port cannot be a member of a multi-link bundle, APS group or IMA bundle.

If the defined SAP exists in the context of another service ID, **mirror-dest** or any other type, an error is generated.

Mirror destination SAPs can be created on Ethernet interfaces that have been defined as an access interface. If the interface is defined as network, the SAP creation returns an error.

When the **no** form of this command is used on a SAP created by a mirror destination service ID, the SAP with the specified port and encapsulation parameters is deleted.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition.

no-endpoint

Removes the association of a SAP or a sdp with an explicit endpoint name.

name

Specifies the name of the endpoint associated with the SAP.

Platforms

7705 SAR Gen 2

sap

Syntax

sap *sap-id* {[**egress**] [**ingress**]}

no sap *sap-id* [**egress**] [**ingress**]

Context

[Tree] (config>mirror>mirror-source sap)

Full Context

configure mirror mirror-source sap

Description

This command enables mirroring of traffic ingressing or egressing a service access port (SAP). A SAP that is defined within a mirror destination cannot be used in a mirror source. The mirror source SAP referenced by the *sap-id* is owned by the service ID of the service in which it was created. The SAP is only referenced in the mirror source name for mirroring purposes. The mirror source association does not need to be removed before deleting the SAP from its service ID. If the SAP is deleted from its service ID, the mirror association is removed from the mirror source.

More than one SAP can be associated within a single **mirror-source**. Each SAP has its own **ingress** and **egress** parameter keywords to define which packets are mirrored to the mirror destination.

The SAP must be valid and properly configured. If the associated SAP does not exist, an error occurs and the command will not execute.

The same SAP cannot be associated with multiple mirror source definitions for ingress packets.

The same SAP cannot be associated with multiple mirror source definitions for egress packets.

If a particular SAP is not associated with a mirror source name, then that SAP will not have mirroring enabled for that mirror source.

Note that the ingress and egress options cannot be supported at the same time on a CEM encap-type SAP. The options must be configured in either the ingress **or** egress contexts.

The **no** form of this command disables mirroring for the specified SAP. All mirroring for that SAP on ingress and egress is terminated. Mirroring of packets on the SAP can continue if more specific mirror criteria is configured. If the **egress** or **ingress** parameter keywords are specified in the **no** command, only the ingress or egress mirroring condition is removed.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition.

egress

Specifies that packets egressing the SAP should be mirrored. Egress packets are mirrored to the mirror destination after egress packet modification.

ingress

Specifies that packets ingressing the SAP should be mirrored. Ingress packets are mirrored to the mirror destination before the ingress packet modification.

Platforms

7705 SAR Gen 2

sap

Syntax

sap *sap-id* {[**egress**] [**ingress**]}

no sap *sap-id* [**egress**] [**ingress**]

Context

[Tree] (debug>mirror-source sap)

Full Context

debug mirror-source sap

Description

This command enables mirroring of traffic ingressing or egressing a service access port (SAP). A SAP that is defined within a mirror destination cannot be used in a mirror source. The mirror source SAP referenced by the *sap-id* is owned by the service ID of the service in which it was created. The SAP is only referenced in the mirror source name for mirroring purposes. The mirror source association does not need to be removed before deleting the SAP from its service ID. If the SAP is deleted from its service ID, the mirror association is removed from the mirror source.

More than one SAP can be associated within a single **mirror-source**. Each SAP has its own **ingress** and **egress** parameter keywords to define which packets are mirrored to the mirror destination.

The SAP must be valid and properly configured. If the associated SAP does not exist, an error occurs and the command does not execute.

The same SAP cannot be associated with multiple mirror source definitions for ingress packets.

The same SAP cannot be associated with multiple mirror source definitions for egress packets.

If a particular SAP is not associated with a mirror source name, then that SAP does not have mirroring enabled for that mirror source.

Note that the ingress and egress options cannot be supported at the same time on a CEM encap-type SAP. The options must be configured in either the ingress **or** egress contexts,

The **no** form of this command disables mirroring for the specified SAP. All mirroring for that SAP on ingress and egress is terminated. Mirroring of packets on the SAP can continue if more specific mirror criteria is configured. If the **egress** or **ingress** parameter keywords are specified in the **no** command, only the ingress or egress mirroring condition is removed.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition.

egress

Specifies that packets egressing the SAP should be mirrored. Egress packets are mirrored to the mirror destination after egress packet modification.

ingress

Specifies that packets ingressing the SAP should be mirrored. Ingress packets are mirrored to the mirror destination before the ingress packet modification.

Platforms

7705 SAR Gen 2

sap

Syntax

sap *sap-id*

no sap

Context

[Tree] (config>service>vpls>proxy-nd>dynamic sap)

[Tree] (config>service>vpls>proxy-arp>dynamic sap)

Full Context

configure service vpls proxy-nd dynamic sap

configure service vpls proxy-arp dynamic sap

Description

This command configures the proxy ARP or ND entry for creation when the ARP or neighbor advertisement (NA) packet for the configured IP address is received on the configured SAP. This command can be configured in combination with the **configure service vpls proxy-arp dynamic mac-list** or **configure service vpls proxy-nd dynamic mac-list** command for the entry. In this case, the MAC of the ARP or NA message and the SAP on which the ARP or NA packet is received are both checked before creating the entry.

The **no** form of this command removes the SAP as the match criterion.

Default

no sap

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition.

Values

null	*port-id* | *lag-id* | *eth-sat-id*	
dot1q	*port-id* | *lag-id* | *pw-id* | *eth-sat-id*:[qtag1 **cp**-*conn-prof-id*]	
qinq	*port-id* | *lag-id* | *pw-id* | *eth-sat-id*:[qtag1 **cp**-*conn-prof-id*].[qtag2 | **cp**-*conn-prof-id*]	
	cp	keyword
	conn-prof-id	1 to 8000
port-id	*slot/mda/port*[.channel]	
eth-tunnel	eth-tunnel-*id*[:*eth-tun-sap-id*]	
	id	1 to 1024
	eth-tun-sap-id	0 to 4094
lag-id	**lag**-*id* | **lag**-*string*	
	lag	keyword
	id	1 to 800
	string	up to 23 characters
pw-id	**pw**-*id*	
	pw	keyword
	id	1 to 32767
qtag1	* | null | 0 to 4094	
qtag2	* | null | 0 to 4094	
tunnel-id	**tunnel**-*id*.*private* | *public*:*tag*	
	tunnel	keyword
	id	1 to 64
	tag	0 to 4094
eth-sat-id	**esat**-*id*/*slot*/*port*	
	esat	keyword
	id	1 to 20

Platforms

7705 SAR Gen 2

27.8 sap-egress

sap-egress

Syntax

sap-egress {*policy-id* | *policy-name*} [**create**] [**name** *name*]

no sap-egress {*policy-id* | *policy-name*}

Context

[Tree] (config>qos sap-egress)

Full Context

configure qos sap-egress

Description

This command is used to create or edit a Service Egress QoS policy. The egress policy defines the SLA for service packets as they egress on the SAP.

Policies are templates that can be applied to multiple services as long as the scope of the policy is template. The queues defined in the policy are not instantiated until a policy is applied to a service.

Sap-egress policies determine queue mappings based on ingress DSCP, IP precedence, dot1p, and IPv4 or IPv6 match criteria. Multiple queues can be created per forwarding class and each queue can have different CIR or PIR parameters.

Egress SAP QoS policies allow the definition of queues and the mapping of forwarding classes to those queues. Each queue needs to have a relative CIR for determining its allocation of QoS resources during periods of congestion. A PIR can also be defined that forces a hard limit on the packets transmitted through the queue. When the forwarding class is mapped to the queue, a DSCP, IP precedence, or dot1p value can optionally be specified.

The sap-egress policy with *policy-id* 1 is the default sap-egress QoS policy and is applied to service egress SAPs when an explicit policy is not specified or removed. The default sap-egress policy cannot be modified or deleted.

By default, all forwarding classes map to queue 1.

Any changes made to an existing policy, using any of the sub-commands, will be applied immediately to all egress SAPs where this policy is applied. For this reason, when many changes are required on a policy, it is highly recommended that the policy be copied to a work area *policy-id*. That work-in-progress policy can be modified until complete, then written over the original policy-id. Use the **config qos copy** command to maintain policies in this manner.

The **no** form of this command deletes the sap-egress policy. A policy cannot be deleted until it is removed from all service SAPs where it is applied. When a sap-egress policy is removed from a SAP, the SAP will revert to the default sap-egress *policy-id* 1.

Parameters

policy-id

The *policy-id* uniquely identifies the policy on the router.

Values 1 to 65535

policy-name

The *policy-name* uniquely identifies the policy.

Values 64 characters maximum.

create

Required parameter when creating a SAP QoS egress policy.

name

Configures an optional policy name which adds a name identifier to a specific policy to then use that policy name in configuration references as well as display and use policy names in show commands throughout the system. This helps the service provider or administrator to identify and manage sap-egress policies within the SR OS platforms.

All sap-egress policies are required to assign a policy ID to initially create a policy. However, either the policy ID or the policy name can be used to identify and reference a specific policy once it is initially created.

If a name is not specified at creation time, then SR OS assigns a string version of the *policy-id* as the name.

Values 64 characters maximum

Platforms

7705 SAR Gen 2

27.9 sap-id

sap-id

Syntax

sap-id *sap-id*

no sap-id

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>host-ident sap-id)

Full Context

configure subscriber-mgmt local-user-db ipoe host host-identification sap-id

Description

This command specifies the SAP ID to match for a host lookup. When the LUDB is accessed using a DHCPv4 server, the SAP-ID is matched against the Nokia vendor-specific sub-option in DHCP Option 82.

Note:
This command is used only when **sap-id** is configured as one of the **match-list** parameters.

The **no** form of this command removes the SAP ID from the configuration.

Parameters

sap-id

Specifies a SAP ID, up to 255 characters.

Platforms

7705 SAR Gen 2

sap-id

Syntax

[**no**] **sap-id**

Context

[Tree] (config>service>vpls>sap>dhcp>option>vendor sap-id)

[Tree] (config>service>vprn>if>dhcp>option>vendor sap-id)

Full Context

configure service vpls sap dhcp option vendor-specific-option sap-id

configure service vprn interface dhcp option vendor-specific-option sap-id

Description

This command enables the sending of the SAP ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the SAP ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

Platforms

7705 SAR Gen 2

27.10 sap-ingress

sap-ingress

Syntax

sap-ingress {*policy-id* | *policy-name*} [**create**] [**name** *name*]

no sap-ingress {*policy-id* | *policy-name*}

Context

[Tree] (config>qos sap-ingress)

Full Context

configure qos sap-ingress

Description

This command is used to create or edit the ingress policy. The ingress policy defines the SLA enforcement that service packets receive as they ingress a SAP. SLA enforcement is accomplished through the definition of queues that have Forwarding Class (FC), Fair Information Rate (FIR), Committed Information Rate (CIR), Peak Information Rate (PIR), Committed Burst Size (CBS), and Maximum Burst Size (MBS) characteristics.

Policies in effect are templates that can be applied to multiple services as long as the **scope** of the policy is template. Queues defined in the policy are not instantiated until they are assigned to at least one forwarding class and a policy is applied to a service SAP.

It is possible that a SAP ingress policy will include the **dscp** map command, the **dot1p** map command, and an IP or MAC match criteria. When multiple matches occur for the traffic, the order of precedence will be used to arrive at the final action. The order of precedence is as follows:

1. 802.1p bits
2. DSCP
3. IP quintuple or MAC headers

The SAP ingress policy with *policy-id* 1 is a system-defined policy applied to services when no other policy is explicitly specified. The system SAP ingress policy cannot be modified or deleted. The default SAP ingress policy defines one unicast and one multipoint queue associated with all forwarding classes, with an FIR of zero, a CIR of zero, and a PIR of line rate.

Any changes made to the existing policy, using any of the sub-commands, are applied immediately to all services where this policy is applied. For this reason, when many changes are required on a policy, it is recommended that the policy be copied to a work area policy ID. That work-in-progress policy can be modified until complete, then written over the original policy-id. Use the **config>qos>copy** command to maintain policies in this manner.

The **no** form of this command deletes the SAP ingress policy. A policy cannot be deleted until it is removed from all services where it is applied.

Parameters

policy-id

The *policy-id* uniquely identifies the policy.

Values 1 to 65535

policy-name

The *policy-name* uniquely identifies the policy.

Values 64 characters maximum

create

Required parameter when creating a SAP QoS ingress policy.

name *name*

Configures an optional policy name which adds a name identifier to a specific policy to then use that policy name in configuration references as well as display and use policy names in show commands throughout the system. This helps the service provider and administrator to identify and manage sap-ingress policies within the SR OS platforms.

All sap-ingress policies are required to assign a policy ID to initially create a policy. However, either the policy ID or the policy name can be used to identify and reference a specific policy after it is initially created.

If a name is not specified at creation time, then SR OS assigns a string version of the *policy-id* as the name.

Values 64 characters

Platforms

7705 SAR Gen 2

27.11 sap-template-binding

sap-template-binding

Syntax

sap-template-binding *name/id*

no sap-template-binding

Context

[Tree] (config>service>vpls>vpls-group sap-template-binding)

Full Context

configure service vpls vpls-group sap-template-binding

Description

This command configures the binding to a SAP template to be used to instantiate SAPs in the data VPLS using as input variables the VLAN IDs generated by the vid-range command.

The **no** form of this command removes the binding and deletes the related SAP instances. The command will fail if any of the affected VPLS instances have either a provisioned SAP or an active MVRP declaration/registration or if the related vpls-group is in no shutdown state. Any changes to the **sap-template-binding** require the **vpls-group** to be in **shutdown** state. New control SAP additions to the management VPLS are allowed as long as data VPLS instantiations/removals for vpls-groups are not in progress. Control SAPs can be removed at any time generating the removal of related data SAPs from the data VPLS. The **shutdown** or **no shutdown** state for the control SAPs does not have any effect on data SAPs instantiated with this command.

Default

no sap-template-binding

Parameters

name

Specifies the name of the VPLS template

Values ASCII character string

id

Specifies the ID of the VPLS template

Values 1 to 8196

Platforms

7705 SAR Gen 2

27.12 save

save

Syntax

save [*cflash-id*]

Context

[Tree] (bof save)

Full Context

bof save

Description

This command uses the boot option parameters currently in memory and writes them from the boot option file to the specified compact flash.

The BOF must be located in the root directory of the internal or external compact flash drives local to the system and have the mandatory filename of *bof.cfg*.

If a location is not specified, the BOF is saved to the default compact flash drive (cf3:) on the active CPM (typically the CPM in slot A, but the CPM in slot B could also be acting as the active CPM). The slot name is not case-sensitive. You can use upper or lowercase "A" or "B".

Command usage:

- **bof save** — saves the BOF to the default drive (cf3:) on the active CPM (either in slot A or B)
- **bof save cf3:** — saves the BOF to cf3: on the active CPM (either in slot A or B)

To save the BOF to a compact flash drive on the standby CPM (for example, the redundant (standby) CPM is installed in slot B), specify -A or -B option.

Command usage:

- **bof save cf3-A:** — saves the BOF to cf3: on CPM in slot A whether it is active or standby
- **bof save cf3-B:** — saves the BOF to cf3: on CPM in slot B whether it is active or standby

The slot name is not case-sensitive. You can use upper or lowercase "A" or "B".

The **bof save** and **show bof** commands allow you to save to or read from the compact flash of the standby CPM. Use the **show card** command to determine the active and standby CPM (A or B).

Default

Saves must be explicitly executed. The BOF is saved to cf3: if a location is not specified.

Parameters

flash-id

Specifies the compact flash ID where the *bof.cfg* is to be saved.

Values cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Default cf3:

Platforms

7705 SAR Gen 2

save

Syntax

save *file-url*

Context

[Tree] (candidate save)

Full Context

candidate save

Description

This command saves the current candidate to a file.

Parameters

file-url

Specifies the directory and filename.

Platforms

7705 SAR Gen 2

save

Syntax

save [**comment** *comment*] [**rescue**]

Context

[Tree] (admin>rollback save)

Full Context

admin rollback save

Description

If the optional **rescue** keyword is not used, this command saves a rollback checkpoint at the location and with the filename specified by the rollback-location with a suffix of .rb. The previously saved checkpoints will have their suffixes incremented by one (.rb.1 becomes .rb.2, and so on). If there are already as many checkpoint files as the maximum number supported, then the last checkpoint file is deleted.

If the **rescue** keyword is used, then this command saves the current operational configuration as a rescue configuration at the location and with the filename specified by the rescue location. The filename will have the suffix .rc appended.

Parameters

comment-string

Specifies a comment, up to 255 characters, that is associated with the checkpoint.

rescue

Saves the rescue checkpoint instead of a normal rollback checkpoint.

Platforms

7705 SAR Gen 2

save

Syntax

save [*file-url*] [**detail**] [**index**]

Context

[Tree] (admin save)

Full Context

admin save

Description

This command saves the running configuration to a configuration file. For example:

```
A:ALA-1>admin# save ftp://test:test@192.168.x.xx/./100.cfg
Saving configuration .........Completed.
```

By default, the running configuration is saved to the primary configuration file.

Parameters

file-url

Specifies the file URL location to save the configuration file.

Values		
	local-url | *remote-url*	
	local-url	[*cflash-id*/][*file-path*] 200 chars max, including cflash-id
		directory length 99 chars max each
	remote-url	[{ftp:// | tftp://}*login:pswd@remote-locn*/][*file-path*]
		243 chars max
		directory length 99 chars max each
	remote-locn	[*hostname* | *ipv4-address* | *ipv6-address*]
	ipv4-address	*a.b.c.d*
	ipv6-address	*x:x:x:x:x:x:x:x*[*-interface*]
		x:x:x:x:x:x:d.d.d.d[*-interface*]
		x - [0 to FFFF]H
		d - [0 to 255]D
		interface - 32 chars max, for link local addresses

cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Default the primary configuration file location

detail

Saves both default and non-default configuration parameters.

index

Forces a save of the persistent index file regardless of the persistent status in the BOF file. The index option can also be used to avoid an additional boot required while changing your system to use the persistence indexes.

Platforms

7705 SAR Gen 2

27.13 save-when-restricted

save-when-restricted

Syntax

[no] save-when-restricted

Context

[Tree] (config>system>security>user save-when-restricted)

[Tree] (config>system>security>user-template save-when-restricted)

Full Context

configure system security user save-when-restricted

configure system security user-template save-when-restricted

Description

This command specifies whether the system permits configuration save operations for all configuration regions (bof, debug, configure, li) via any management interface (such as CLI and NETCONF) even if **restricted-to-home** is enabled.

The configuration for a region can be saved with CLI commands such as **bof save**, **admin debug-save**, **admin save**, or **configure li save**.

The **no** form of this command denies saving the configuration when **restricted-to-home** is enabled.

Default

save-when-restricted

Platforms

7705 SAR Gen 2

27.14 saved-ind-prompt

saved-ind-prompt

Syntax

[**no**] **saved-ind-prompt**

Context

[Tree] (environment saved-ind-prompt)

Full Context

environment saved-ind-prompt

Description

This command enables saved indicator in the prompt. When changes are made to the configuration file a "*" appears in the prompt string indicating that the changes have not been saved. When an **admin save** command is executed the "*" disappears.

```
*A:ALA-48# admin save
Writing file to ftp://192.0.2.43/./sim48/sim48-config.cfg
Saving configuration .... Completed.
A:ALA-48#
```

Platforms

7705 SAR Gen 2

27.15 schedule

schedule

Syntax

[**no**] **schedule** *schedule-name* [**owner** *schedule-owner*]

Context

[Tree] (config>system>cron schedule)

Full Context

configure system cron schedule

Description

This command configures the type of schedule to run, including one-time only (oneshot), periodic or calendar-based runs. All runs are determined by month, day of month or weekday, hour, minute and interval (seconds).

The **no** form of the command removes the context from the configuration.

Parameters

schedule-name

Specifies the name of the schedule. The name can be up to 32 characters.

schedule-owner

Specifies the owner name of the schedule. The name can be up to 32 characters.

Default TiMOS CLI

Platforms

7705 SAR Gen 2

27.16 schedule-type

schedule-type

Syntax

schedule-type *schedule-type*

Context

[Tree] (config>system>security>pki>ca-prof>auto-crl-update schedule-type)

Full Context

configure system security pki ca-profile auto-crl-update schedule-type

Description

This command specifies the schedule type for auto CRL update. The system supports two types:

- **periodic**: — The system will download a CRL periodically at the interval configured via the **periodic-update-interval** command. For example, if the periodic-update-interval is 1 day, then the system will download a CRL every 1 day. The minimal periodic-update-interval is 1 hour.
- **next-update-based** — The system will download a CRL at the time = Next_Update_of_existing_CRL *minus* pre-update-time. For example, if the Next-Update of the existing CRL is 2015-06-30 06:00 and pre-update-time is 1 hour, then the system will start downloading at 2015-06-30, 05:00.

Default

schedule-type next-update-based

Parameters

schedule-type

Specifies the type of time scheduler to update the CRL.

Values periodic, next-update-based

Platforms

7705 SAR Gen 2

27.17 scheduler

scheduler

Syntax

scheduler *scheduler-name* [**create**]

no scheduler *scheduler-name*

Context

[Tree] (config>service>vpls>sap>egress>sched-override scheduler)

Full Context

configure service vpls sap egress scheduler-override scheduler

Description

This command overrides specific attributes of the specified scheduler name.

A scheduler defines a bandwidth control that limits each child (other schedulers, policers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created has policers, queues or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.

Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If *scheduler-name* already exists within the policy tier level (regardless of the inclusion of the keyword **create**), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause policers, queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If the *scheduler-name* exists within the policy on a different tier (regardless of the inclusion of the keyword create), an error occurs and the current CLI context does not change.

If the *scheduler-name* does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

1. The maximum number of schedulers has not been configured.
2. The provided *scheduler-name* is valid.
3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command does not execute, nor does the CLI context change.

If the provided scheduler-name is invalid according to the criteria below, a name syntax error occurs, the command does not execute, and the CLI context does not change.

The **no** form of this command removes the scheduler name from the configuration.

Parameters

scheduler-name

Specifies name of the scheduler

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

create

This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable create is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

Platforms

7705 SAR Gen 2

scheduler

Syntax

scheduler *scheduler-name* [**create**]

no scheduler *scheduler-name*

Context

[Tree] (config>port>ethernet>access>egr>qgrp>sched-override scheduler)

[Tree] (config>port>ethernet>access>ing>qgrp>sched-override scheduler)

Full Context

configure port ethernet access egress queue-group scheduler-override scheduler

configure port ethernet access ingress queue-group scheduler-override scheduler

Description

This command can be used to override specific attributes of the specified scheduler name. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers. The *scheduler-name* must exist in the applied scheduler policy.

The **no** form of this command removes the scheduler overrides for the specified scheduler and returns the scheduler's parent weight and CIR weight, and its PIR and CIR to the values configured in the applied scheduler policy.

Parameters

scheduler-name

Specifies the name of the scheduler.

Values Valid names consist of any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

create

Creates a new scheduler for this port.

Platforms

7705 SAR Gen 2

scheduler

Syntax

scheduler *scheduler-name* [**create**]

no scheduler *scheduler-name*

Context

[Tree] (config>service>epipe>sap>egress>sched-override scheduler)

[Tree] (config>service>epipe>sap>ingress>sched-override scheduler)

Full Context

configure service epipe sap egress scheduler-override scheduler

configure service epipe sap ingress scheduler-override scheduler

Description

This command can be used to override specific attributes of the specified scheduler name. A scheduler defines bandwidth controls that limit each child (other schedulers, policers, and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have policers, queues or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.

Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If *scheduler-name* already exists within the policy tier level (regardless of the inclusion of the keyword create), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause policers, queues, or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If the *scheduler-name* exists within the policy on a different tier (regardless of the inclusion of the keyword create), an error occurs and the current CLI context will not change.

If the *scheduler-name* does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

1. The maximum number of schedulers has not been configured.
2. The provided *scheduler-name* is valid.
3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command will not execute, nor will the CLI context change.

If the provided scheduler-name is invalid according to the following criteria, a name syntax error will occur, the command will not execute, and the CLI context will not change.

Parameters

scheduler-name

The name of the scheduler. Each scheduler must be explicitly created.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

create

This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable create is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

Platforms

7705 SAR Gen 2

scheduler

Syntax

scheduler *scheduler-name* [**create**]

no scheduler *scheduler-name*

Context

[Tree] (config>service>vprn>if>sap>ingress>sched-override scheduler)

[Tree] (config>service>vprn>if>sap>egress>sched-override scheduler)

Full Context

configure service vprn interface sap ingress scheduler-override scheduler

configure service vprn interface sap egress scheduler-override scheduler

Description

This command can be used to override specific attributes of the specified scheduler name.

A scheduler defines a bandwidth controls that limit each child (other schedulers, policers, and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have policers, queues, or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.

Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If *scheduler-name* already exists within the policy tier level (regardless of the inclusion of the keyword create), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If the *scheduler-name* exists within the policy on a different tier (regardless of the inclusion of the keyword create), an error occurs and the current CLI context will not change.

If the *scheduler-name* does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

1. The maximum number of schedulers has not been configured.
2. The provided *scheduler-name* is valid.
3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command will not execute, nor will the CLI context change.

If the provided scheduler-name is invalid according to the criteria below, a name syntax error will occur, the command will not execute, and the CLI context will not change.

Parameters

scheduler-name

Specifies the name of the scheduler.

Values Valid names consist of any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

create

Specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable create is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

Platforms

7705 SAR Gen 2

scheduler

Syntax

scheduler *scheduler-name* [**create**]

no scheduler *scheduler-name*

Context

[Tree] (config>service>ies>if>sap>egress>sched-override scheduler)

[Tree] (config>service>ies>if>sap>ingress>sched-override scheduler)

Full Context

configure service ies interface sap egress scheduler-override scheduler

configure service ies interface sap ingress scheduler-override scheduler

Description

This command can be used to override specific attributes of the specified scheduler name.

A scheduler defines a bandwidth controls that limit each child (other schedulers, policers, and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have policers, queues, or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.

Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If *scheduler-name* already exists within the policy tier level (regardless of the inclusion of the keyword create), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed

on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If the *scheduler-name* exists within the policy on a different tier (regardless of the inclusion of the keyword create), an error occurs and the current CLI context will not change.

If the *scheduler-name* does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

1. The maximum number of schedulers has not been configured.
2. The provided *scheduler-name* is valid.
3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command will not execute, nor will the CLI context change.

If the provided scheduler-name is invalid according to the criteria below, a name syntax error will occur, the command will not execute, and the CLI context will not change.

Parameters

scheduler-name

The name of the scheduler. Each scheduler must be explicitly created.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

create

This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable create is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

Platforms

7705 SAR Gen 2

scheduler

Syntax

scheduler *scheduler-name* [**create**]

no scheduler *scheduler-name*

Context

[Tree] (config>qos>scheduler-policy>tier scheduler)

Full Context

configure qos scheduler-policy tier scheduler

Description

This command creates a new scheduler or edits an existing scheduler within the scheduler policy tier. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.

Each scheduler must have a unique name within the context of the scheduler policy; however, the same name can be reused in multiple scheduler policies. If *scheduler-name* already exists within the policy tier level (regardless of the inclusion of the keyword create), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce SLAs.

If the *scheduler-name* exists within the policy on a different tier (regardless of the inclusion of the keyword create), an error occurs and the current CLI context will not change.

If the *scheduler-name* does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

1. The maximum number of schedulers has not been configured.
2. The provided *scheduler-name* is valid.
3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs, the command will not execute, nor will the CLI context change.

If the provided scheduler-name is invalid according to the criteria below, a name syntax error will occur, the command will not execute, and the CLI context will not change.

Parameters

scheduler-name

Specifies the scheduler name.

Values Valid names consist of any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

create

This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created

when the system environment variable create is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

Platforms

7705 SAR Gen 2

scheduler

Syntax

scheduler *scheduler-name* [**create**]

no scheduler *scheduler-name*

Context

[Tree] (config>service>cust>multi-service-site>egress>sched-override scheduler)

[Tree] (config>service>cust>multi-service-site>ingress>sched-override scheduler)

Full Context

configure service customer multi-service-site egress scheduler-override scheduler

configure service customer multi-service-site ingress scheduler-override scheduler

Description

This command override specifics attributes of the specified scheduler name.

A scheduler defines bandwidth controls that limit each child (other schedulers, policers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have policers, queues or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.

Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If *scheduler-name* already exists within the policy tier level (regardless of the inclusion of the keyword create), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause policer, queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If the *scheduler-name* exists within the policy on a different tier (regardless of the inclusion of the keyword create), an error occurs and the current CLI context will not change.

If the *scheduler-name* does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

1. The maximum number of schedulers has not been configured.
2. The provided *scheduler-name* is valid.

3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command will not execute, nor will the CLI context change.

If the provided scheduler-name is invalid according to the criteria below, a name syntax error will occur, the command will not execute, and the CLI context will not change.

The **no** form of the command disables the scheduler override.

Parameters

scheduler-name

Specifies the name of the scheduler.

Values Valid names consist of any string up to 32 characters in length, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

create

This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable create is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

Platforms

7705 SAR Gen 2

27.18 scheduler-override

scheduler-override

Syntax

[**no**] **scheduler-override**

Context

[Tree] (config>service>vpls>sap>ingress scheduler-override)

[Tree] (config>service>vpls>sap>egress scheduler-override)

Full Context

configure service vpls sap ingress scheduler-override

configure service vpls sap egress scheduler-override

Description

Commands in this context configure the set of attributes whose values have been overridden via management on this virtual scheduler. Clearing a given flag returns the corresponding overridden attribute to the value defined on the SAP's ingress scheduler policy.

The **no** form of this command removes scheduler parameters from the configuration.

Platforms

7705 SAR Gen 2

scheduler-override

Syntax

[**no**] **scheduler-override**

Context

[Tree] (config>port>ethernet>access>ing>qgrp scheduler-override)

[Tree] (config>port>ethernet>access>egr>qgrp scheduler-override)

Full Context

configure port ethernet access ingress queue-group scheduler-override

configure port ethernet access egress queue-group scheduler-override

Description

This command specifies the set of attributes whose values have been overridden by management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the ingress or egress queue group template.

The **no** form of this command removes all of the scheduler overrides and returns the scheduler's parent weight and CIR weight, and its PIR and CIR to the values configured in the applied scheduler policy.

Platforms

7705 SAR Gen 2

scheduler-override

Syntax

[**no**] **scheduler-override**

Context

[Tree] (config>service>epipe>sap>ingress scheduler-override)

[Tree] (config>service>epipe>sap>egress scheduler-override)

Full Context

configure service epipe sap ingress scheduler-override

configure service epipe sap egress scheduler-override

Description

This command specifies the set of attributes whose values have been overridden by management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the SAP's ingress scheduler policy.

Platforms

7705 SAR Gen 2

scheduler-override

Syntax

[**no**] **scheduler-override**

Context

[Tree] (config>service>ies>if>sap>ingress scheduler-override)

[Tree] (config>service>ies>if>sap>egress scheduler-override)

Full Context

configure service ies interface sap ingress scheduler-override

configure service ies interface sap egress scheduler-override

Description

This command specifies the set of attributes whose values have been overridden via management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the SAP's ingress scheduler policy.

Platforms

7705 SAR Gen 2

scheduler-override

Syntax

[**no**] **scheduler-override**

Context

[Tree] (config>service>vprn>if>sap>egress scheduler-override)

[Tree] (config>service>vprn>if>sap>ingress scheduler-override)

Full Context

configure service vprn interface sap egress scheduler-override

configure service vprn interface sap ingress scheduler-override

Description

This command specifies the set of attributes whose values have been overridden via management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the SAP's ingress scheduler policy.

Platforms

7705 SAR Gen 2

scheduler-override

Syntax

[**no**] **scheduler-override**

Context

[Tree] (config>service>cust>multi-service-site>ingress scheduler-override)

[Tree] (config>service>cust>multi-service-site>egress scheduler-override)

Full Context

configure service customer multi-service-site ingress scheduler-override

configure service customer multi-service-site egress scheduler-override

Description

This command specifies the set of attributes whose values have been overridden by management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the SAP's ingress and egress scheduler policy.

The **no** form of the command disables the override.

Platforms

7705 SAR Gen 2

27.19 scheduler-policy

scheduler-policy

Syntax

scheduler-policy *scheduler-policy-name*

no scheduler-policy

Context

[Tree] (config>service>vprn>if>sap>ingress scheduler-policy)

[Tree] (config>service>ies>if>sap>ingress scheduler-policy)

[Tree] (config>service>vprn>if>sap>egress scheduler-policy)

[Tree] (config>service>vpls>sap>ingress scheduler-policy)

[Tree] (config>service>ies>if>sap>egress scheduler-policy)

[Tree] (config>service>vpls>sap>egress scheduler-policy)

Full Context

configure service vprn interface sap ingress scheduler-policy

configure service ies interface sap ingress scheduler-policy

configure service vprn interface sap egress scheduler-policy

configure service vpls sap ingress scheduler-policy

configure service ies interface sap egress scheduler-policy

configure service vpls sap egress scheduler-policy

Description

This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created once the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the **config>qos>scheduler-policy** *scheduler-policy-name* context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the ingress SAP queues and egress SAP policers and queues associated with the customer site. Policers and queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have policers or queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more policers or queues. When the **no scheduler-policy** command is executed, the customer site's ingress or egress node will not contain an applied scheduler policy.

Parameters

scheduler-policy-name

Specifies that the *scheduler-policy-name* is applied to an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress or egress queues created on associated SAPs.

Values Any existing valid scheduler policy name.

Platforms

7705 SAR Gen 2

scheduler-policy

Syntax

scheduler-policy *scheduler-policy-name*

no scheduler-policy

Context

[Tree] (config>port>ethernet>network>egress>queue-group scheduler-policy)

Full Context

configure port ethernet network egress queue-group scheduler-policy

Description

This command configures a scheduler policy for the egress queue group.

Parameters

scheduler-policy-name

Specifies the scheduler policy name, up to 32 characters.

Platforms

7705 SAR Gen 2

scheduler-policy

Syntax

scheduler-policy *scheduler-policy-name*

no scheduler-policy

Context

[Tree] (config>service>epipe>sap>ingress scheduler-policy)

[Tree] (config>service>epipe>sap>egress scheduler-policy)

Full Context

configure service epipe sap ingress scheduler-policy

configure service epipe sap egress scheduler-policy

Description

This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created when the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the **config>qos>scheduler-policy** *scheduler-policy-name* context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the ingress SAP queues associated with the customer site. Policers or queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have policers or queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more policers or queues. When the **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

Parameters

scheduler-policy-name

The *scheduler-policy-name* parameter applies an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress or egress queues and to egress policers managed by HQoS created on associated SAPs.

Platforms

7705 SAR Gen 2

scheduler-policy

Syntax

scheduler-policy *scheduler-policy-name* [**create**]

no scheduler-policy *scheduler-policy-name*

Context

[Tree] (config>qos scheduler-policy)

Full Context

configure qos scheduler-policy

Description

Each scheduler policy is divided up into groups of schedulers based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations.

The **scheduler-policy** command creates a scheduler policy or allows editing of an existing policy. The policy defines the hierarchy and operating parameters for virtual schedulers. Creating a policy does not create the schedulers; it only provides a template for the schedulers to be created when the policy is associated with a SAP or multiservice site.

Each scheduler policy must have a unique name within the context of the system. Modifications made to an existing policy are executed on all schedulers that use the policy. This can cause queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce SLAs.

If a *scheduler-policy-name* does not exist, it is assumed that an attempt is being made to create a new policy. The success of the command execution is dependent on the following:

1. The maximum number of scheduler policies has not been configured.
2. The provided scheduler-policy-name is valid.
3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of scheduler policies has been exceeded, a configuration error occurs, the command will not execute, and the CLI context will not change.

If the provided scheduler-policy-name is invalid according to the criteria below, a name syntax error occurs, the command will not execute, and the CLI context will not change.

Parameters

scheduler-policy-name

The name of the scheduler policy.

Values Valid names consist of any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

scheduler-policy

Syntax

scheduler-policy *scheduler-policy-name*

no scheduler-policy

Context

[Tree] (config>service>cust>multi-service-site>egress scheduler-policy)

[Tree] (config>service>cust>multi-service-site>ingress scheduler-policy)

Full Context

configure service customer multi-service-site egress scheduler-policy

configure service customer multi-service-site ingress scheduler-policy

Description

This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues or, at egress only, policers associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created once the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the **config>qos>scheduler-policy** *scheduler-policy-name* context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the SAP policers and queues associated with the customer site. Policers and queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler.

The SAPs that have ingress queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more policers and queues. When the **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

Parameters

scheduler-policy-name

Applies an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress or egress queues and egress policers managed by HQoS created on associated SAPs.

Values Any existing valid scheduler policy name up to 32 characters in length.

Platforms

7705 SAR Gen 2

27.20 schema-path

schema-path

Syntax

schema-path *url-string*

no schema-path

Context

[Tree] (config>system>management-interface schema-path)

Full Context

configure system management-interface schema-path

Description

This command specifies the schema path where the SR OS YANG modules can be placed by the user before using a <get-schema> request. Nokia recommends that the URL string not exceed 135 characters for the <get-schema> request to work correctly with all schema files.

If this command is not configured, the software upgrade process manages the YANG schema files to ensure the schema files are synchronized with the software image on both the primary and standby CPM.

The **no** form of this command reverts to the default value.

Default

no schema-path

Parameters

url-string

Specifies the schema path URL up to 180 characters. However, Nokia recommends that the string shall not exceed 135 characters to ensure that the <get-schema> request works properly with *all* schema files.

Platforms

7705 SAR Gen 2

27.21 scope

scope

Syntax

scope {**exclusive** | **template**}

no scope

Context

[Tree] (config>qos>sap-ingress scope)

Full Context

configure qos sap-ingress scope

Description

This command configures the Service Ingress QoS policy scope as exclusive or template.

The policy's scope cannot be changed if the policy is applied to a service.

The **no** form of this command sets the scope of the policy to the default of **template**.

Default

scope template

Parameters

exclusive

When the scope of a policy is defined as exclusive, the policy can only be applied to one SAP. If a policy with an exclusive scope is assigned to a second SAP, an error message

is generated. If the policy is removed from the exclusive SAP, it will become available for assignment to another exclusive SAP.

The system default policies cannot be put into the exclusive scope. An error will be generated if scope exclusive is executed in any policies with a policy-id equal to 1.

template

When the scope of a policy is defined as template, the policy can be applied to multiple SAPs on the router.

Default QoS policies are configured with template scopes. An error is generated when the template scope parameter to exclusive scope on default policies is modified.

Platforms

7705 SAR Gen 2

scope

Syntax

scope {**exclusive** | **template**}

no scope

Context

[Tree] (config>qos>sap-egress scope)

Full Context

configure qos sap-egress scope

Description

Enter the scope of this policy. The scope of the policy cannot be changed if the policy is applied to one or more services.

The no form of this command sets the scope of the policy to the default of template.

Default

scope template

Parameters

exclusive

When the scope of a policy is defined as exclusive, the policy can only be applied to a single SAP. Attempting to assign the policy to a second SAP will result in an error message. If the policy is removed from the exclusive SAP, it will become available for assignment to another exclusive SAP.

The system default policies cannot be put into the exclusive scope. An error will be generated if scope exclusive is executed in any policies with a policy-id equal to 1.

template

When the scope of a policy is defined as template, the policy can be applied to multiple SAPs on the router.

Platforms

7705 SAR Gen 2

scope

Syntax

scope {**exclusive** | **template**}

no scope

Context

[Tree] (config>qos>network scope)

Full Context

configure qos network scope

Description

This command configures the network policy scope as exclusive or template. The policy's scope cannot be changed if the policy is applied to an interface.

The **no** form of this command sets the scope of the policy to the default of **template**.

Default

scope template

Parameters

exclusive

When the scope of a policy is defined as exclusive, the policy can only be applied to one interface. If a policy with an exclusive scope is assigned to a second interface, an error message is generated. If the policy is removed from the exclusive interface, it will become available for assignment to another exclusive interface.

The system default policies cannot be put into the exclusive scope. An error will be generated if the **scope exclusive** command is executed in any policies with a policy-id equal to 1.

template

When the scope of a policy is defined as template, the policy can be applied to multiple interfaces on the router.

Default QoS policies are configured with template scopes. An error is generated if the template scope parameter is modified to exclusive scope on default policies.

Platforms

7705 SAR Gen 2

scope

Syntax

scope {**exclusive** | **template** | **embedded** | **system**}

scope {**exclusive** | **template**}

no scope

Context

[Tree] (config>filter>ip-exception scope)

[Tree] (config>filter>ipv6-filter scope)

[Tree] (config>filter>ip-filter scope)

Full Context

configure filter ip-exception scope

configure filter ipv6-filter scope

configure filter ip-filter scope

Description

This command configures the filter policy scope as exclusive, template, embedded or system.

The scope of the policy cannot be changed when:

- the scope is **template** and the policy is applied to one or more services or network interfaces
- the scope is **embedded** and the policy is embedded by another policy

Changing the scope to/from system is only allowed when a policy is not active and the policy has no entries configured.

The **no** form of the command sets the scope of the policy to the default of **template**.

Default

scope template

Parameters

exclusive

Specifies that the policy can only be applied to a single entity. Attempting to assign the policy to a second entity will result in an error message.

template

Specifies that the policy can be applied to multiple entities.

embedded

Specifies that the policy cannot be applied directly. The policy defines embedded filter rules, which are embedded by other exclusive/template/system filter policies. The **embedded** scope is supported for IPv4 and IPv6 filter policies only.

system

Specifies that the policy defines system-wide filter rules. To apply system policy rules, activate system filter and chain exclusive/template ACL filter policy to the system filter. The **system** scope is supported for IPv4 and IPv6 filter policies only.

Platforms

7705 SAR Gen 2

27.22 scp

scp

Syntax

scp *local-file-url destination-file-url* [**router** *router-instance*] [**force**]

scp *local-file-url destination-file-url* [**force**] **service** *service-name*

Context

[Tree] (file scp)

Full Context

file scp

Description

This command copies a local file to a remote host file system. It uses `ssh` for data transfer, and uses the same authentication and provides the same security as `ssh`. The following prompt appears:

"Are you sure (y/n)?" The destination must specify a user and a host.

Parameters

local-file-url

Specifies the local source file or directory.

Values

[*cflash-id/*] *file-path*	up to 200 characters
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

destination-file-url

Specifies the destination file.

Values

*destination-file-**: *user@hostname:file-path* - up to 255 characters

user	up to 32 characters
hostname	[*dns-name* | *ipv4-address* | "["*ipv6-address*"]"]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*] *x:x:x:x:x:x:d.d.d.d*[-*interface*] *x* - [0 to FFFF]H *d* - [0 to 255]D *interface* - up to 32 characters, mandatory for link local addresses
dns-name	up to 128 characters
file-path	up to 200 characters, directory length up to 99 characters

user

Specifies the SSH user.

hostname

Specifies the remote host IP address of DNS name.

file-path

Specifies the destination path.

router-instance

Specifies the router name or service ID used to specify the router instance.

Values

router-name	"Base", "management", "vpls-management"
vprn-service-id	1 to 2147483647

Default Base

force

Forces an immediate copy of the specified file. The command **file scp** *local-file-url destination-file-url* [**router** *router-instance*] **force** executes the command without displaying a user prompt message.

service-name

Specifies the service name used to identify the router instance. The service name can be a maximum of 64 characters long.

Platforms

7705 SAR Gen 2

27.23 script

script

Syntax

script *script-name* [**owner** *script-owner*]

no script

Context

[Tree] (config>system>script-control>script-policy script)

[Tree] (config>system>script-control script)

Full Context

configure system script-control script-policy script

configure system script-control script

Description

This command is used to configure a script to be run.

The **no** form of the command removes the script.

Default

no script

Parameters

script-name

Specifies the name of the script. Can be up to 32 characters.

script-owner

Specifies the name of the script owner. Can be up to 32 characters.

The owner is an arbitrary name and not necessarily a user name. Commands in the scripts are not authorized against the owner. The **configure system security cli-script authorization** *x* **cli-user** command determines the user context against which commands in the scripts are authorized.

Default "TiMOS CLI”

Platforms
7705 SAR Gen 2

27.24 script-control

script-control

Syntax
script-control

Context
[Tree] (config>system script-control)

Full Context
configure system script-control

Description
Commands in this context configure command script parameters.

Platforms
7705 SAR Gen 2

27.25 script-policy

script-policy

Syntax
script-policy *policy-name* [**owner** *policy-owner*]
no script-policy

Context
[Tree] (config>system>cron>schedule script-policy)

Full Context
configure system cron schedule script-policy

Description
This command is used to configure the CLI script policy.

Parameters

policy-name

Specifies the name of the policy. Can be up to 32 characters.

policy-owner

Specifies the name of the policy owner. Can be up to 32 characters.

The owner is an arbitrary name and not necessarily a user name. Commands in the scripts are not authorized against the owner. The **configure system security cli-script authorization** *x* **cli-user** command determines the user context against which commands in the scripts are authorized.

Default "TiMOS CLI"

Platforms

7705 SAR Gen 2

script-policy

Syntax

[**no**] **script-policy** *policy-name* [**owner** *policy-owner*]

Context

[Tree] (config>system>script-control script-policy)

Full Context

configure system script-control script-policy

Description

This command is used to configure the CLI script policy.

Parameters

policy-name

Specifies the name of the policy, up to 32 characters.

policy-owner

Specifies the name of the policy owner, up to 32 characters.

The owner is an arbitrary name and not necessarily a user name. Commands in the scripts are not authorized against the owner. The **configure system security cli-script authorization** *x* **cli-user** command determines the user context against which commands in the scripts are authorized.

Default "TiMOS CLI"

Platforms

7705 SAR Gen 2

script-policy

Syntax

script-policy *policy-name* [**owner** *policy-owner*]

no script-policy

Context

[Tree] (config>log>event-handling>handler>action-list>entry script-policy)

Full Context

configure log event-handling handler action-list entry script-policy

Description

This command configures the script policy parameters to use for this EHS handler action-list entry. The associated script is launched when the handler is triggered.

Default

no script-policy

Parameters

policy-name

Specifies the script policy name. Can be up to 32 characters maximum.

owner policy-owner

Specifies the script policy owner. Can be up to 32 characters maximum.

Default "TiMOS CLI"

Platforms

7705 SAR Gen 2

27.26 sd-offset

sd-offset

Syntax

sd-offset *offset-value*

no sd-offset

Context

[Tree] (config>service>vprn>isis>if>level sd-offset)

Full Context

configure service vprn isis interface level sd-offset

Description

If the pre-FEC error rate of the associated DWDM port crosses the configured **sd-threshold**, this offset-value is added to the IS-IS interface metric. This parameter is only effective if the interface is associated with a DWDM port and the **sd-threshold** value is configured under that port.

The **no** form of this command reverts the offset value to 0.

Default

no sd-offset

Parameters

offset-value

Specifies the amount the interface metric is increased by if the **sd-threshold** is crossed.

Values 0 to 16777215

Platforms

7705 SAR Gen 2

sd-offset

Syntax

sd-offset *sd-offset*

no sd-offset

Context

[Tree] (config>router>isis>if>level sd-offset)

Full Context

configure router isis interface level sd-offset

Description

If the pre-FEC error rate of the associated DWDM port crosses the configured **sd-threshold**, this offset-value is added to the IS-IS interface metric. This parameter is only effective if the interface is associated with a DWDM port and the **sd-threshold** value is configured under that port.

The **no** form of this command reverts the offset value to 0.

Default

no sd-offset

Parameters

sd-offset

Specifies the amount the interface metric is increased by if the **sd-threshold** is crossed.

Values 0 to 16777215

Platforms

7705 SAR Gen 2

27.27 sd-threshold

sd-threshold

Syntax

sd-threshold *threshold* [**multiplier** *multiplier*]

no sd-threshold

Context

[Tree] (config>port>ethernet>crc-monitor sd-threshold)

Full Context

configure port ethernet crc-monitor sd-threshold

Description

This command specifies the error rate at which to declare the Signal Degrade condition on an Ethernet interface. The value represents M*10E-N a ratio of errored frames over total frames received over W seconds of the sliding window. The CRC errors on the interface are sampled once per second. A default of 10 seconds is used when there is no additional window-size configured. The multiplier keyword is optional. If the multiplier keyword is omitted or **no sd-threshold** is specified the multiplier will return to the default value of 1.

Default

no sd-threshold

Parameters

threshold

Specifies the threshold value.

Values 1 to 9

multiplier

Specifies the multiplier value.

Values 1 to 9

Platforms

7705 SAR Gen 2

27.28 sdp

sdp

Syntax

[**no**] **sdp** *sdp-id:vc-id*

Context

[Tree] (debug>service>id sdp)

[Tree] (debug>service>id>stp sdp)

[Tree] (debug>service>id>dhcp sdp)

Full Context

debug service id sdp

debug service id stp sdp

debug service id dhcp sdp

Description

This command enables STP debugging for a specific SDP.

The **no** form of the command disables debugging.

Parameters

sdp-id:vc-id

Specifies the SDP ID and VC ID.

Values sdp-id: 1 to 17407

vc-id: 1 to 4294967295

Platforms

7705 SAR Gen 2

sdp

Syntax

[**no**] **sdp** *sdp-id:vc-id*

Context

[Tree] (debug>service>id>igmp-snooping sdp)

Full Context

debug service id igmp-snooping sdp

Description

This command shows IGMP packets for a specific SDP.

The **no** form of this command disables the debugging for the SDP.

Parameters

sdp-id

Displays only IGMP snooping entries associated with the specified mesh SDP or spoke-SDP. For a spoke-SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id

Displays information for the specified virtual circuit ID on the SDP ID

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

sdp

Syntax

[**no**] **sdp** *sdp-id:vc-id*

Context

[Tree] (debug>service>id>mld sdp)

Full Context

debug service id mld-snooping sdp

Description

This command shows MLD packets for a specific SDP.

The **no** form of this command disables the debugging for the SDP.

Parameters

sdp-id

Displays only MLD entries associated with the specified mesh SDP or spoke-SDP

Values 1 to 17407

vc-id

Displays information for the specified virtual circuit ID on the SDP ID

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

sdp

Syntax

sdp *sdp-id* [*delivery-type*] [**create**]

no sdp *sdp-id*

Context

[Tree] (config>service sdp)

Full Context

configure service sdp

Description

This command creates or edits a service destination point (SDP). SDPs must be explicitly configured.

An SDP is a logical mechanism that ties a far-end router to a particular service without having to specifically define far-end SAPs. Each SDP represents a method to reach another router.

One method is IP Generic Router Encapsulation (GRE), which has no state in the core of the network. GRE does not specify a specific path to the far-end router. A GRE-based SDP uses the underlying IGP routing table to find the best next hop to the far-end router.

The second method is Multi-Protocol Label Switching (MPLS) encapsulation. A router supports both signaled and non-signaled Label Switched Paths (LSPs) through the network. Non-signaled paths are defined at each hop through the network. Signaled paths are communicated by protocol from end-to-end using Resource Reservation Protocol (RSVP). Paths may be manually defined or a constraint-based routing protocol (such as OSPF-TE or CSPF) can be used to determine the best path with specific constraints. An LDP LSP can also be used for an SDP when the encapsulation is MPLS. The use of an

LDP LSP type or an RSVP/Static LSP type are mutually exclusive except when the mixed-lsp option is enabled on the SDP.

Segment routing is another MPLS tunnel type and is used to allow service binding to an SR tunnel programmed in TTM by OSPF or IS-IS. The SDP of type **sr-isis** or **sr-ospf** can be used with the **far-end** option. The **tunnel-far-end** option is not supported. In addition, the **mixed-lsp-mode** option does not support the **sr-isis** and **sr-ospf** tunnel types.

L2TPv3-over-IPv6 transport is also an option for 7705 SAR Gen 2 Ethernet Pipe (Epipe) Services. Like GRE, L2TPv3 is stateless in the core of the network, as well as on the service nodes as the L2TPv3 control plane functionality is disabled for this SDP type. A unique source and destination IPv6 address combined with TX and RX Cookie values are used to ensure that the SDP is bound to the correct service.

SDPs are created and then bound to services. Many services may be bound to a single SDP. The operational and administrative state of the SDP controls the state of the SDP binding to the service.

If the *sdp-id* does not exist, a new SDP is created. When creating an SDP, either the **gre, mpls, or l2tpv3** keyword must be specified. SDPs are created in the admin down state (**shutdown**) and the **no shutdown** command must be executed once all relevant parameters are defined and before the SDP can be used.

If *sdp-id* exists, the current CLI context is changed to that SDP for editing and modification. For editing an existing SDP, neither the **gre, mpls**, or **l2tpv3** keyword is specified. If a keyword is specified for an existing *sdp-id*, an error is generated and the context of the CLI will not be changed to the specified *sdp-id*.

The **no** form of this command deletes the specified SDP. Before an SDP can be deleted, it must be administratively down (shutdown) and not bound to any services. If the specified SDP is bound to a service, the **no sdp** command will fail generating an error message specifying the first bound service found during the deletion process. If the specified *sdp-id* does not exist an error will be generated.

Parameters

sdp-id

Specifies the SDP identifier.

Values 1 to 32767

gre

Specifies the SDP will use GRE to reach the far-end router. The GRE encapsulation of the MPLS service packet uses the base 4-byte header as per RFC 2890. The optional fields Checksum (plus Reserved field), Key, and Sequence Number are not inserted. Only one GRE SDP can be created to a given destination address. Multiple GRE SDPs to a single destination address serve no purpose as the path taken to reach the far end is determined by the IGP which will be the same for all SDPs to a given destination and there is no bandwidth reservation in GRE tunnels.

mpls

Specifies the SDP will use MPLS encapsulation and one or more LSP tunnels to reach the far-end device. Multiple MPLS SDPs may be created to a given destination device. Multiple MPLS SDPs to a single destination device are helpful when they use divergent paths.

l2tpv3

Specifies the SDP will use L2TPv3-over-IPv6 encapsulation. One SDP is created per service, regardless of whether the far-end node is common or not. Unique local and far-end addresses are configured for every L2TPv3 SDP type. The local address must exist on the local node.

eth-gre-bridged

Configures the SDP as an L2oGRE tunnel that is terminated on an FPE-based PW port. Only the end-points of such a tunnel (the far-end IPv4/IPv6 address or local-end IPv4/IPv6 address) are allowed to be configured under this SDP.

Platforms

7705 SAR Gen 2

27.29 sdp-exclude

sdp-exclude

Syntax

[**no**] **sdp-exclude** *group-name*

Context

[Tree] (config>service>pw-template sdp-exclude)

Full Context

configure service pw-template sdp-exclude

Description

This command configures SDP admin group constraints for a pseudowire template.

The admin group name must have been configured or the command is failed. The user can execute the command multiple times to include or exclude more than one admin group. The sdp-include and sdp-exclude commands can only be used with the **use-provisioned-sdp** or **prefer-provisioned-sdp** options. If the same group name is included and excluded within the same pseudowire template, only the exclude option will be enforced.

Any changes made to the admin group sdp-include and sdp-exclude constraints will only be reflected in existing spoke-sdps after the following command has been executed:

tools>perform>service>eval-pw-template>allow-service-impact

When the service is bound to the pseudowire template, the SDP selection rules will enforce the admin group constraints specified in the sdp-include and sdp-exclude commands.

In the SDP selection process, all provisioned SDPs with the correct far-end IP address, the correct tunnel-far-end IP address, and the correct service label signaling are considered. The SDP with the lowest admin metric is selected. If more than one SDP with the same lowest metric are found then the SDP with the highest sdp-id is selected. The type of SDP, GRE or MPLS (BGP/RSVP/LDP) is not a criterion in this selection.

The selection rule with SDP admin groups is modified such that the following admin-group constraints are applied upfront to prune SDPs that do not comply:

- if one or more **sdp-include** statement is part of the PW template, then an SDP that is a member of one or more of the included groups will be considered. With the **sdp-include** statement, there is no

preference for an SDP that belongs to all included groups versus one that belongs to one or fewer of the included groups. All SDPs satisfying the admin-group constraint will be considered and the selection above based on the lowest metric and highest sdp-id is applied.

- if one or more **sdp-exclude** statement is part of the PW template, then an sdp that is a member of any of the excluded groups will not be considered.

SDP admin group constraints can be configured on all router services that makes use of the pseudowire template (BGP-AD VPLS service, BGP-VPLS service, and FEC129 VLL service). In the latter case, only support at a T-PE node is provided.

The **no** form of this command removes the SDP admin group constraints from the pseudowire template.

Parameters

group-name

Specifies the name of the SDP admin group. A maximum of 32 characters can be entered.

Platforms

7705 SAR Gen 2

27.30 sdp-group

sdp-group

Syntax

sdp-group

Context

[Tree] (config>service sdp-group)

Full Context

configure service sdp-group

Description

This command configures the SDP membership in admin groups.

The user can enter a maximum of one (1) admin group name at once. The user can execute the command multiple times to add membership to more than one admin group. The admin group name must have been configured or the command is failed. Admin groups are supported on an SDP of type GRE and of type MPLS (BGP/RSVP/LDP). They are also supported on an SDP with the mixed-lsp-mode option enabled.

The **no** form of this command removes this SDP membership to the specified admin group.

Platforms

7705 SAR Gen 2

27.31 sdp-include

sdp-include

Syntax

[**no**] **sdp-include** *group-name*

Context

[Tree] (config>service>pw-template sdp-include)

Full Context

configure service pw-template sdp-include

Description

This command configures SDP admin group constraints for a pseudowire template.

The admin group name must have been configured or the command is failed. The user can execute the command multiple times to include or exclude more than one admin group. The sdp-include and sdp-exclude commands can only be used with the **use-provisioned-sdp** or **prefer-provisioned-sdp** options. If the same group name is included and excluded within the same pseudowire template, only the exclude option will be enforced.

Any changes made to the admin group sdp-include and sdp-exclude constraints will only be reflected in existing spoke-sdps after the following command has been executed:

tools>perform>service>eval-pw-template>allow-service-impact

When the service is bound to the pseudowire template, the SDP selection rules will enforce the admin group constraints specified in the sdp-include and sdp-exclude commands.

In the SDP selection process, all provisioned SDPs with the correct far-end IP address, the correct tunnel-far-end IP address, and the correct service label signaling are considered. The SDP with the lowest admin metric is selected. If more than one SDP with the same lowest metric are found then the SDP with the highest sdp-id is selected. The type of SDP, GRE or MPLS (BGP/RSVP/LDP) is not a criterion in this selection.

The selection rule with SDP admin groups is modified such that the following admin-group constraints are applied upfront to prune SDPs that do not comply:

- if one or more **sdp-include** statement is part of the PW template, then an SDP that is a member of one or more of the included groups will be considered. With the **sdp-include** statement, there is no preference for an SDP that belongs to all included groups versus one that belongs to one or fewer of the included groups. All SDPs satisfying the admin-group constraint will be considered and the selection above based on the lowest metric and highest sdp-id is applied.
- if one or more **sdp-exclude** statement is part of the PW template, then an sdp that is a member of any of the excluded groups will not be considered.

SDP admin group constraints can be configured on all router services that make use of the pseudowire template (BGP-AD VPLS service, BGP-VPLS service, and FEC129 VLL service). In the latter case, only support at a T-PE node is provided.

The **no** form of this command removes the SDP admin group constraints from the pseudowire template.

Parameters

group-name

Specifies the name of the SDP admin group. A maximum of 32 characters can be entered.

Platforms

7705 SAR Gen 2

27.32 sdp-mtu

sdp-mtu

Syntax

sdp-mtu *orig-sdp-id* **size-inc** *start-octets end-octets* [**step** *step-size*] [**timeout** *timeout*] [**interval** *interval*]

Context

[Tree] (oam sdp-mtu)

Full Context

oam sdp-mtu

Description

Performs MTU Path tests on an SDP to determine the largest path-mtu supported on an SDP. The **size-inc** parameter can be used to easily determine the **path-mtu** of a given SDP-ID. The forwarding class is assumed to be Best-Effort Out-of-Profile. The message reply is returned with IP/GRE encapsulation from the far-end router. OAM request messages sent within an IP/GRE SDP must have the 'DF' IP header bit set to 1 to prevent message fragmentation.

To terminate an **sdp-mtu** in progress, use the CLI break sequence <Ctrl-C>.

Parameters

orig-sdp-id

Specifies the *sdp-id* to be used by **sdp-ping,** expressed as a decimal integer. The far-end address of the specified **sdp-id** is the expected *responder-id* within each reply received. The specified *sdp-id* defines the encapsulation of the SDP tunnel encapsulation used to reach the far end. This can be IP/GRE or MPLS. If *orig-sdp-id* is invalid or administratively down or unavailable, the SDP echo request message is not sent and an appropriate error message is displayed (once the **interval** timer expires, **sdp-ping** attempts to send the next request, if required).

Values 1 to 32767

start-octets

Specifies the beginning size in octets of the first message sent for an incremental MTU test, expressed as a decimal integer.

Values 40 to 9786

end-octets

Specifies the ending size in octets of the last message sent for an incremental MTU test, expressed as a decimal integer. The specified value must be greater than *start-octets*.

Values 40 to 9786

step-size

Specifies the number of octets to increment the message size request for each message sent for an incremental MTU test, expressed as a decimal integer. The next size message is not sent until a reply is received or three messages have timed out at the current size.

If the incremented size exceeds the *end-octets* value, no more messages are sent.

Values 1 to 512

Default 32

timeout

Specifies the *timeout* parameter in seconds, expressed as a decimal integer. This value is used to override the default *timeout* value and is the amount of time that the router waits for a message reply after sending the message request. Upon the expiration of the message time out, the requesting router assumes that the message response is not received. A **request timeout** message is displayed by the CLI for each message request sent that expires. Any response received after the request times out is silently discarded.

Values 1 to 10

Default 5

interval

Specifies the *interval* parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the *interval* is set to 1 second, and the *timeout* value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Values 1 to 10

Default 1

Platforms

7705 SAR Gen 2

Output

Output Example: SDP MTU Path Test

```
*A:Dut-A# oam sdp-mtu 1201 size-inc 512 3072 step 256
Size    Sent    Response
---------------------------
512     .       Success
768     .       Success
1024    .       Success
1280    .       Success
1536    .       Success
1792    .       Success
2048    .       Success
2304    .       Success
2560    .       Success
2816    .       Success
3072    .       Success

Maximum Response Size: 3072
*A:Dut-A#
```

27.33 sdp-ping

sdp-ping

Syntax

sdp-ping *orig-sdp-id* [**resp-sdp** *resp-sdp-id*] [**fc** *fc-name* [**profile** { **in** | **out**}]] [**size** *octets*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]

Context

[Tree] (oam sdp-ping)

[Tree] (config>saa>test>type sdp-ping)

Full Context

oam sdp-ping

configure saa test type sdp-ping

Description

This command tests SDPs for uni-directional or round trip connectivity and performs SDP MTU Path tests.

The **sdp-ping** command accepts an originating SDP-ID and an optional responding SDP-ID. The size, number of requests sent, message time-out and message send interval can be specified. All **sdp-ping** requests and replies are sent with PLP OAM-Label encapsulation, as a *service-id* is not specified.

For round trip connectivity testing, the **resp-sdp** keyword must be specified. If **resp-sdp** is not specified, a uni-directional SDP test is performed.

To terminate an **sdp-ping** in progress, use the CLI break sequence <Ctrl-C>.

An **sdp-ping** response message indicates the result of the **sdp-ping** message request. When multiple response messages apply to a single SDP echo request/reply sequence, the response message with the highest precedence is displayed. Table 85: sdp-ping Response Messages shows the response messages sorted by precedence.

Table 85: sdp-ping Response Messages

Result of Request	Displayed Response Message	Precedence
Request time out without reply	Request Timeout	1
Request not sent due to non-existent *orig-sdp-id*	Orig-SDP Non-Existent	2
Request not sent due to administratively down *orig-sdp-id*	Orig-SDP Admin-Down	3
Request not sent due to operationally down *orig-sdp-id*	Orig-SDP Oper-Down	4
Request terminated by user before reply or time out	Request Terminated	5
Reply received, invalid *origination-id*	Far End: Originator-ID Invalid	6
Reply received, invalid *responder-id*	Far End: Responder-ID Error	7
Reply received, non-existent *resp-sdp-id*	Far End: Resp-SDP Non-Existent	8
Reply received, invalid *resp-sdp-id*	Far End: Resp-SDP Invalid	9
Reply received, *resp-sdp-id* down (admin or oper)	Far-end: Resp-SDP Down	10
Reply received, No Error	Success	11

Parameters

orig-sdp-id

Specifies the SDP ID to be used by **sdp-ping,** expressed as a decimal integer. The far-end address of the specified SDP-ID is the expected *responder-id* within each reply received. The specified SDP-ID defines the encapsulation of the SDP tunnel encapsulation used to reach the far end. This can be IP/GRE or MPLS. If *orig-sdp-id* is invalid or administratively down or unavailable for some reason, the SDP Echo Request message is not sent and an appropriate error message is displayed (once the **interval** timer expires, **sdp-ping** attempts to send the next request if required).

Values 1 to 32767

resp-sdp-id

Specifies the return SDP-ID to be used by the far-end router for the message reply for round trip SDP connectivity testing. If *resp-sdp-id* does not exist on the far-end router, terminates on another router different than the originating router, or another issue prevents

the far-end router from using *resp-sdp-id*, the SDP echo reply is sent using generic IP/GRE OAM encapsulation. The received forwarding class (as mapped on the ingress network interface for the far end) defines the forwarding class encapsulation for the reply message.

Values 1 to 32767

Default null. Use the non-SDP return path for message reply.

fc-name

Specifies the parameter to be used to indicate the forwarding class of the SDP encapsulation. The actual forwarding class encoding is controlled by the network egress DSCP or LSP-EXP mappings.

The DSCP or LSP-EXP mappings on the receive network interface controls the mapping back to the internal forwarding class used by the far-end router that receives the message request. The egress mappings of the egress network interface on the far-end router controls the forwarding class markings on the return reply message.

The DSCP or LSP-EXP mappings on the receive network interface controls the mapping of the message reply at the originating router. This is displayed in the response message output upon receipt of the message reply.

Values be, l2, af, l1, h2, ef, h1, nc

Default be

profile {in | out}

Specifies the profile state of the SDP encapsulation.

Default out

octets

Specifies the **size** parameter in octets. This parameter is used to override the default message size for the **sdp-ping** request. Changing the message size is a method of checking the ability of an SDP to support a **path-mtu**. The size of the message does not include the SDP encapsulation, VC-Label (if applied) or any DLC headers or trailers.

When the OAM message request is encapsulated in an IP/GRE SDP, the IP 'DF' (Do Not Fragment) bit is set. If any segment of the path between the sender and receiver cannot handle the message size, the message is discarded. MPLS LSPs are not expected to fragment the message either, as the message contained in the LSP is not an IP packet.

Values 72 to 9786

Default 72

send-count

Specifies the number of messages to send. The **count** parameter is used to override the default number of message requests sent. Each message request must either time out or receive a reply before the next message request is sent. The message **interval** value must have expired before the next message request is sent.

Values 1 to 100

Default 1

timeout

Specifies the time, in seconds, used to override the default *timeout* value and is the amount of time that the router waits for a message reply after sending the last probe for a specific test. Upon the expiration of time out, the test is marked complete and no more packets is processed for any of those request probes.

Values 1 to 10

Default 5

interval

Specifies the time, in seconds, used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

Values 1 to 10

Default 1

Platforms

7705 SAR Gen 2

Output

Single Response Round Trip Connectivity Test Output Example

```
A:router1> sdp-ping 10 resp-sdp 22 fc ef
Request Result: Sent - Reply Received
RTT:30ms

Err  SDP-ID Info            Local     Remote
__   SDP-ID:                10        22
__   Administrative State:  Up        Up
__   Operative State:       Up        Up
__   Path MTU               4470      4470
__   Response SDP Used:                Yes

Err  System IP Interface Info
Local Interface Name: "ESR-System-IP-Interface (Up to 32 chars)…"
__   Local IP Interface State:          Up
__   Local IP Address:                  10.10.10.11
__   IP Address Expected By Remote:     10.10.10.11
__   Expected Remote IP Address:        10.10.10.10
__   Actual Remote IP Address:          10.10.10.10

Err   FC Mapping Info      Local        Remote
__    Forwarding Class     Assured      Assured
__    Profile              In           In
```

Multiple Response Connectivity Tests — When the connectivity test count is greater than one (1), a single line is displayed per SDP echo request send attempt.

The request number is a sequential number starting with 1 and ending with the last request sent, incrementing by one (1) for each request. This should not be confused with the *message-id* contained in each request and reply message.

A response message indicates the result of the message request. Following the response message is the round trip time value. If any reply is received, the round trip time is displayed.

After the last reply has been received or response timed out, a total is displayed for all messages sent and all replies received. A maximum, minimum and average round trip time is also displayed. Error response and timed out requests do not apply towards the average round trip time.

Multiple Response Round Trip Connectivity Test Output Example

```
A:router1> sdp-ping 6 resp-sdp 101size 1514 count 5
Request        Response       RTT
----------     ----------     -------
     1         Success        10ms
     2         Success        15ms
     3         Success        10ms
     4         Success        20ms
     5         Success        5ms
Sent:    5    Received:    5
Min: 5ms       Max: 20ms      Avg: 12ms
```

27.34 search

search

Syntax

search *base-dn*

no search

Context

[Tree] (config>system>security>ldap>server search)

Full Context

configure system security ldap server search

Description

This command configures the LDAP **search** command. The search *base-dn* tells the server which part of the external directory tree to search. The search DN uses the same LDAP attribute as *root-dn*. For example, to search a public-key for an SSH generated for a Nokia vendor, one might use "dc=public-key,dc=nokia,dc=com”.

The **no** version of this command removes the search DN; as such, no search is possible on the LDAP server.

Parameters

base-dn

Specifies the base domain name used in the search, up to 512 characters.

Platforms

7705 SAR Gen 2

27.35 secondary

secondary

Syntax

secondary *ip-address[/mask]* [*netmask* **]** [**broadcast** {**all-ones** | **host-ones**}] [**igp-inhibit**] [**track-srrp** *srrp-instance*]

no secondary *ip-address[/mask]*

Context

[Tree] (config>service>ies>if secondary)

Full Context

configure service ies interface secondary

Description

This command assigns a secondary IP address or IP subnet/broadcast address format to the interface.

The **no** form of this command reverts to the default.

Parameters

ip-address

The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

mask

The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that is used in a logical and function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.254.

Note:
A mask of 255.255.255.255 is reserved for system IP addresses.

netmask

Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.

broadcast

Overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed.

This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) is received by the IP interface. (Default: host-ones)

all-ones

Specifies the broadcast address used by the IP interface for this IP address is 255.255.255.255, also known as the local broadcast.

host-ones

Specifies that the broadcast address used by the IP interface for this IP address is the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface. The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

igp-inhibit

Signals that the given secondary IP interface should not be recognized as a local interface by the running IGP. For OSPF and IS-IS, this means that the specified secondary IP interfaces are not injected and used as passive interfaces and are not advertised as internal IP interfaces into the IGP's link state database. For RIP, this means that these secondary IP interfaces do not source RIP updates.

track-srrp *srrp-instance*

Specifies the SRRP instance ID that this interface route needs to track.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

secondary

Syntax

secondary *ip-address*[/*mask*] [*netmask*] [**broadcast** {**all-ones** | **host-ones**}] [**igp-inhibit**] [**track-srrp** *srrp-instance*]

no secondary *ip-address*[/*mask*]

Context

[Tree] (config>service>vprn>nw-if secondary)

[Tree] (config>service>vprn>if secondary)

Full Context

configure service vprn network-interface secondary

configure service vprn interface secondary

Description

This command assigns a secondary IP address to the interface. Up to 16 total primary and secondary IPv4 and IPv6 addresses can be assigned to network interfaces, and up to 256 to access interfaces. Each address can be configured in an IP address, IP subnet or broadcast address format.

Caution:
Configurations must not exceed 16 secondary IP addresses when IPsec, GRE, L2TPv3, or IP in IP protocols are active on an access interface.

Parameters

ip-address

The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

mask

The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.254. A mask of 255.255.255.255 is reserved for system IP addresses.

netmask

Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.

broadcast

The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is

specified for the IP address, the default value is **host-ones** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed. This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) will be received by the IP interface. *(Default: host-ones)*

all-ones

The **all-ones** keyword following the **broadcast** parameter specifies the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.

host-ones

The **host-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default used by an IP interface.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

igp-inhibit

The optional **igp-inhibit** parameter signals that the given secondary IP interface should not be recognized as a local interface by the running IGP. For OSPF and IS-IS, this means that the specified secondary IP interfaces will not be injected and used as passive interfaces and will not be advertised as internal IP interfaces into the IGP's link state database. For RIP, this means that these secondary IP interfaces will not source RIP updates.

track-srrp *srrp-instance*

Specifies the SRRP instance ID that this interface route needs to track.

Platforms

7705 SAR Gen 2

secondary

Syntax

[**no**] **secondary** *path-name*

Context

[Tree] (config>router>mpls>lsp secondary)

Full Context

configure router mpls lsp secondary

Description

This command specifies an alternative path that the LSP uses if the primary path is not available. This command is optional and is not required if the **config router mpls lsp** *lsp-name* **primary** *path-name* command is specified. After the switch over from the primary to the secondary, the system continuously tries to revert to the primary path. The switch back to the primary path is based on the **retry-timer** interval.

For RSVP-TE LSPs, up to eight secondary paths can be specified (or seven if a primary is configured). For SR-TE LSPs, up to three paths of any type (with a maximum of one primary) can be configured. By default, a secondary path is non-standby unless the **standby** keyword is configured. All non-standby secondary paths are considered equal and the first available path is used.

The system does not switch among secondary paths. The system starts the signaling (RSVP-TE) or programming (SR-TE) of all non-standby secondary paths at the same time. Retry counters are maintained for each unsuccessful attempt. After the retry limit is reached on a path, the system does not attempt to signal the path and administratively shuts down the path. The first successfully established non-standby secondary path is made the active path for the LSP.

The **no** form of this command removes the association between this *path-name* and *lsp-name*. All specific configurations for this association are deleted. The secondary path must be shut down prior to deleting it. The **no secondary** *path-name* command does not result in any action except a warning message on the console indicating that the secondary path is administratively up.

Parameters

path-name

Specifies the case-sensitive alphanumeric name label for the LSP path, up to 64 characters.

Platforms

7705 SAR Gen 2

secondary

Syntax

secondary {*ip-address/mask* | *ip-address netmask*} [**broadcast** {**all-ones** | **host-ones**}] [**igp-inhibit**] [**track-srrp** *srrp-instance*]

no secondary {*ip-address/mask* | *ip-address netmask*}

Context

[Tree] (config>router>if secondary)

Full Context

configure router interface secondary

Description

This command assigns additional IP addresses to the interface. Up to 16 total primary and secondary IPv4 and IPv6 addresses can be assigned to network interfaces, and up to 256 to access interfaces. Each address can be configured in an IP address, IP subnet, or broadcast address format.

Caution:
Configurations must not exceed 16 secondary IP addresses when IPsec, GRE, L2TPv3, or IP in IP protocols are active on an access interface.

Parameters

ip-address

Specifies the IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values 1.0.0.0 to 223.255.255.255

/

The forward slash is a parameter delimiter that separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-addr*, the "*/*" and the *mask-length* parameter. If a forward slash does not immediately follow the *ip-addr*, a dotted decimal mask must follow the prefix.

mask

Specifies the subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask* parameter. The *mask* parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 1 to 32. A mask length of 32 is reserved for system IP addresses.

Values 1 to 32

netmask

Specifies the subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical 'AND' function to derive the local subnet of the IP address. A mask of 255.255.255.255 is reserved for system IP addresses.

Values 128.0.0.0 to 255.255.255.255

broadcast

The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones,** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed.

This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) will be received by the IP interface.

all-ones

The **all-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.

host-ones

The **host-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask* with all the host bits set to binary 1. This is the default broadcast address used by an IP interface.

igp-inhibit

The secondary IP address should not be recognized as a local interface by the running IGP.

srrp-instance

Specifies the SRRP instance ID that this interface route needs to track.

Platforms

7705 SAR Gen 2

27.36 secondary-config

secondary-config

Syntax

secondary-config *file-url*

no secondary-config

Context

[Tree] (bof secondary-config)

Full Context

bof secondary-config

Description

This command specifies the name and location of the secondary configuration file.

The system attempts to use the configuration as specified in **secondary-config** if the primary config cannot be located. If the **secondary-config** file cannot be located, the system attempts to obtain the configuration from the location specified in the **tertiary-config**.

Note that if an error in the configuration file is encountered, the boot process aborts.

The **no** form of this command removes the **secondary-config** configuration.

Parameters

file-url

Specifies the secondary configuration file location, expressed as a file URL.

Values

file-url	[*local-url* | *remote-url*] (up to 180 characters)
local-url	[*cflash-id*/][*file-path*]
remote-url	[{ftp://| tftp://} *login:pswd@remote-locn*/][*file-path*]
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

27.37 secondary-dns

secondary-dns

Syntax

secondary-dns *ip-address*

no secondary-dns

Context

[Tree] (config>service>vprn>dns secondary-dns)

Full Context

configure service vprn dns secondary-dns

Description

This command configures the secondary DNS server for DNS name resolution. The secondary DNS server is used only if the primary DNS server does not respond.

DNS name resolution can be used when executing ping, traceroute, and service-ping, and also when defining file URLs. DNS name resolution is not supported when DNS names are embedded in configuration files.

The **no** form of this command removes the secondary DNS server from the configuration.

Default

no secondary-dns — No secondary DNS server is configured.

Parameters

ip-address

The IP or IPv6 address of the secondary DNS server.

Values

ipv4-address -a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x[-interface]

x:x:x:x:x:x:d.d.d.d[-interface]

x: [0 to FFFF]H

d: [0 to 255]D

interface - 32 characters max, for link local addresses.

Platforms

7705 SAR Gen 2

secondary-dns

Syntax

secondary-dns *ip-address*

no secondary-dns [*ip-address*]

Context

[Tree] (bof secondary-dns)

Full Context

bof secondary-dns

Description

This command configures the secondary DNS server for DNS name resolution. The secondary DNS server is used only if the primary DNS server does not respond.

DNS name resolution can be used when executing ping, traceroute, and service-ping, and also when defining file URLs. DNS name resolution is not supported when DNS names are embedded in configuration files.

The **no** form of this command removes the secondary DNS server from the configuration.

Default

no secondary-dns

Parameters

ip-address

Specifies the IP or IPv6 address of the secondary DNS server.

Values

ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-interface]
	x:x:x:x:x:x:d.d.d.d[-interface]
	x: [0 to FFFF]H
	d: [0 to 255]D
interface	up to 32 characters for link local addresses

Platforms

7705 SAR Gen 2

27.38 secondary-fast-retry-timer

secondary-fast-retry-timer

Syntax

secondary-fast-retry-timer *seconds*

no secondary-fast-retry-timer

Context

[Tree] (config>router>mpls secondary-fast-retry-timer)

Full Context

configure router mpls secondary-fast-retry-timer

Description

This command specifies the value used as the fast retry timer for a secondary path. If the first attempt to set up a secondary path fails due to a path error, the fast retry timer will be started for the secondary path so that the path can be retried sooner. If the next attempt also fails, further retries for the path will use the configured value for LSP retry timer.

If retry-timer for the LSP is configured to be less than the MPLS secondary-fast-retry-timer, all retries for the secondary path will use the LSP retry-timer.

The **no** form of this command reverts to the default.

Default

no secondary-fast-retry-timer

Parameters

seconds

Specifies the value (in seconds), used as the fast retry timer for a secondary path

Values 1 to 10

Platforms

7705 SAR Gen 2

27.39 secondary-image

secondary-image

Syntax

secondary-image *file-url*

no secondary-image

Context

[Tree] (bof secondary-image)

Full Context

bof secondary-image

Description

This command specifies the secondary directory location for runtime image file loading.

The system attempts to load all runtime image files configured in the **primary-image** first. If this fails, the system attempts to load the runtime images from the location configured in the **secondary-image**. If the secondary image load fails, the tertiary image specified in **tertiary-image** is used.

All runtime image files (*.tim files) must be located in the same directory.

The **no** form of this command removes the **secondary-image** configuration.

Parameters

file-url

Specifies the file URL; can be either local (this CPM) or a remote FTP server.

Values

file-url	{*local-url* | *remote-url*} (up to 180 characters)
local-url	[*cflash-id*/][*file-path*]

remote-url	[{ftp://| tftp://} *login:pswd@remote-locn/*][*file-path*]
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

27.40 secondary-ip-address

secondary-ip-address

Syntax

secondary-ip-address *ipv4-address*

no secondary-ip-address

Context

[Tree] (config>router>bgp>orr>location secondary-ip-address)

Full Context

configure router bgp optimal-route-reflection location secondary-ip-address

Description

This command specifies the secondary IP address of a reference location used for BGP optimal route reflection. Up to three IPv4 addresses and three IPv6 addresses can be specified per location.

If the TE DB is unable to find a node in its topology database that matches the primary address, then the TE DB tries to find a node with the matching secondary address. If this attempt also fails, the TE DB then tries to find a node with the matching tertiary address.

The IP addresses specified for a location should be topologically "close” to a set of clients that should all receive the same optimal path for that location.

The **no** form of this command removes the secondary IP address information.

Default

no secondary-ip-address

Parameters

ipv4-address

Specifies the secondary IPv4 address of a location, expressed in dotted decimal notation.

Values a.b.c.d

Platforms

7705 SAR Gen 2

27.41 secondary-ipv6-address

secondary-ipv6-address

Syntax

secondary-ipv6-address *ipv6-address*

no secondary-ipv6-address

Context

[Tree] (config>router>bgp>orr>location secondary-ipv6-address)

Full Context

configure router bgp optimal-route-reflection location secondary-ipv6-address

Description

This command specifies the secondary IPv6 address of a reference location used for BGP optimal route reflection. Up to three IPv4 addresses and three IPv6 addresses can be specified per location.

If the TE DB is unable find a node in its topology database that matches a primary address of the location, then it tries to find a node matching a secondary address. If this attempt also fails, the TE DB tries to find a node matching a tertiary address.

The IP addresses specified for a location should be topologically "close" to a set of clients that should all receive the same optimal path for that location.

The **no** form of this command removes the secondary IPv6 address information.

Default

no secondary-ipv6-address

Parameters

ipv6-address

Specifies the secondary IPv6 address of a location.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

27.42 secondary-ports

secondary-ports

Syntax

secondary-ports

Context

[Tree] (config>service>template>vpls-template>mac-move secondary-ports)

[Tree] (config>service>vpls>mac-move secondary-ports)

Full Context

configure service template vpls-template mac-move secondary-ports

configure service vpls mac-move secondary-ports

Description

This command opens configuration context for defining secondary vpls-ports. VPLS ports that were declared as primary prior to the execution of this command will be moved from primary port-level to secondary port-level. Changing a port to the tertiary level can only be done by first removing it from the primary port-level.

Platforms

7705 SAR Gen 2

27.43 secure-boot

secure-boot

Syntax

secure-boot

Context

[Tree] (admin>system>security secure-boot)

Full Context

admin system security secure-boot

Description

Commands in this context administratively provision secure boot.

Platforms

7705 SAR Gen 2

27.44 secure-nd

secure-nd

Syntax

[**no**] **secure-nd**

Context

[Tree] (config>service>ies>if>ipv6 secure-nd)

Full Context

configure service ies interface ipv6 secure-nd

Description

This command enables Secure Neighbor Discovery (SeND) on the IPv6 interface.

The **no** form of this command reverts to the default and disabled SeND.

Platforms

7705 SAR Gen 2

secure-nd

Syntax

[**no**] **secure-nd**

Context

[Tree] (config>service>vprn>if>ipv6 secure-nd)

Full Context

configure service vprn interface ipv6 secure-nd

Description

This command enables Secure Neighbor Discovery (SeND) on the IPv6 interface.

The **no** form of this command reverts to the default and disabled SeND.

Platforms
7705 SAR Gen 2

secure-nd

Syntax
[**no**] **secure-nd**

Context
[Tree] (config>router>if>ipv6 secure-nd)

Full Context
configure router interface ipv6 secure-nd

Description
This command enables Secure Neighbor Discovery (SeND) on the IPv6 interface.

The **no** form of this command reverts to the default and disabled SeND.

Platforms
7705 SAR Gen 2

27.45 secure-nd-export

secure-nd-export

Syntax
secure-nd-export

Context
[Tree] (admin>certificate secure-nd-export)

Full Context
admin certificate secure-nd-export

Description
This command exports IPv6 Secure Neighbor Discovery (SeND) certificates to the file cf[1..3]:\system-pki \secureNdKey in PKCS #7 DER format.

Platforms
7705 SAR Gen 2

27.46 secure-nd-import

secure-nd-import

Syntax

secure-nd-import input *url-string* **format** *input-format* [**password** *password*] [**key-rollover**]

Context

[Tree] (admin>certificate secure-nd-import)

Full Context

admin certificate secure-nd-import

Description

This command imports IPv6 Secure Neighbor Discovery (SeND) certificates from a file, and saves them to cf[1..3]:\system-pki\secureNdKey in PKCS #7 DER format.

Parameters

url-string

Specifies the name of an input file up to 99 characters.

Values

local-url	<cflash-id>\<file-path>
cflash-id	cf1:| cf2:| cf3:

input-format

Specifies the input file format.

Values pkcs12, pem, or der

password

Specifies the password to decrypt the input file if it is an encrypted PKCS#12 file.

Values 32 characters maximum

Platforms

7705 SAR Gen 2

27.47 security

security

Syntax

security

Context

[Tree] (config>system security)

Full Context

configure system security

Description

Commands in this context configure a number of central security settings, such as DDoS protection, users, authorization profiles, and certificates. Access to these commands should be restricted to highly trusted users and device administrators.

Platforms

7705 SAR Gen 2

27.48 security-association

security-association

Syntax

security-association *security-entry-id* **authentication-key** *hex-string* **encryption-key** *hex-string* **spi** *spi* **transform** *transform-id* **direction** *direction*

no security-association *security-entry-id* **direction** *direction*

Context

[Tree] (config>router>if>ipsec>ipsec-tunnel>manual-keying security-association)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>manual-keying security-association)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>manual-keying security-association)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel>manual-keying security-association)

Full Context

configure router interface ipsec ipsec-tunnel manual-keying security-association

configure service ies interface ipsec ipsec-tunnel manual-keying security-association

configure service vprn interface ipsec ipsec-tunnel manual-keying security-association

configure service vprn interface sap ipsec-tunnel manual-keying security-association

Description

This command configures the information required for manual keying SA creation.

The **no** form of this command removes the **security-association** parameters from the configuration.

Parameters

security-entry-id

Specifies the ID of an SA entry.

Values 1 to 16

authentication-key *hex-string*

Specifies an authentication key.

Values none or 0x0 to 0xFFFFFFFF...(max 128 hex nibbles)

encryption-key *hex-string*

Specifies the key used for the encryption algorithm.

Values none or 0x0 to 0xFFFFFFFF...(max 64 hex nibbles)

spi *spi*

Specifies the Security Parameter Index (SPI) used to look up the instruction to verify and decrypt the incoming IPsec packets when the direction is inbound. When the direction is outbound, the SPI that will be used in the encoding of the outgoing packets. The remote node can use this SPI to lookup the instruction to verify and decrypt the packet.

Values 256 to 16383

transform *transform-id*

Specifies the transform entry that will be used by this SA entry. This object should be specified for all the entries created which are manual SAs. If the value is dynamic, then this value is irrelevant and will be zero.

Values 1 to 2048

direction

Specifies the direction of an IPsec tunnel.

Platforms

7705 SAR Gen 2

security-association

Syntax

security-association spi *spi* **authentication-key** *authentication-key* **encryption-key** *encryption-key* [**crypto**]

no security-association spi *spi*

Context

[Tree] (config>grp-encryp>encryp-keygrp security-association)

Full Context

configure group-encryption encryption-keygroup security-association

Description

This command is used to create a security association for a specific SPI value in a key group. The command is also used to enter the authentication and encryption key values for the security association, or to delete a security association.

The SPI value used for the security association is a node-wide unique value, meaning that no two security associations in any key group on the node may share the same SPI value.

Keys are entered in cleartext. After configuration, they are never displayed in their original, cleartext form. Keys are displayed in an encrypted form, which is indicated by the system-appended **crypto** keyword when an **info** or an **admin>save** command is run. For security reasons, keys encrypted on one node are not usable on other nodes (that is, keys are not exchangeable between nodes).

The **no** form of the command removes the security association and related key values from the list of security associations for the key group. If the **no** form of the command is attempted using the same SPI value that is configured for **active-outbound-sa**, then a warning is issued and the command is blocked. If the **no** form of the command is attempted on the last SPI in the key group and the key group is configured on a service, then the command is blocked.

Parameters

spi

Specifies the SPI ID of the SPI being referenced for the security association.

Values 1 to 127

authentication-key

Specifies the authentication key for the SPI, in hexadecimal format. The number of characters in the hexadecimal string must be 64 or 128, depending on whether the authentication algorithm is set to sha256 or sha512, respectively.

encryption-key

Specifies the encryption key for the SPI, in hexadecimal format. The number of characters in the hexadecimal string must be 32 or 64, depending on whether the encryption algorithm is set to aes128 or aes256, respectively.

crypto

Displays the keys showing on the CLI **info** display in an encrypted form.

Platforms

7705 SAR Gen 2

27.49 security-parameter

security-parameter

Syntax

security-parameter *sec*

no security-parameter

Context

[Tree] (config>service>ies>if>ipv6>secure-nd security-parameter)

Full Context

configure service ies interface ipv6 secure-nd security-parameter

Description

This command configures the security parameter used in the generation of a Cryptographically Generated Address (CGA).

Parameters

sec

Specifies the security parameter.

Values 0 to 1

Platforms

7705 SAR Gen 2

security-parameter

Syntax

security-parameter *sec*

[**no**] **security-parameter**

Context

[Tree] (config>service>vprn>if>send security-parameter)

Full Context

configure service vprn interface ipv6 secure-nd security-parameter

Description

This command configures the security parameter used in the generation of a Cryptographically Generated Address (CGA).

Parameters

sec

Specifies the security parameter.

Values 0 to 1

Platforms

7705 SAR Gen 2

security-parameter

Syntax

security-parameter *sec*

no security-parameter

Context

[Tree] (config>router>if>ipv6>secure-nd security-parameter)

Full Context

configure router interface ipv6 secure-nd security-parameter

Description

This command configures the security parameter used in the generation of a Cryptographically Generated Address (CGA).

Parameters

sec

Specifies the security parameter.

Values 0 to 1

Platforms

7705 SAR Gen 2

27.50 security-policy

security-policy

Syntax

security-policy *security-policy-id* [**create**]

no security-policy *security-policy-id*

Context

[Tree] (config>router>ipsec security-policy)

[Tree] (config>service>vprn>ipsec security-policy)

Full Context

configure router ipsec security-policy

configure service vprn ipsec security-policy

Description

This command configures a security policy to use for an IPsec tunnel.

The **no** form of this command removes the security policy ID from the configuration.

Parameters

security-policy-id

specifies a value to be assigned to a security policy.

Values 1 to 32768

create

Keyword used to create the security policy instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

security-policy

Syntax

security-policy *security-policy-id* [**strict-match**]

no security-policy

Context

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel security-policy)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel security-policy)

[Tree] (config>router>if>ipsec>ipsec-tunnel security-policy)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel security-policy)

Full Context

configure service ies interface ipsec ipsec-tunnel security-policy

configure service vprn interface ipsec ipsec-tunnel security-policy

configure router interface ipsec ipsec-tunnel security-policy

configure service vprn interface sap ipsec-tunnel security-policy

Description

This command configures an IPsec security policy. The policy may then be associated with static IPsec tunnels defined in the same routing instance.

With **strict-match** parameter enabled, when a CREATE_CHILD exchange request is received for a static IPsec tunnel, and this request is not a re-key request, then ISA matches the received TSi and TSr with the configured security policy. This can be a match only when a received TS (in TSi or TSr) address range matches exactly with the subnet in a security policy entry.

If there is no match, then the setup fails, and TS_UNACCEPTABLE is sent.

If there is a match, but there is an existing CHILD_SA for the matched security policy, then the setup fails, and NO_PROPOSAL_CHOSEN.

If there is a match, and there is not CHILD_SA for the matched entry, then the subnet is sent in the matched security-policy entry as TSi and TSr, and the CHILD_SA is created.

Default

no security-policy

Parameters

security-policy-id

Specifies the IPsec security policy entry that the tunnel will use.

Values 1 to 32768

strict-match

Enables strict match of security-policy entry.

Platforms

7705 SAR Gen 2

27.51 segment

segment

Syntax

segment [*segment-id*] [**create**]

no segment *segment-id*

Context

[Tree] (conf>router>segment-routing>sr-policies>policy>seg-list segment)

Full Context

configure router segment-routing sr-policies static-policy segment-list segment

Description

This command creates the context to configure a segment inside a segment-list of a statically-defined segment routing policy candidate path.

A segment list of a statically-defined SR policy candidate path of type **sr-mpls** can only accept a segment of type **mpls-label**.

A segment list of a statically-defined SR policy candidate path of type **srv6** can only accept a segment of type **srv6-sid**. However, you can mix SRv6 segments derived from both classic SRv6 and micro-segment SRv6 locators.

The **no** form of this command deletes the segment context.

Default

no segment

Parameters

segment-id

Specifies the segment ID number.

Values 1 to 11 (for segment ID type **mpls-label**)

1 to 7 (for segment ID type **srv6-sid**) in a classic SRv6 policy candidate path

1 to 24 (for segment ID type **srv6-sid**) in a micro-segment SRv6 policy candidate path

create

Keyword used to create the list.

Platforms

7705 SAR Gen 2

27.52 segment-list

segment-list

Syntax

segment-list *segment-list*

no segment-list

Context

[Tree] (config>saa>test>type-multi-line>lsp-ping>sr-policy segment-list)

Full Context

configure saa test type-multi-line lsp-ping sr-policy segment-list

Description

This command configures the segment list ID.

The **no** form of this command removes the configuration.

Parameters

segment-list

Specifies the segment list number.

Values 1 to 32

Platforms

7705 SAR Gen 2

segment-list

Syntax

segment-list [1..32] [**create**]

no segment-list *list*

Context

[Tree] (conf>router>segment-routing>sr-policies>policy segment-list)

Full Context

configure router segment-routing sr-policies static-policy segment-list

Description

This command creates the context to configure a segment list for the statically-defined segment routing policy candidate path.

Up to 32 segment lists are supported per policy.

The **no** form of this command deletes the segment list.

Parameters

create

Keyword used to create the segment list.

Platforms

7705 SAR Gen 2

segment-list

Syntax

segment-list *segment-list-id*

no segment-list

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls>sr-policy segment-list)

Full Context

configure oam-pm session ip tunnel mpls sr-policy segment-list

Description

This command configures the segment list ID for the specific policy.

The **no** form of this command removes segment list ID.

Default

no segment-list

Parameters

segment-list-id

Specifies the segment list ID.

Values 1 to 32

Platforms

7705 SAR Gen 2

27.53 segment-routing

segment-routing

Syntax

segment-routing

Context

[Tree] (config>router>bgp segment-routing)

Full Context

configure router bgp segment-routing

Description

Commands in this context configure options related to BGP segment routing (prefix SID support).

Platforms

7705 SAR Gen 2

segment-routing

Syntax

segment-routing

no segment-routing

Context

[Tree] (config>router>isis segment-routing)

Full Context

configure router isis segment-routing

Description

Commands in this context configure segment routing parameters within a given IGP instance.

Segment routing adds to IS-IS and OSPF routing protocols the ability to perform shortest path routing and source routing using the concept of abstract segment. A segment can represent a local prefix of a node, a specific adjacency of the node (interface or next-hop), a service context, or a specific explicit path over the network. For each segment, the IGP advertises an identifier referred to as Segment ID (SID).

When segment routing is used together with MPLS data plane, the SID is a standard MPLS label. A router forwarding a packet using segment routing will thus push one or more MPLS labels.

Segment routing using MPLS labels can be used in both shortest path routing applications and in traffic engineering applications. This feature implements the shortest path forwarding application.

After segment routing is successfully enabled in the IS-IS or OSPF instance, the router will perform the following operations:

1. Advertise the Segment Routing Capability Sub-TLV to routers in all areas/levels of this IGP instance. However, only neighbors with which it established an adjacency interprets the SID or label range information and use it for calculating the label to swap to or push for a given resolved prefix SID.
2. Advertise the assigned index for each configured node SID in the new prefix SID sub-TLV with the N-flag (node-SID flag) set. Then the segment routing module programs the incoming label map (ILM) with a pop operation for each local node SID in the data path.
3. Assign and advertise automatically an adjacency SID label for each formed adjacency over a network IP interface in the new adjacency SID sub-TLV. The segment routing module programs the incoming label map (ILM) with a pop operation, in effect with a swap to an implicit null label operation, for each advertised adjacency SID.
4. Resolve received prefixes and if a prefix SID sub-TLV exists, the Segment Routing module programs the ILM with a swap operation and also an LTN with a push operation both pointing to the primary/LFA NHLFE. An SR tunnel is also added to the TTM.

When the user enables segment routing in a given IGP instance, the main SPF and LFA SPF are computed normally and the primary next-hop and LFA backup next-hop for a received prefix are added to RTM without the label information advertised in the prefix SID sub-TLV.

Platforms

7705 SAR Gen 2

segment-routing

Syntax

[**no**] **segment-routing**

Context

[Tree] (config>router>ospf segment-routing)

Full Context

configure router ospf segment-routing

Description

Commands in this context configure segment routing parameters within an IGP instance.

Segment routing adds to IS-IS, OSPF, or OSPF3 routing protocols the ability to perform shortest path routing and source routing using the concept of abstract segment. A segment can represent a local prefix of a node, a specific adjacency of the node (interface or next hop), a service context, or a specific explicit path over the network. For each segment, the IGP advertises an identifier referred to as a segment ID (SID).

When segment routing is used together with the MPLS data plane, the SID is a standard MPLS label. A router forwarding a packet using segment routing will thus push one or more MPLS labels.

Segment routing using MPLS labels can be used in both shortest path routing applications and traffic engineering applications. This feature implements the shortest path forwarding application.

After segment routing is successfully enabled in the IS-IS, OSPF, or OSPF3 instance, the router will perform the following operations:

- Advertise the Segment Routing Capability sub-TLV to routers in all areas or levels of the IGP instance. However, only neighbors with which the IGP instance established an adjacency will interpret the SID and label range information and use it for calculating the label to swap to or push for a particular resolved prefix SID.
- Advertise the assigned index for each configured node SID in the new prefix SID sub-TLV with the N-flag (node SID flag) set. The segment routing module then programs the incoming label map (ILM) with a pop operation for each local node SID in the data path.
- Automatically assign and advertise an adjacency SID label for each formed adjacency over a network IP interface in the new adjacency SID sub-TLV. The segment routing module programs the incoming label map (ILM) with a pop operation, in effect with a swap to an implicit null label operation, for each advertised adjacency SID.
- Resolve received prefixes, and if a prefix SID sub-TLV exists, the segment routing module programs the ILM with a swap operation and programs an LSP ID to NHLFE (LTN) with a push operation, both pointing to the primary/LFA NHLFE. An SR tunnel is also added to the TTM.

When the user enables segment routing in an IGP instance, the main SPF and LFA SPF are computed normally and the primary next hop and LFA backup next hop for a received prefix are added to the RTM without the label information advertised in the prefix SID sub-TLV.

Platforms

7705 SAR Gen 2

segment-routing

Syntax

segment-routing

Context

[Tree] (config>router segment-routing)

Full Context

configure router segment-routing

Description

This command creates a context to configure protocol-independent parameters relating to segment routing.

Platforms

7705 SAR Gen 2

27.54 sel-mcast-advertisement

sel-mcast-advertisement

Syntax

[**no**] **sel-mcast-advertisement**

Context

[Tree] (config>service>vpls>bgp-evpn sel-mcast-advertisement)

Full Context

configure service vpls bgp-evpn sel-mcast-advertisement

Description

This command enables the advertisement of BGP EVPN Selective Multicast Ethernet Tag (SMET) routes.

The **no** form of this command disables the advertisement of BGP EVPN SMET routes.

Default

no sel-mcast-advertisement

Platforms

7705 SAR Gen 2

27.55 selection-criteria

selection-criteria

Syntax

selection-criteria [**best-port** | **highest-count** | **highest-weight**] [**slave-to-partner**] [**subgroup-hold-time** *hold-time*]

no selection-criteria

Context

[Tree] (config>lag selection-criteria)

Full Context

configure lag selection-criteria

Description

This command specifies which selection criteria should be used to select the active sub-group. If there is a tie for highest-count or highest-weight, the LAG will prefer the port with the lowest priority. If that does not break the tie, the currently active subgroup will stay active (that is, non-revertive behavior).

The **no** form of this command reverts to the default value.

Default

selection-criteria highest-count

Parameters

highest-count

Selects a sub-group with the highest number of eligible members as an active sub-group (not applicable to "power-off" mode of operations).

highest-weight

Selects a sub-group with the highest aggregate weight as an active subgroup (not applicable to "power-off" mode of operations). Aggregate weight is calculated as the sum of (65535 - port priority) all ports within a sub-group.

best-port

Selects a sub-group containing the port with highest priority port as an active subgroup. In case of equal port priorities, the sub-group containing the port with the lowest port-id is chosen.

slave-to-partner

The **slave-to-partner** keyword specifies that it, together with the selection criteria, should be used to select the active sub-group. An eligible member is a LAG-member link which can potentially become active. This means it is operationally up (not disabled) for use by the remote side. The **slave-to-partner** keyword can be used to control whether or not this latter condition is taken into account.

hold-time

Applicable with LACP enabled. Specifies the optional delay timer for switching to a newly selected active sub-group from the existing active sub-group. The timer delay applies only if the existing sub-group remains operationally up.

Values

not specified	Equivalent to specifying a value of 0. Specifies no delay and to switchover immediately to a new candidate active sub-group.
0 to 2000	Integer specifying the timer value in 10ths of a second.
infinite	Do not switchover from existing active sub-group if the subgroup remains UP. Manual switchover possible using tools perform lag force command.

Platforms

7705 SAR Gen 2

27.56 selective-label-ip

selective-label-ip

Syntax

selective-label-ip {**no-install** | **route-table-install-only**}

no selective-label-ip

Context

[Tree] (config>router>bgp selective-label-ip)

Full Context

configure router bgp selective-label-ip

Description

This command configures **selective-label-ip** for the BGP level.

The **no-install** option conserves labeled route table space on BGP-LU **next-hop-self** route reflectors. This option causes BGP-LU routes to be reflected downstream via the ABR with the **next-hop-self** update. BGP-LU routes are not installed to local MPLS tables or routing tables for use by local services.

The **route-table-install-only** option conserves labeled route table space on BGP-LU **next-hop-self** route reflectors and allows these routes to be used for IP transport, unlike the **no-install** option. When the **route-table-install-only** option is used, learned BGP-LU routes are also reflected downstream via the ABR with the **next-hop-self** update. BGP-LU routes are not installed to local MPLS tables for use by local services. These routes are installed to the RTM and used for the best route selection process.

Note: If local services need to use BGP-LU routes, the **no-install** and **route-table-install-only** options should not be used.

The default **no** form of this command installs BGP-LU routes to the datapath for local services and makes them available to the RTM for IP next-hop selection.

Default

no selective-label-ip

Parameters

no-install

Specifies that BGP-LU routes are not installed to local MPLS tables or routing tables.

route-table-install-only

Specifies the installation of BGP-LU routes to the RTM. BGP-LU routes are not installed to local MPLS tables for use by local services.

Platforms

7705 SAR Gen 2

27.57 selective-label-ip-prioritization

selective-label-ip-prioritization

Syntax

[**no**] **selective-label-ip-prioritization**

Context

[Tree] (config>router>bgp selective-label-ip-prioritization)

Full Context

configure router bgp selective-label-ip-prioritization

Description

This command enables selective-label IP prioritization for BGP labeled IPv4 and IPv6 routes.

When this command is configured, every received labeled IPv4 and IPv6 route that is potentially usable by a local service is automatically prioritized for fast control plane reconvergence. When the reachability of a BGP next-hop changes, these labeled IPv4 and IPv6 routes are updated into the route table first, along with other routes manually tagged as high priority by import policies.

A /32 or /128 labeled unicast route (and associated BGP-LU tunnel) is determined to be potentially usable by a local service if one of the following conditions is met:

- the route matches the far-end address of a user-provisioned SDP of an Layer 2 service and the SDP is configured to use BGP tunnels as transport
- the route matches the BGP next-hop address of a BGP-EVPN or IP VPN route, and this VPN route is either imported into a local service or readvertised by the router acting as a next-hop-self route-reflector or a model-B ASBR

The **no** form of this command disables selective-label IP prioritization for BGP.

Default

no selective-label-ip-prioritization

Platforms

7705 SAR Gen 2

27.58 selective-label-ipv4-install

selective-label-ipv4-install

Syntax

[**no**] **selective-label-ipv4-install**

Context

[Tree] (config>router>bgp>group>neighbor selective-label-ipv4-install)

[Tree] (config>router>bgp>group selective-label-ipv4-install)

[Tree] (config>router>bgp selective-label-ipv4-install)

Full Context

configure router bgp group neighbor selective-label-ipv4-install

configure router bgp group selective-label-ipv4-install

configure router bgp selective-label-ipv4-install

Description

This command enables selective download for BGP label-ipv4 routes.

When this command is configured so that it applies to a BGP session, label-ipv4 routes received on this session are marked as invalid if they are not needed for any eligible service. A /32 label-ipv4 route is determined to be required if one of the following applies:

1. It matches the far-end address of a manually configured or auto-created SDP Layer 2 VLL or VPLS service and the SDP is configured to use BGP tunnels as transport.
2. It matches the IPv4 BGP next hop of a BGP-EVPN route and this EVPN route is either imported into a VPLS service or re-advertised by the router acting as a next-hop-self route-reflector or a model-B ASBR.
3. It matches the IPv4 BGP next hop of a VPN-IPv4 route and this VPN-IP route is either imported into a VPRN service or re-advertised by the router acting as a next-hop-self route-reflector or a model-B ASBR.
4. It matches the IPv4 address in the IPv4-mapped IPv6 address of a VPN IPv6 route and this VPN-IP route is either imported into a VPRN service or re-advertised by the router acting as a next-hop-self route-reflector or a model-B ASBR.

The **no** form of this command at the top (**config>router>bgp**) level disables the selective installation functionality. The **no** form of this command at the **group** or **neighbor** level causes the setting to be inherited from a higher level configuration.

Default

no selective-label-ipv4-install

Platforms

7705 SAR Gen 2

27.59 selective-learned-fdb

selective-learned-fdb

Syntax

[**no**] **selective-learned-fdb**

Context

[Tree] (config>service>vpls selective-learned-fdb)

Full Context

configure service vpls selective-learned-fdb

Description

This command determines which line cards FDB entries are allocated on for MAC addresses in the VPLS service in which the command is configured.

By default, FDB entries for MAC addresses in VPLS services are allocated on all line cards in the system. Enabling **selective-learned-fdb** causes FDB entries to be allocated only on the line cards on which the service has a configured object, which includes all line cards:

- on which a SAP is configured
- which have ports configured in a LAG SAP
- which have ports configured in an Ethernet tunnel SAP
- which have ports configured on a network interface (which also may be on a LAG) when the service has a mesh or spoke-SDP, VXLAN or EVPN-MPLS configured

Only MAC addresses with a type "L" or "Evpn" in the **show** output displaying the FDB can be allocated selectively, unless a MAC address configured as a conditional static MAC address is learned dynamically on an object other than its monitored object; this can be displayed with type "L" or "Evpn" but is allocated as global because of the conditional static MAC configuration.

The **no** form of this command returns the FDB MAC address entry allocation mode to its default where FDB entries for MAC addresses are allocated on all line cards in the system.

Default

no selective-learned-fdb

Platforms

7705 SAR Gen 2

27.60 send

send

Syntax

send {**broadcast** | **multicast** | **none** | **version-1** | **both**}

no send

Context

[Tree] (config>service>vprn>ripng>group send)

[Tree] (config>service>vprn>rip>group send)

[Tree] (config>service>vprn>ripng send)

[Tree] (config>service>vprn>rip send)

[Tree] (config>service>vprn>ripng>group>neighbor send)

[Tree] (config>service>vprn>rip>group>neighbor send)

Full Context

configure service vprn ripng group send

configure service vprn rip group send

configure service vprn ripng send

configure service vprn rip send

configure service vprn ripng group neighbor send

configure service vprn rip group neighbor send

Description

This command configures the type of RIP messages sent to RIP neighbors. This control can be issued at the global, group or interface level. The default behavior sends RIPv2 messages with the multicast (224.0.0.9) destination address.

If **version-1** is specified, the router only listens for and accepts packets sent to the broadcast address.

The **no** form of this command resets the type of messages sent back to the default value.

Default

no send

Parameters

broadcast

Send RIPv2 formatted messages to the broadcast address.

multicast

Send RIPv2 formatted messages to the multicast address.

none

Do not send any RIP messages (i.e. silent listener).

version-1

Send RIPv1 formatted messages to the broadcast address.

both

Send both RIP v1 & RIP v2 updates to the broadcast address.

Platforms

7705 SAR Gen 2

send

Syntax

send

Context

[Tree] (config>system>security>keychain>direction>uni send)

Full Context

configure system security keychain direction uni send

Description

This command specifies the send nodal context to sign TCP segments that are being sent by the router to another device.

Platforms

7705 SAR Gen 2

send

Syntax

send *option-number*

no send

Context

[Tree] (config>system>security>keychain>tcp-option-number send)

Full Context

configure system security keychain tcp-option-number send

Description

This command configures the TCP option number accepted in TCP packets sent.

Default

send 254

Parameters

option-number

Specifies an enumerated integer that indicates the TCP option number to be used in the TCP header.

Values 253, 254, tcp-ao

Platforms

7705 SAR Gen 2

send

Syntax

send {**broadcast** | **multicast** | **none** | **version-1**}

no send

Context

[Tree] (config>router>rip>group send)

[Tree] (config>router>rip send)

[Tree] (config>router>rip>group>neighbor send)

Full Context

configure router rip group send

configure router rip send

configure router rip group neighbor send

Description

This command specifies the type of RIP messages sent to RIP neighbors.

If **version-1** is specified, the router need only listen for and accept packets sent to the broadcast address.

This control can be issued at the global, group or interface level.

The **no** form of the command reverts to the default value.

Default

send version-1

Parameters

broadcast

Specifies send RIPv2 formatted messages to the broadcast address.

multicast

Specifies send RIPv2 formatted messages to the multicast address.

none

Specifies not to send any RIP messages (i.e. silent listener).

version-1

Specifies send RIPv1 formatted messages to the broadcast address.

Platforms

7705 SAR Gen 2

send

Syntax

send {**none** | **ripng** | **unicast**}

no send

Context

[Tree] (config>router>ripng>group>neighbor send)

[Tree] (config>router>ripng>group send)

[Tree] (config>router>ripng send)

Full Context

configure router ripng group neighbor send

configure router ripng group send

configure router ripng send

Description

This command specifies if RIPng are sent to RIP neighbors or not and what type of IPv6 address is to be used to deliver the messages.

This control can be issued at the global, group or interface level.

The **no** form of the command reverts to the default value.

Default

send ripng

Parameters

ripng

Specifies RIPng messages to be sent to the standard multicast address (FF02::9).

none

Specifies not to send any RIPng messages (i.e. silent listener).

unicast

Specifies to send RIPng updates as unicast messages to the defined unicast address configured through the **unicast-address** command. This option is only allowed within the neighbor context.

Platforms

7705 SAR Gen 2

27.61 send-chain

send-chain

Syntax

[**no**] **send-chain**

Context

[Tree] (config>ipsec>cert-profile>entry send-chain)

Full Context

configure ipsec cert-profile entry send-chain

Description

Commands in this context configure the send-chain in the **cert-profile entry**.

The configuration of this command is optional, by default system will only send the certificate specified by **cert** command in the selected entry to the peer. This command allows system to send additional CA certificates to the peer. These additional CA certificates must be in the certificate chain of the certificate specified by the **cert** command in the same entry.

Platforms

7705 SAR Gen 2

send-chain

Syntax

[**no**] **send-chain**

Context

[Tree] (config>system>security>tls>cert-profile>entry send-chain)

Full Context

configure system security tls cert-profile entry send-chain

Description

This command enables the sending of certificate authority (CA) certificates, and enters the context to configure send-chain information.

By default, the system only sends the TLS server certificate or TLS client certificate specified by the **cert** command. If CA certificates are to be sent using send-chain, they must be in the chain of certificates specified by the **config**>**system**>**security**>**pki**>**ca-profile** command. The specification of the send-chain is not necessary for a working TLS profile if the TLS peer has the CA certificate used to sign the server or client certificate in its own trust anchor.

For example, given a TLS client running on SR OS, the ROOT CA certificate resides on the TLS server, but the subsequent SUB-CA certificate needed to complete the chain resides within SR OS. The **send-chain** command allows these SUB-CA certificates to be sent from SR OS to the peer to be authenticated using the ROOT CA certificate that resides on the peer.

The **no** form of the command disables the send-chain.

Default

no send-chain

Platforms

7705 SAR Gen 2

27.62 send-count

send-count

Syntax

send-count *send-count*

no send-count

Context

[Tree] (config>saa>test>type-multi-line>lsp-ping send-count)

[Tree] (config>saa>test>type-multi-line>lsp-ping>sr-policy send-count)

Full Context

configure saa test type-multi-line lsp-ping send-count

configure saa test type-multi-line lsp-ping sr-policy send-count

Description

This command configures the number of messages to send. The *send-count* value is used to override the default number of message requests sent. Each message request must either time out or receive a reply before the next message request is sent. The message *interval* value must be expired before the next message request is sent.

The **no** form of this command reverts to the default value.

Default

send-count 1

Parameters

send-count

Specifies the send count in number of packets.

Values 1 to 100

Default 1

Platforms

7705 SAR Gen 2

27.63 send-default

send-default

Syntax

send-default [**ipv4**] [**ipv6**] [**export-policy** *export-policy*]

no send-default

Context

[Tree] (config>router>bgp send-default)

[Tree] (config>router>bgp>group>neighbor send-default)

[Tree] (config>router>bgp>group send-default)

Full Context

configure router bgp send-default

configure router bgp group neighbor send-default

configure router bgp group send-default

Description

This command enables the advertisement of a default route. When this command is configured to apply to an IBGP or EBGP session, the default route for IPv4 or IPv6 is automatically added to the Adj_RIB-OUT of that peer. The advertised default routes are unrelated to any default routes installed in the FIB of the local router.

If a BGP export policy allows an active default route in the FIB of the local router to be advertised and conflict with this command, the artificially generated default route overrides the advertisement of the installed default route.

The artificially generated default route is not matched by BGP export policies. To modify its attributes or decide whether it should be advertised (based on a conditional expression), a route policy must be created and referenced by the **export-policy** parameter. Only conditional entries with an action and no from or to criteria are parsed. If there are no such entries, only the default action is applied.

The **no** form of this command restores the default behavior. At the group and neighbor levels, the default behavior is to inherit the configuration from a higher level. At the instance level, the default behavior is to neither generate nor inject a default route.

Default

no send-default

Parameters

ipv4

Generates and advertises an IPv4 default route (0/0).

ipv6

Generates and advertises an IPv6 default route (::/0).

export-policy

Specifies the name of a route policy, up to 64 characters. Only the route modifications in the matching conditional-expression entry or the default action are applied. These modifications change the attributes of the advertised default routes.

Platforms

7705 SAR Gen 2

27.64 send-flush-on-failure

send-flush-on-failure

Syntax

[**no**] **send-flush-on-failure**

Context

[Tree] (config>service>vpls send-flush-on-failure)

Full Context

configure service vpls send-flush-on-failure

Description

This command enables sending out flush-all-from-me messages to all LDP peers included in affected VPLS, in the event of physical port failures or "operationally down" events of individual SAPs. This feature provides an LDP-based mechanism for recovering a physical link failure in a dual-homed connection to a VPLS service. This method provides an alternative to RSTP solutions where dual homing redundancy and recovery, in the case of link failure, is resolved by RSTP running between a PE router and CE devices. If the endpoint is configured within the VPLS and send-flush-on-failure is enabled, flush-all-from-me messages will be sent out only when all spoke-SDPs associated with the endpoint go down.

This feature cannot be enabled on management VPLS.

Default

no send-flush-on-failure

Platforms

7705 SAR Gen 2

27.65 send-idr-after-eap-success

send-idr-after-eap-success

Syntax

[**no**] **send-idr-after-eap-success**

Context

[Tree] (config>ipsec>ike-policy send-idr-after-eap-success)

Full Context

configure ipsec ike-policy send-idr-after-eap-success

Description

This command enables the system to add the Identification Responder (IDr) payload in the last IKE authentication response after an Extensible Authentication Protocol (EAP) Success packet is received. When disabled, the system will not include IDr payload.

The **no** form of this command disables sending the IDr payload in the last IKE.

Default

send-idr-after-eap-success

Platforms

7705 SAR Gen 2

27.66 send-orf

send-orf

Syntax

send-orf [*comm-id*]

no send-orf [*comm-id*]

Context

[Tree] (config>router>bgp>group>outbound-route-filtering>extended-community send-orf)

[Tree] (config>router>bgp>group>neighbor>outbound-route-filtering>extended-community send-orf)

[Tree] (config>router>bgp>outbound-route-filtering>extended-community send-orf)

Full Context

configure router bgp group outbound-route-filtering extended-community send-orf

configure router bgp group neighbor outbound-route-filtering extended-community send-orf

configure router bgp outbound-route-filtering extended-community send-orf

Description

This command instructs the router to negotiate the send capability in the BGP outbound route filtering (ORF) negotiation with a peer.

This command also causes the router to send a community filter, prefix filter, or AS path filter configured as an inbound filter on the BGP session to its peer as an ORF Action ADD.

The **no** form of this command causes the router to remove the send capability in the BGP ORF negotiation with a peer.

The **no** form also causes the router to send an ORF remove action for a community filter, prefix filter, or AS path filter configured as an inbound filter on the BGP session to its peer.

If the *comm-id* parameters are not exclusively route target communities then the router will extract appropriate route targets and use those. If, for some reason, the *comm-id* parameters specified contain no route targets, then the router will not send an ORF.

Default

no send-orf

Parameters

comm-id

Specifies up to 32 community policies, which must consist exclusively of route target extended communities. If it is not specified, then the ORF policy is automatically generated from configured route target lists, accepted client route target ORFs and locally configured route targets.

Values [**target:** {*ip-address*:*comm-val* | *2byte-asnumber*:*ext-comm-val* | *4byte-asnumber*:*comm-val*}

where:

- *ip-address* — a.b.c.d
- *comm-val* — 0 to 65535
- *2byte-asnumber* — 0 to 65535
- *ext-comm-val* — 0 to 4294967295
- *4byte-asnumber* — 0 to 4294967295

Platforms

7705 SAR Gen 2

27.67 send-queries

send-queries

Syntax

[**no**] **send-queries**

Context

[Tree] (config>service>vpls>sap>igmp-snooping send-queries)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping send-queries)

[Tree] (config>service>vpls>mesh-sdp>mld-snooping send-queries)

[Tree] (config>service>vpls>sap>mld-snooping send-queries)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping send-queries)

[Tree] (config>service>vpls>spoke-sdp>mld-snooping send-queries)

Full Context

configure service vpls sap igmp-snooping send-queries

configure service vpls spoke-sdp igmp-snooping send-queries

configure service vpls mesh-sdp mld-snooping send-queries

configure service vpls sap mld-snooping send-queries

configure service vpls mesh-sdp igmp-snooping send-queries

configure service vpls spoke-sdp mld-snooping send-queries

Description

This command specifies whether to send IGMP general query messages on the SAP or SDP.

When send-queries is configured, all type of queries generate ourselves are of the configured version. If a report of a version higher than the configured version is received, the report will get dropped and a new wrong version counter will get incremented. If send-queries is not configured, the version command has no effect. The version used will be the version of the querier. This implies that, for example, when we have a v2 querier, we will never send out a v3 group or group-source specific query when a host wants to leave a certain group.

If **mrouter-port** is enabled on this SAP or spoke SDP, the **send-queries** command parameter cannot be set.

The **no** form of this command disables the IGMP general query messages.

Default

no send-queries

Platforms

7705 SAR Gen 2

send-queries

Syntax

[**no**] **send-queries**

Context

[Tree] (config>service>pw-template>igmp-snooping send-queries)

Full Context

configure service pw-template igmp-snooping send-queries

Description

This command specifies whether to send IGMP general query messages.

When **send-queries** is configured, all type of queries generated are of the configured version. If a report of a version higher than the configured version is received, the report will get dropped and a new wrong version counter will get incremented.

If **send-queries** is not configured, the version command has no effect. The version used on that SAP or SDP will be the version of the querier. This implies that, for example, when we have a v2 querier, we will never send out a v3 group or group-source specific query when a host wants to leave a certain group.

Default

no send-queries

Platforms

7705 SAR Gen 2

27.68 send-refresh

send-refresh

Syntax

send-refresh *seconds*

no send-refresh

Context

[Tree] (config>service>vpls>proxy-nd send-refresh)

[Tree] (config>service>vpls>proxy-arp send-refresh)

Full Context

configure service vpls proxy-nd send-refresh

configure service vpls proxy-arp send-refresh

Description

If enabled, this command will make the system send a refresh at the configured time. A refresh message is an ARP-request message that uses 0s as sender's IP for the case of a proxy-ARP entry. For proxy-ND entries, a refresh is a regular NS message using the chassis-mac as MAC source-address.

Default

no send-refresh

Parameters

seconds

Specifies the send-refresh in seconds.

Values 120 to 86400

Platforms

7705 SAR Gen 2

27.69 send-release

send-release

Syntax

[**no**] **send-release**

Context

[Tree] (config>service>ies>if>sap>ipsec-gw>dhcp send-release)

[Tree] (config>service>ies>if>sap>ipsec-gw>dhcp6 send-release)

[Tree] (config>service>vprn>if>sap>ipsec-gw>dhcp6 send-release)

[Tree] (config>service>vprn>if>sap>ipsec-gw>dhcp send-release)

Full Context

configure service ies interface sap ipsec-gw dhcp send-release

configure service ies interface sap ipsec-gw dhcp6 send-release

configure service vprn interface sap ipsec-gw dhcp6 send-release

configure service vprn interface sap ipsec-gw dhcp send-release

Description

This command enables the system to send a DHCPv4/v6 release message when the IPsec tunnel is removed.

Default

no send-release

Platforms

7705 SAR Gen 2

27.70 send-to-ebgp

send-to-ebgp

Syntax

send-to-ebgp *family* [*family*]

no send-to-ebgp

Context

[Tree] (config>service>vprn>bgp>group>link-bandwidth send-to-ebgp)

[Tree] (config>service>vprn>bgp>group>neighbor>link-bandwidth send-to-ebgp)

Full Context

configure service vprn bgp group link-bandwidth send-to-ebgp

configure service vprn bgp group neighbor link-bandwidth send-to-ebgp

Description

This command configures BGP to allow link-bandwidth extended community to be sent in routes advertised to EBGP peers in the scope of the command, as long the routes belong to one of the listed address families.

The link-bandwidth extended community is encoded as a non-transitive type. This means that by default it should not be attached to any route advertised to an EBGP peer and it should be discarded when received in any route from an EBGP peer. This command overrides the standard behavior.

Up to three families may be configured.

The **no** form of this command restores the default behavior of stripping the link-bandwidth extended community from any route advertised to an EBGP peer.

Default

no send-to-ebgp

Parameters

family

Specifies the address families for which receiving the link-bandwidth extended community from EBGP peers should be supported.

Values ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes.

label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes.

ipv6 — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes.

Platforms

7705 SAR Gen 2

send-to-ebgp

Syntax

send-to-ebgp *family* [*family*]

no send-to-ebgp

Context

[Tree] (config>router>bgp>group>neighbor>link-bandwidth send-to-ebgp)

[Tree] (config>router>bgp>group>link-bandwidth send-to-ebgp)

Full Context

configure router bgp group neighbor link-bandwidth send-to-ebgp

configure router bgp group link-bandwidth send-to-ebgp

Description

This command configures BGP to allow link-bandwidth extended community to be sent in routes advertised to EBGP peers in the scope of the command, as long the routes belong to one of the listed address families.

The link-bandwidth extended community is encoded as a non-transitive type. This means that by default it should not be attached to any route advertised to an EBGP peer and it should be discarded when received in any route from an EBGP peer. This command overrides the standard behavior.

Up to six families may be configured.

The **no** form of this command restores the default behavior of stripping the link-bandwidth extended community from any route advertised to an EBGP peer.

Default

no send-to-ebgp

Parameters

family

Specifies the address families for which receiving the link-bandwidth extended community from EBGP peers should be supported.

Values ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes.

label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes.

vpn-ipv4 — Adds a link-bandwidth extended community to IPv4 VPN (SAFI 128) routes.

ipv6 — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes.

label-ipv6 — Adds a link-bandwidth extended community to labeled-unicast IPv6 routes.

vpn-ipv6 — Adds a link-bandwidth extended community to IPv6 VPN (SAFI 128) routes.

Platforms

7705 SAR Gen 2

27.71 send-tunnel-encap

send-tunnel-encap

Syntax

send-tunnel-encap [**mpls**] [**mplsoudp**]

no send-tunnel-encap

Context

[Tree] (config>service>epipe>bgp-evpn>mpls send-tunnel-encap)

[Tree] (config>service>vpls>bgp-evpn>mpls send-tunnel-encap)

[Tree] (config>service>vprn>bgp-evpn>mpls send-tunnel-encap)

Full Context

configure service epipe bgp-evpn mpls send-tunnel-encap

configure service vpls bgp-evpn mpls send-tunnel-encap

configure service vprn bgp-evpn mpls send-tunnel-encap

Description

This command configures the encapsulation to be advertised with the EVPN routes for the service. The encapsulation is encoded in RFC 5512-based tunnel encapsulation extended communities.

When used in the **bgp-evpn**>**mpls** context, the supported options are none (**no send-tunnel-encap**), **mpls**, **mplsoudp** or both.

When used in the **bgp-evpn**>**vxlan** context, the supported options are **send-tunnel-encap** (the router signals a VXLAN value) or **no send-tunnel-encap** (no encapsulation extended community is sent).

Default

send-tunnel-encap mpls (in the **config**>**service**>**vpls**>**bgp-evpn**>**mpls** context)

send-tunnel-encap (in the **config**>**service**>**vpls**>**bgp-evpn**>**vxlan** context)

Parameters

mpls

Specifies the MPLS-over-UDP encapsulation value in the RFC 5512 encapsulation extended community.

mplsoudp

Specifies the MPLS encapsulation value in the RFC 5512 encapsulation extended community.

Platforms

7705 SAR Gen 2

27.72 sensor-group

sensor-group

Syntax

sensor-group *name* [**create**]

no sensor-group *name*

Context

[Tree] (config>system>telemetry>sensor-groups sensor-group)

Full Context

configure system telemetry sensor-groups sensor-group

Description

Commands in this context configure sensor-related commands.

The **no** form of this command removes the configuration.

Parameters

name

Specifies the sensor group name, up to 32 characters.

create

Keyword used to create a sensor group.

Platforms

7705 SAR Gen 2

sensor-group

Syntax

sensor-group *name*

no sensor-group

Context

[Tree] (config>system>telemetry>persistent-subscriptions>subscription sensor-group)

Full Context

configure system telemetry persistent-subscriptions subscription sensor-group

Description

This command assigns an existing sensor group to the specified persistent subscription. If no valid paths exist in the sensor group, the configuration is accepted; however, no gRPC connection is established when persistent subscription is activated.

The **no** form of this command removes the configuration.

Parameters

name

Specifies the sensor group name, up to 32 characters.

Platforms

7705 SAR Gen 2

27.73 sensor-groups

sensor-groups

Syntax

sensor-groups

Context

[Tree] (config>system>telemetry sensor-groups)

Full Context

configure system telemetry sensor-groups

Description

Commands in this context configure a sensor group.

Platforms

7705 SAR Gen 2

27.74 serial-notify

serial-notify

Syntax

[**no**] **serial-notify**

Context

[Tree] (debug>router>rpki-session>packet serial-notify)

Full Context

debug router rpki-session packet serial-notify

Description

This command enables debugging for serial notify RPKI packets.

The **no** form of this command disables debugging for serial notify RPKI packets.

Platforms

7705 SAR Gen 2

27.75 serial-query

serial-query

Syntax

[**no**] **serial-query**

Context

[Tree] (debug>router>rpki-session>packet serial-query)

Full Context

debug router rpki-session packet serial-query

Description

This command enables debugging for serial query RPKI packets.

The **no** form of this command disables debugging for serial query RPKI packets.

Platforms

7705 SAR Gen 2

27.76 server

server

Syntax

server *ipv6z-address* [*ipv6z-address*]

no server [*ipv6z-address*]

Context

[Tree] (config>service>ies>if>ipv6>dhcp6-relay server)

Full Context

configure service ies interface ipv6 dhcp6-relay server

Description

This command specifies a list of servers where DHCP6 requests are forwarded. The list of servers can be entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP6 relay to work. If there are multiple servers, the request is forwarded to all servers in the list.

The **no** form of this command reverts to the default.

Parameters

ipv6z-address

Specifies up to eight non-global IPv4 addresses including a zone index as defined by the InetAddressIPv4z textual convention.

Values

ipv6z-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

Platforms

7705 SAR Gen 2

server

Syntax

server *server1* [*server2*]

Context

[Tree] (config>service>vprn>if>dhcp server)

[Tree] (config>service>ies>if>dhcp server)

Full Context

configure service vprn interface dhcp server

configure service ies interface dhcp server

Description

This command specifies a list of servers where requests are forwarded. The list of servers can be entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all servers in the list.

There can be a maximum of 8 DHCP servers configured.

The **no** form of this command reverts to the default.

Parameters

server

Specifies up to eight DHCP server IP addresses.

Platforms

7705 SAR Gen 2

server

Syntax

server *server-index* **name** *server-name*

no server *server-index*

Context

[Tree] (config>aaa>radius-srv-plcy>servers server)

Full Context

configure aaa radius-server-policy servers server

Description

This command adds a RADIUS server.

The **no** form of this command removes a RADIUS server.

Parameters

index

Specifies the index for the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index.

Values 1 to 5

server-name

Specifies the server name, up to 32 characters.

Platforms

7705 SAR Gen 2

server

Syntax

server *server-name* [**address** *ip-address*] [**secret** *key*] [**hash** | **hash2**| **custom**] [**create**]

no server *server-name*

Context

[Tree] (config>router>radius-server server)

[Tree] (config>service>vprn>radius-server server)

Full Context

configure router radius-server server

configure service vprn radius-server server

Description

This command either specifies an external RADIUS server in the corresponding routing instance or enters configuration context of an existing server. The configured server could be referenced in the radius-server-policy.

The **no** form of this command removes the parameters from the server configuration.

Parameters

server-name

Specifies the name of the external RADIUS server.

ip-address

Specifies the IPv4 or IPv6 IP address of the external RADIUS server.

key

Specifies the shared secret key of the external RADIUS server, up to 64 characters.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

Platforms

7705 SAR Gen 2

server

Syntax

server *index* **address** *ip-address* **secret** *key* [{**hash** | **hash2** | **custom**}] [**port** *port*]

no server *index*

Context

[Tree] (config>service>vprn>aaa>rmt-srv>tacplus server)

[Tree] (config>system>security>tacplus server)

Full Context

configure service vprn aaa remote-servers tacplus server

configure system security tacplus server

Description

This command adds a TACACS+ server and configures the TACACS+ server IP address, index, and key values.

Up to five TACACS+ servers can be configured at any one time. TACACS+ servers are accessed in order from lowest index to the highest index for authentication requests.

The **no** form of this command removes the server from the configuration.

Default

No TACACS+ servers are configured.

Parameters

index

Specifies the index for the TACACS+ server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from the lowest index to the highest index.

Values 1 to 5

ip-address

Specifies the IP address of the TACACS+ server. Two TACACS+ servers cannot have the same IP address. An error message is generated if the server address is a duplicate.

Values

ipv4-address	a.b.c.d (host bits must be 0)
ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0..FFFF]H
	d: [0..255]D

key

Specifies the secret key, up to 128 characters, for access to the TACACS+ server. This secret key must match the password on the TACACS+ server.

Values Up to 128 characters in length.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

port

Specifies the port ID.

Values 0 to 65535

Platforms

7705 SAR Gen 2

server

Syntax

server *ipv6-address* [*ipv6-address*]

no server

Context

[Tree] (config>service>vprn>router-advert>dns-options server)

[Tree] (config>service>vprn>router-advert>if>dns-options server)

Full Context

configure service vprn router-advertisement dns-options server

configure service vprn router-advertisement interface dns-options server

Description

This command specifies the IPv6 DNS servers to include in the RDNSS option in Router Advertisements. When specified at the router advertisement level this applies to all interfaces that have **include-dns** enabled, unless the interfaces have more specific **dns-options** configured.

Parameters

ipv6-address

Specifies the IPv6 address of the DNS server(s), up to a maximum of four, specified as eight 16-bit hexadecimal pieces.

Platforms

7705 SAR Gen 2

server

Syntax

server *ip-address* [*ip-address*] **router** *router-instance*

server *ip-address* [*ip-address*] **service-name** *service-name*

no server

Context

[Tree] (config>service>vprn>if>sap>ipsec-gw>dhcp server)

[Tree] (config>service>ies>if>sap>ipsec-gw>dhcp server)

Full Context

configure service vprn interface sap ipsec-gw dhcp server

configure service ies interface sap ipsec-gw dhcp server

Description

This command specifies up to eight DHCPv4 server addresses for DHCPv4-based address assignment. If multiple server addresses are specified, the first advertised DHCPv4 address received is chosen.

Default

no server

Parameters

ip-address

Specifies up to eight unicast IPv4 addresses.

Values ipv4-address a.b.c.d

router-instance

Specifies the router instance ID used to reach the configured server address.

This variant of this command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **server** *ip-address* **service-name** *service-name* variant can be used in all configuration modes.

Values {router-name | *vprn-svc-id*}

vprn-svc-id: 1 to 2147483647

router-name: *router-name* is an alias for input only. The *router-name* gets replaced with an id automatically by SR OS in the configuration).

Default Base

service-name

Specifies the name of the IES or VPRN service used to reach the configured server address, up to 64 characters.

Platforms

7705 SAR Gen 2

server

Syntax

server *ipv6-address* [*ipv6-address*] **router** *router-instance*

server *ipv6-address* [*ipv6-address*] **service-name** *service-name*

no server

Context

[Tree] (config>service>vprn>if>sap>ipsec-gw>dhcp6 server)

[Tree] (config>service>ies>if>sap>ipsec-gw>dhcp6 server)

Full Context

configure service vprn interface sap ipsec-gw dhcp6 server

configure service ies interface sap ipsec-gw dhcp6 server

Description

This command specifies up to eight DHCPv6 server addresses for DHCPv6-based address assignment. If multiple server addresses are specified, the first advertised DHCPv6 address received is chosen.

Default

no server

Parameters

ipv6-address

Specifies up to eight unicast global unicast IPv6 addresses.

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0..FFFF]H
	d - [0..255]D

router-instance

Specifies the router instance ID used to reach the configured server address.

This variant of this command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **server** *ip-address* **service-name** *service-name* variant can be used in all configuration modes.

Values {router-name | *vprn-svc-id*}

vprn-svc-id:	1 to 2147483647
router-name:	*router-name* is an alias for input only. The *router-name* gets replaced with an id automatically by SR OS in the configuration).

Default Base

service-name

Specifies the name of the IES or VPRN service used to reach the configured server address, up to 64 characters.

Platforms

7705 SAR Gen 2

server

Syntax

server

Context

[Tree] (config>test-oam>twamp server)

Full Context

configure test-oam twamp server

Description

This command configures the node for TWAMP server functionality.

Platforms

7705 SAR Gen 2

server

Syntax

server *server* [*server*]

Context

[Tree] (config>router>if>dhcp server)

Full Context

configure router interface dhcp server

Description

This command specifies a list of servers where requests will be forwarded. The list of servers can be entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list. There can be a maximum of eight DHCP servers configured.

The flood command is applicable only in the VPLS case. There is a scenario with VPLS where the VPLS node only wants to add Option 82 information to the DHCP request to provider per-subscriber information, but it does not do full DHCP relay. In this case, the server is set to "flood". This means the DHCP request is still a broadcast and is sent through the VPLS domain. A node running at Layer 3 further upstream then can perform the full Layer 3 DHCP relay function.

Default

no server

Parameters

server

Specifies the DHCP server IP address. A maximum of eight servers can be specified in a single statement.

Platforms

7705 SAR Gen 2

server

Syntax

server *ipv6-address* [*ipv6-address*]

no server

Context

[Tree] (config>router>router-advert>if>dns-options server)

[Tree] (config>router>router-advert>dns-options server)

Full Context

configure router router-advertisement interface dns-options server

configure router router-advertisement dns-options server

Description

This command specifies the IPv6 DNS servers to include in the RDNSS option in Router Advertisements. When specified at the router advertisement level this applies to all interfaces that have **include-dns** enabled, unless the interfaces have more specific **dns-options** configured.

Parameters

ipv6-address

Specifies the IPv6 address of the DNS servers as eight 16-bit hexadecimal pieces. A maximum of four ipv6 addresses can be specified in a single statement.

Platforms

7705 SAR Gen 2

server

Syntax

server [**router** *router-instance* | **service-name** *service-name*] {*ip-address* | *ipv6-address* | **ptp**} [**key-id** *key-id* | **authentication-keychain** *keychain-name*] [**version** *version*] [**prefer**]

no server [**router** *router-instance* | **service-name** *service-name*] {*ip address* | *ipv6-address* | **ptp**}

Context

[Tree] (config>system>time>ntp server)

Full Context

configure system time ntp server

Description

This command configures the node to operate in client mode with the NTP server specified in the address field of this command.

If the internal PTP process is used as a source of time for System Time and OAM time then it must be specified as a server for NTP. If PTP is specified, the **prefer** parameter must be specified. After PTP has established a UTC traceable time from an external grandmaster it is always the source for time into NTP, even if PTP goes into time holdover.

Using the internal PTP time source for NTP promotes the internal NTP server to stratum 1 level, which may impact the NTP network topology.

The **no** form of this command removes the server with the specified address from the configuration.

Parameters

router-instance

Specifies the routing context that contains the interface in the form of *router-name* or *service-id*.

Values *router-name* — Base | Management

service-id — 1 to 2147483647

Default Base

service name

Specifies the service name for the VPRN, up to 64 characters. CPM routing instances are not supported.

ip-address

Configures the IPv4 address of an external NTP server.

Values a.b.c.d

ipv6-address

Configures the IPv6 address of an external NTP server.

Values
- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF] H
- d: [0 to 255] D

key-id

Specifies the key ID that identifies the configured authentication key and authentication type used by this node to transmit NTP packets to an NTP server. If an NTP packet is received by this node, the authentication key-id, type, and key value must be valid, otherwise the packet is rejected and an event/trap generated. This is an optional parameter.

Values 1 to 255

keychain-name

Identifies the keychain name, up to 32 characters.

version

Configures the NTP version number that is expected by this node. This is an optional parameter.

Values 2 to 4

Default 4

ptp

Configures the internal PTP process as a time server into the NTP process. The **prefer** parameter is mandatory with this server option.

prefer

Specifies that, when configuring more than one peer, one remote system can be configured as the preferred peer. When a second peer is configured as preferred, the new entry overrides the old entry.

Platforms

7705 SAR Gen 2

server

Syntax

server

Context

[Tree] (config>system>security>ssh>key-re-exchange server)

Full Context

configure system security ssh key-re-exchange server

Description

This command enables the key re-exchange context for the SSH server.

Platforms

7705 SAR Gen 2

server

Syntax

server *index* **address** *ip-address* **secret** *key* [**hash** | **hash2** | **custom**] [**tls-client-profile** *profile*] [**authenticator** {**md5** | **sm3**}]

no server *index*

Context

[Tree] (config>service>vprn>aaa>rmt-srv>radius server)

[Tree] (config>system>security>radius server)

Full Context

configure service vprn aaa remote-servers radius server

configure system security radius server

Description

This command adds a RADIUS server and configures the IP address, index, and key values.

Up to five RADIUS servers can be configured at any one time. For authentication requests, RADIUS servers are accessed in order from the lowest to highest index until a response from a server is received. A higher indexed server is only queried if no response is received from a lower indexed server (which implies that the server is not available). If a response from a server is received, no other RADIUS servers are queried. It is assumed that there are multiple identical servers configured as backups and that the servers do not have redundant data.

The **no** form of this command removes the server from the configuration.

Default

no server

Parameters

index

Specifies the index for the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index.

Values 1 to 5

ip-address

Specifies the IP address of the RADIUS server. Two RADIUS servers cannot have the same IP address. An error message is generated if the server address is a duplicate.

Values		
	ipv4-address	a.b.c.d (host bits must be 0)
	ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x: [0..FFFF]H
		d: [0..255]D

key

Specifies the secret key to access the RADIUS server, up to 64 characters. This secret key must match the password on the RADIUS server.

hash

Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, cleartext form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

custom

Specifies the custom encryption to management interface.

tls-client-profile

Specifies the TLS profile for the RADIUS server.

profile

Specifies the TLS profile name, up to 32 characters.

md5

Specifies the MD5 hash algorithm for the RADIUS server.

sm3

Specifies the SM3 hash algorithm for the RADIUS server.

Platforms

7705 SAR Gen 2

server

Syntax

server *server-index* [**create**]

no server *server-index*

Context

[Tree] (config>system>security>ldap server)

Full Context

configure system security ldap server

Description

This command configures an LDAP server. Up to five servers can be configured, which can then work in a redundant manner.

The **no** version of this command removes the server connection.

Parameters

server-index

Specifies a unique LDAP server connection.

Values 1 to 5

Platforms

7705 SAR Gen 2

server

Syntax

server [*ip-address* | **fqdn**] [**port** *port*]

no server

Context

[Tree] (config>system>security>pki>est-profile server)

Full Context

configure system security pki est-profile server

Description

Commands in this context configure EST server parameters.

The **no** form of the command reverts to the default value.

Parameters

ip-address

Specifies the IP address of the server.

Values		
	ipv4-address	a.b.c.d (host bits must be 0)
	ipv6-address	x:x:x:x:x:x:x:x
		x:x:x:x:x:x:d.d.d.d
		x: [0..FFFF]H
		d: [0..255]D

fqdn

Specifies to use the Fully Qualified Domain Name (FQDN) of the EST server, up to 255 characters.

port

Specifies the port number of the EST server.

Values 1 to 65535

Default 443

Platforms

7705 SAR Gen 2

server

Syntax

server

Context

[Tree] (config>system>security>ssh>authentication-method server)

Full Context

configure system security ssh authentication-method server

Description

Commands in this context configure, at the system level, the authentication method that the SSH server accepts for the session.

Platforms

7705 SAR Gen 2

server

Syntax

server

Context

[Tree] (config>system>security>user>ssh-auth-method server)

Full Context

configure system security user ssh-authentication-method server

Description

Commands in this context configure, at the user level, the authentication method accepted by the SSH server for the session. The user-level configuration overrides the system-level configuration.

Platforms

7705 SAR Gen 2

27.77 server-address

server-address

Syntax

server-address *ip-address* [**version** *version-number*] [**normal** | **preferred**]

[**interval** *seconds*]

no server-address *ip-address*

Context

[Tree] (config>system>time>sntp server-address)

Full Context

configure system time sntp server-address

Description

This command creates an SNTP server for unicast client mode.

Parameters

ip-address

Specifies the IP address of the SNTP server.

Values a.b.c.d

version-number

Specifies the SNTP version supported by this server.

Values 1 to 3

Default 3

normal | preferred

Specifies the preference value for this SNTP server. When more than one time-server is configured, one server can have preference over others. The value for that server should be set to **preferred**. Only one server in the table can be a preferred server.

Default normal

seconds

Specifies the frequency at which this server is queried.

Values 64 to 1024

Default 64

Platforms

7705 SAR Gen 2

27.78 server-cipher-list

server-cipher-list

Syntax

server-cipher-list

Context

[Tree] (config>system>security>ssh server-cipher-list)

Full Context

configure system security ssh server-cipher-list

Description

Commands in this context configure a list of allowed ciphers by the SSH server.

Platforms

7705 SAR Gen 2

server-cipher-list

Syntax

server-cipher-list *name* [**create**]

no server-cipher-list *name*

Context

[Tree] (config>system>security>tls server-cipher-list)

Full Context

configure system security tls server-cipher-list

Description

This command creates the cipher list that is compared against cipher lists sent by the client to the server in the client hello message. The list contains all ciphers that are supported and desired by SR OS for use in the TLS session. The first common cipher found in both the server and client cipher lists will be chosen. As such, the most desired ciphers should be added at the top of the list.

The **no** form of the command removes the cipher list.

Parameters

name

Specifies the name of the server cipher list, up to 32 characters in length.

create

Keyword used to create the server cipher list.

Platforms

7705 SAR Gen 2

27.79 server-group-list

server-group-list

Syntax

server-group-list *name* [**create**]

no server-group-list *name*

Context

[Tree] (config>system>security>tls server-group-list)

Full Context

configure system security tls server-group-list

Description

This command configures a list of TLS 1.3-supported group suite codes that the server sends in a server Hello message.

The **no** form of this command removes the server group list.

Parameters

name

Specifies the name of the server group list, up to 32 characters.

create

Keyword used to create the server group list.

Platforms

7705 SAR Gen 2

27.80 server-host-key-list

server-host-key-list

Syntax

server-host-key-list

Context

[Tree] (config>system>security>ssh server-host-key-list)

Full Context

configure system security ssh server-host-key-list

Description

Commands in this context configure the list of host key algorithms negotiated by the SR OS acting as the SSH server.

Platforms

7705 SAR Gen 2

27.81 server-id

server-id

Syntax

server-id duid-en hex *hex-string*

server-id duid-en string *ascii-string*

server-id duid-ll

no server-id

Context

[Tree] (config>service>vprn>dhcp6>server server-id)

[Tree] (config>router>dhcp6>server server-id)

Full Context

configure service vprn dhcp6 local-dhcp-server server-id

configure router dhcp6 local-dhcp-server server-id

Description

This command allows the operator to customize the **server-id** attribute of a DHCPv6 message (such as DHCPv6 advertise and reply). By default, the **server-id** uses DUID-ll derived from the chassis link layer address. Operators have the option to use a unique identifier by using the **duid-en** (vendor based on an enterprise number). There is a maximum length associated with the customizable hex-string and ascii-string.

The **no** form of this command reverts to the default.

Default

server-id duid-ll

Parameters

hex-string

Specifies a DUID system ID in a hex format.

Values 0x0 to 0xFFFFFFFF (maximum 116 hex nibbles)

ascii-string

Specifies a DUID system ID in an ASCII format, up to 58 characters.

duid-ll

Specifies that the DUID system ID is derived from the system link layer address.

duid-en

Specifies the enterprise number.

Platforms

7705 SAR Gen 2

27.82 server-kex-list

server-kex-list

Syntax

server-kex-list

Context

[Tree] (config>system>security>ssh server-kex-list)

Full Context

configure system security ssh server-kex-list

Description

This command configures SSH KEX algorithms for SR OS as an SSH server.

An empty list is the default list that the SSH KEX advertises. The default list contains the following:

ecdh-sha2-nistp512

ecdh-sha2-nistp384

ecdh-sha2-nistp256

diffie-hellman-group16-sha512

diffie-hellman-group14-sha256

diffie-hellman-group14-sha1

diffie-hellman-group-exchange-sha1

diffie-hellman-group1-sha1

Platforms

7705 SAR Gen 2

27.83 server-mac-list

server-mac-list

Syntax

server-mac-list

Context

[Tree] (config>system>security>ssh server-mac-list)

Full Context

configure system security ssh server-mac-list

Description

This command allows the user to configure SSH MAC algorithms for SR OS as an SSH server.

Platforms

7705 SAR Gen 2

27.84 server-shutdown

server-shutdown

Syntax

[**no**] **server-shutdown**

Context

[Tree] (config>system>security>ssh server-shutdown)

Full Context

configure system security ssh server-shutdown

Description

This command enables the SSH servers running on the system.

Default

no server-shutdown

Platforms

7705 SAR Gen 2

27.85 server-signature-list

server-signature-list

Syntax

server-signature-list *name* [**create**]

no server-signature-list *name*

Context

[Tree] (config>system>security>tls server-signature-list)

Full Context

configure system security tls server-signature-list

Description

This command configures a list of TLS 1.3-supported signature suite codes for the digital signature that the server sends in a server Hello message.

The **no** form of this command removes the server signature list.

Parameters

name

Specifies the name of the server signature list, up to 32 characters.

create

Keyword used to create the server signature list.

Platforms

7705 SAR Gen 2

27.86 server-timeout

server-timeout

Syntax

server-timeout *seconds*

no server-timeout

Context

[Tree] (config>port>ethernet>dot1x server-timeout)

Full Context

configure port ethernet dot1x server-timeout

Description

This command configures the period during which the router waits for the RADIUS server to respond to its access request message. When this timer expires, the router will re-send the access request message, up to the specified number times.

The **no** form of this command returns the value to the default.

Default

server-timeout 30

Parameters

seconds

Specifies the server timeout period, in seconds.

Values 1 to 300

Platforms

7705 SAR Gen 2

27.87 server-tls-profile

server-tls-profile

Syntax

server-tls-profile *name* [**create**]

no server-tls-profile *name*

Context

[Tree] (config>system>security>tls server-tls-profile)

Full Context

configure system security tls server-tls-profile

Description

This command creates a TLS server profile. This profile can be used by applications that support TLS for encryption. The applications should not send any PDUs until the TLS handshake has been successful.

The **no** form of the command removes the TLS server profile.

Parameters

name

Specifies the name of the TLS server profile, up to 32 characters in length.

create

Keyword used to create the TLS server profile.

Platforms

7705 SAR Gen 2

27.88 servers

servers

Syntax

servers

Context

[Tree] (config>aaa>radius-srv-plcy servers)

Full Context

configure aaa radius-server-policy servers

Description

Commands in this context configure radius-server-policy parameters.

Platforms

7705 SAR Gen 2

27.89 service

service

Syntax

service *service-id*

no service

Context

[Tree] (config>service>vpls>sap>msap-defaults service)

Full Context

configure service vpls sap msap-defaults service

Description

This command sets default service for all subscribers created based on trigger packets received on the given capture SAP in case the corresponding VSA is not included in the RADIUS authentication response. This command is applicable to capture SAP only.

The **no** form of this command reverts to the default.

Parameters

service-id

Specifies the service ID as an integer or a name.

Values *service-id* - 1 to 2147483648

service-name - up to 64 characters

Platforms

7705 SAR Gen 2

service

Syntax

service *service-id* **preference** *preference*

no service *service-id*

Context

[Tree] (config>router>dns>redirect-vprn service)

Full Context

configure router dns redirect-vprn service

Description

This command configures the VPRN DNS redirection for the specified service.

The **no** form of this command removes the service from the VPRN DNS resolution configuration.

Parameters

service-id

Specifies the unique service identification number or string identifying the service in the service domain.

Values *service-id*: 1 to 2147483647

svc-name: 64 characters maximum

preference

Specifies the service preference.

Values 0 to 255

Platforms

7705 SAR Gen 2

service

Syntax

[**no**] **service** *service-id*

Context

[Tree] (config>log>services-all-events service)

Full Context

configure log services-all-events service

Description

This command enables access to the entire system-wide set of log events (VPRN and non-VPRN) in the logs configured within the management VPRN specified by the service ID.

The **no** form of the command enables the display of VPRN events only.

Parameters

service-id

Identifies the VPRN.

Values {*id* | *svc-name*}

id: 1 to 2147483647

svc-name: up to 64 characters

Platforms

7705 SAR Gen 2

service

Syntax

service *service-id*

service name *service-name*

no service

Context

[Tree] (config>system>security>pki>ca-profile>ocsp service)

Full Context

configure system security pki ca-profile ocsp service

Description

This command specifies the service or routing instance that used to contact OCSP responder. This applies to OCSP responders that either configured in CLI or defined in AIA extension of the certificate to be verified.

The responder-url will also be resolved by using the DNS server configured in the configured routing instance.

With VPRN services, the system checks whether the specified service ID or service name is an existing VPRN service at the time of CLI configuration. Otherwise the configuration fails.

Parameters

service-id

Specifies an existing service ID to be used in the match criteria.

This variant of this command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **service name** *service-name* variant can be used in all configuration modes.

Values service-id: 1 to 2147483647 base-router: 0

name service-name

Identifies the service, up to 64 characters.

Platforms

7705 SAR Gen 2

27.90 service-id

service-id

Syntax

service-id *service-id*

no service-id

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>host-ident service-id)

Full Context

configure subscriber-mgmt local-user-db ipoe host host-identification service-id

Description

This command specifies the service ID to match for a host lookup. When the LUDB is accessed using a DHCPv4 server, the SAP ID is matched against the Nokia vendor-specific sub-option in DHCP Option 82.

The **no** form of this command removes the service ID from the configuration.

Parameters

service-id

Specifies an existing service ID or service name.

Values *service-id* — 1 to 2147483647

service-name — up to 64 characters

Platforms

7705 SAR Gen 2

service-id

Syntax

[**no**] **service-id**

Context

[Tree] (config>service>vpls>sap>dhcp>option>vendor service-id)

[Tree] (config>service>vprn>if>dhcp>option>vendor service-id)

Full Context

configure service vpls sap dhcp option vendor-specific-option service-id

configure service vprn interface dhcp option vendor-specific-option service-id

Description

This command enables the sending of the service ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the service ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

Platforms

7705 SAR Gen 2

service-id

Syntax

[**no**] **service-id**

Context

[Tree] (config>router>if>dhcp>option>vendor-specific-option service-id)

Full Context

configure router interface dhcp option vendor-specific-option service-id

Description

This command enables the sending of the service ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the service ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

Default

no service-id

Platforms

7705 SAR Gen 2

service-id

Syntax

service-id *service-id*

no service-id

Context

[Tree] (config>redundancy>mc>peer>mcr>l3-ring>ibc service-id)

Full Context

configure redundancy multi-chassis peer mc-ring l3-ring in-band-control-path service-id

Description

This command specifies the service ID if the interface used for the inband control connection belongs to a VPRN service. If not specified, the service-id is zero and the interface must belong to the Base router. This command supersedes the configuration of a service name.

The no form of this command removes the service ID from the IBC configuration.

Parameters

service-id

Specifies a service ID or an existing service name.

Values 1 to 214748364 - Only supported in 'classic' configuration-mode (**configure>system>management-interface>configuration-mode classic**)

Platforms

7705 SAR Gen 2

service-id

Syntax

service-id *service-id*

no service-id

Context

[Tree] (config>redundancy>mc>peer>mcr>l3ring>node>cv service-id)

Full Context

configure redundancy multi-chassis peer mc-ring l3-ring ring-node connectivity-verify service-id

Description

This command specifies the service ID of the SAP used for the ring-node connectivity verification of this ring node. This command supersedes the configuration of a service name.

The **no** form of the command removes the service ID from the CV configuration.

Default

no service-id

Parameters

service-id

Specifies the service ID or an existing service name.

Values 1 to 2147483647- Only supported in "classic" configuration mode (**configure system management-interface configuration-mode classic**)

Platforms

7705 SAR Gen 2

27.91 service-id-lag-hashing

service-id-lag-hashing

Syntax

[**no**] **service-id-lag-hashing**

Context

[Tree] (config>system>load-balancing service-id-lag-hashing)

Full Context

configure system load-balancing service-id-lag-hashing

Description

This command enables enhanced VLL LAG service ID hashing. This command improves the LAG spraying of VLL service packets and is applied only when both ECMP and LAG hashing are performed by the same router. By default, the ECMP interface and LAG link for all packets on the VLL service are selected based on a direct modulo operation of the service ID. This command enhances distribution and hashes the service ID prior to the LAG link modulo operation when an ECMP link modulo operation is performed.

The **no** form of the command preserves the default behavior of VLL LAG service ID hashing.

Default

no service-id-lag-hashing

Platforms

7705 SAR Gen 2

27.92 service-id-range

service-id-range

Syntax

service-id-range start *service-id* **end** *service-id*

no service-id-range

Context

[Tree] (config>service>md-auto-id service-id-range)

Full Context

configure service md-auto-id service-id-range

Description

This command specifies the range of IDs used by SR OS to automatically assign an ID to services that are created in model-driven interfaces without an ID explicitly specified by the user or client.

A service created with an explicitly-specified ID cannot use an ID in this range. In the classic CLI and SNMP, the ID range cannot be changed while objects exist inside the previous or new range. In MD interfaces, the range can be changed, which causes any previously existing objects in the previous ID range to be deleted and re-created using a new ID in the new range.

The **no** form of this command removes the range values.

See the **config>service md-auto-id** command for further details.

Default

no service-id-range

Parameters

start ***service-id***

Specifies the lower value of the ID range. The value must be less than or equal to the **end** value.

Values 1 to 2147483647

end ***service-id***

Specifies the upper value of the ID range. The value must be greater than or equal to the **start** value.

Values 1 to 2147483647

Platforms

7705 SAR Gen 2

27.93 service-mtu

service-mtu

Syntax

service-mtu *octets*

no service-mtu

Context

[Tree] (config>service>vpls service-mtu)

[Tree] (config>service>template>vpls-template service-mtu)

Full Context

configure service vpls service-mtu

configure service template vpls-template service-mtu

Description

This command configures the service payload (Maximum Transmission Unit – MTU), in bytes, for the service. This MTU value overrides the service-type default MTU. The **service-mtu** defines the payload capabilities of the service. It is used by the system to validate the SAP and SDP binding's operational state within the service.

The service MTU and a SAP's service delineation encapsulation overhead (4 bytes for a dot1q tag) is used to derive the required MTU of the physical port or channel on which the SAP was created. If the required payload is larger than the port or channel MTU, then the SAP will be placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP will be able to transition to the operative state.

When binding an SDP to a service, the service MTU is compared to the path MTU associated with the SDP. The path MTU can be administratively defined in the context of the SDP. The default or administrative path MTU can be dynamically reduced due to the MTU capabilities discovered by the tunneling mechanism of the SDP or the egress interface MTU capabilities based on the next hop in the tunnel path. If the service MTU is larger than the path MTU, the SDP binding for the service will be placed in an inoperative state. If the service MTU is equal to or less than the path MTU, then the SDP binding will be placed in an operational state.

If a service MTU, port or channel MTU, or path MTU is dynamically or administratively modified, then all associated SAP and SDP binding operational states are automatically re-evaluated.

For i-VPLS and Epipes bound to a b-VPLS, the service-mtu must be at least 18 bytes smaller than the b-VPLS service MTU to accommodate the PBB header.

The **no** form of this command returns the default **service-mtu** for the indicated service type to the default value.

Default

service-mtu 1514

Parameters

octets

The following table displays MTU values for specific VC types.

Table 86: MTU Values

VC-Type	Example Service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
VPLS	1514	1500

VC-Type	Example Service MTU	Advertised MTU
VPLS (with preserved dot1q)	1518	1504
VLAN (dot1p transparent to MTU value)	1514	1500
VLAN (qinq with preserved bottom qtag)	1518	1504

The size of the MTU in octets, expressed as a decimal integer

Values 1 to 9194

Platforms

7705 SAR Gen 2

service-mtu

Syntax

service-mtu *octets*

no service-mtu

Context

[Tree] (config>service>epipe service-mtu)

Full Context

configure service epipe service-mtu

Description

This command configures the service payload in bytes, for the service. The configured Maximum Transmission Unit (MTU) value overrides the service-type default MTU. The **service-mtu** command defines the payload capabilities of the service. It is used by the system to validate the operational state of the SAP and SDP binding within the service.

The service MTU and a SAP's service delineation encapsulation overhead (4 bytes for a dot1q tag) is used to derive the required MTU of the physical port or channel on which the SAP was created. If the required payload is larger than the port or channel MTU, the SAP is placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP transitions to the operative state.

When binding an SDP to a service, the service MTU is compared to the path MTU associated with the SDP. The path MTU can be administratively defined in the context of the SDP. The default or administrative path MTU can be dynamically reduced due to the MTU capabilities discovered by the tunneling mechanism of the SDP or the egress interface MTU capabilities based on the next hop in the tunnel path. If the service MTU is larger than the path MTU, the SDP binding for the service is placed in an inoperative state. If the service MTU is equal to or less than the path MTU, the SDP binding is placed in an operational state.

If a service MTU, port or channel MTU, or path MTU is dynamically or administratively modified, all associated SAP and SDP binding operational states are automatically reevaluated.

Binding operational states are automatically reevaluated.

For I-VPLS and Epipes bound to a B-VPLS, the service MTU must be at least 18 bytes smaller than the B-VPLS service MTU to accommodate the PBB header.

Because this connects a Layer 2 to a Layer 3 service, adjust the service MTU under the Epipe service. The MTU that is advertised from the Epipe side is service MTU minus EtherHeaderSize.

Note:
In the **configure**>**service**>**epipe** context, the **adv-service-mtu** command can be used to override the configured MTU value used in T-LDP signaling to the far-end of an Epipe spoke-sdp. The **adv-service-mtu** command is also used to validate the value signaled by the far-end PE. For more information, see **adv-service-mtu** command.

The **no** form of this command returns the default **service-mtu** for the indicated service type to the default value.

By default, if **no service-mtu** is configured, the MTU value is (1514 - 14) = 1500.

Default

no service-mtu 1508 (for Apipe, Fpipe)

no service-mtu 1500 (for Ipipe)

no service-mtu 1524 (for Epipe)

Table 87: MTU Values lists the MTU values for specific VC types.

Table 87: MTU Values

SAP VC-Type	Example: Service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
VPLS	1514	1500
VPLS (with preserved dot1q)	1518	1504
VLAN (dot1p transparent to MTU value)	1514	1500
VLAN (qinq with preserved bottom qtag)	1518	1504

Parameters

octets

Specifies the MTU size in octets, expressed as a decimal integer.

Values 1 to 9782

1 to 9800 (for Epipe only)

Platforms

7705 SAR Gen 2

27.94 service-name

service-name

Syntax

service-name *service-name*

no service-name

Context

[Tree] (config>redundancy>mc>peer>mcr>l3ring>node>cv service-name)

Full Context

configure redundancy multi-chassis peer mc-ring l3-ring ring-node connectivity-verify service-name

Description

This command specifies the service name of the SAP used for ring-node connectivity verification of this ring node. This command supersedes the configuration of a service ID.

The **no** form of this command removes the service name from the CV configuration.

Default

no service-name

Parameters

service-name

Specifies a service name, up to 64 characters.

Platforms

7705 SAR Gen 2

27.95 service-range

service-range

Syntax

service-range *startid-endid* [**start-vlan-id** *startvid*]

no service-range

Context

[Tree] (config>service>vpls>vpls-group service-range)

Full Context

configure service vpls vpls-group service-range

Description

This command configures the service ID and implicitly the VLAN ID ranges to be used as input variables for related VPLS and SAP templates to pre-provision "data" VPLS instances and related SAPs using the service ID specified in the command. If the **start-vlan-id** is not specified then the service-range values are used for vlan-ids. The data SAPs will be instantiated on all the ports used to specify SAP instances under the related control VPLS.

Modifications of the service id and vlan ranges are allowed with the following restrictions.

- service-range increase can be achieved in two ways:
  - Allowed when vpls-group is in shutdown state
  - By creating a new vpls-group
- service-range decrease can be achieved in two ways:
  - Allowed when vpls-group is in shutdown state; when **shutdown** command is executed the associated service instances are deleted.
  - Allowed when vpls-group is in no shutdown state and has completed successfully instantiating services.
  - In both cases, only the services that do not have user configured SAPs will be deleted. Otherwise the above commands are rejected. Existing declarations or registrations do not prevent service deletion.
- start-vlan-id change can be achieved in two ways:
  - Allowed when vpls-group is in shutdown state
  - At the time of range decrease by increasing the start-vlan-id which can be done when vpls-group is in no shutdown state and has completed successfully instantiating services

The **no** form of this command removes the specified ranges and deletes the pre-provisioned VPLS instances and related SAPs. The command will fail if any of the VPLS instances in the affected ranges have a provisioned SAP.

Default

no service-range

Parameters

startid-endid

Specifies the range of service IDs

Values 1 to 2147483647

startvid

Specifies the starting VLAN ID; it provides a way to set aside a service ID range that is not the same as the VLAN range and allows for multiple MVRP control-VPLSs to control same VLAN range on different ports.

Values 1 to 4094

Platforms

7705 SAR Gen 2

27.96 service-request

service-request

Syntax

[**no**] **service-request**

Context

[Tree] (config>service>vprn>aaa>remote-servers>tacplus service-request)

[Tree] (config>system>security>tacplus service-request)

Full Context

configure service vprn aaa remote-servers tacplus service-request

configure system security tacplus service-request

Description

This command enables Nokia services to be requested from the TACACS+ server.

The **no** form of this command disables Nokia services from being requested from the TACACS+ server.

Default

no service-request

Platforms

7705 SAR Gen 2

27.97 services-all-events

services-all-events

Syntax

services-all-events

Context

[Tree] (config>log services-all-events)

Full Context

configure log services-all-events

Description

Commands in this context control which log events are present in VPRN logs.

By default, the event streams for VPRN logs contain only events that are associated with the particular VPRN.

Access to the entire system-wide set of events (VPRN and non-VPRN) can be enabled using the **services-all-events** command.

Platforms

7705 SAR Gen 2

27.98 session

session

Syntax

session *session-name* [**test-family** [**ethernet** | **ip** | **mpls**] [**session-type** {**proactive** | **on-demand**}] **create**]

no session *session-name*

Context

[Tree] (config>oam-pm session)

Full Context

configure oam-pm session

Description

This command creates the individual session containers that houses the test specific configuration parameters. Since this session context provides only a container abstract to house the individual test functions, it cannot be shut down. Individual tests sessions within the container may be shut down. No values, parameters, or configuration within this context may be changed if any individual test is active. Changes may only be made when all tests within the context are shut down. The only exception to this is the description value.

The **no** form of this command deletes the session.

Parameters

session-name

Specifies the session name, up to 32 characters.

test-family

Indicates the type family and sets the context for the individual parameters.

Values **ethernet** — Specifies that the test is based on the Ethernet layer.

ip — Specifies that the test is based on the IP layer.

mpls — Specifies that the test is based on the MPLS layer.

session-type

Specifies how to set the Type bit in the Flags byte, and influences how different test criteria may be applied to the individual test. Not all test families carry this information in the PDU.

Values **proactive** — Sets the type to always on, with an immediate start and no stop.

on-demand — Sets the type to on-demand, with an immediate start and no stop, or a stop based on the offset.

Default proactive

create

Creates the PM session.

Platforms

7705 SAR Gen 2

27.99 session-limits

session-limits

Syntax

session-limits

Context

[Tree] (config>service>nat>nat-policy session-limits)

Full Context

configure service nat nat-policy session-limits

Description

Commands in this context configure session limits for the NAT policy.

Platforms

7705 SAR Gen 2

27.100 session-parameters

session-parameters

Syntax

session-parameters

Context

[Tree] (config>router>ldp session-parameters)

Full Context

configure router ldp session-parameters

Description

Commands in this context configure peer specific parameters.

Platforms

7705 SAR Gen 2

27.101 session-sender-type

session-sender-type

Syntax

session-sender-type {**twamp-light** | **stamp**}

Context

[Tree] (config>oam-pm>session>ip>twamp-light session-sender-type)

Full Context

configure oam-pm session ip twamp-light session-sender-type

Description

This command configures the type of test packet format to transmit.

Default

session-sender-type twamp-light

Parameters

twamp-light

Specifies TWAMP-Light transmission, packet formatting, and packet processing. TWAMP-Light test packets do not allow TLVs.

stamp

Specifies STAMP transmission, packet formatting, and packet processing. STAMP test packets support TLVs.

Platforms

7705 SAR Gen 2

27.102 set-time

set-time

Syntax

set-time *date time*

Context

[Tree] (admin set-time)

Full Context

admin set-time

Description

This command sets the local system time.

The time entered should be accurate for the time zone configured for the system. The system will convert the local time to UTC before saving to the system clock which is always set to UTC. This command does not take into account any daylight saving offset if defined.

If SNTP or NTP is enabled (no shutdown) then this command cannot be used.

Parameters

date

Specifies the local date and time accurate to the minute in the YYYY/MM/DD format.

Values *YYYY* is the four-digit year

MM is the two-digit month

DD is the two-digit date

time

Specifies the time (accurate to the second) in the *hh*:*mm[:ss]* format. If no seconds value is entered, the seconds are reset to :00.

Values *hh* is the two-digit hour in 24 hour format (00=midnight, 12=noon)*mm* is the two-digit minute

Default 0

Platforms

7705 SAR Gen 2

27.103 severity

severity

Syntax

severity {**eq** | **neq** | **lt** | **lte** | **gt** | **gte**} *severity-level*

no severity

Context

[Tree] (config>service>vprn>log>filter>entry>match severity)

Full Context

configure service vprn log filter entry match severity

Description

This command adds an event severity level as a match criterion. Only one severity command can be entered per event filter entry. The latest severity command overwrites the previous command.

The **no** form of this command removes the severity match criterion.

Default

no severity

Parameters

eq | neq | lt | lte | gt | gte

Specifies the type of match. Valid operators are listed below.

Values

Table 88: Valid Operators

Operator	Notes
eq	equal to
neq	not equal to
lt	less than
lte	less than or equal to
gt	greater than
gte	greater than or equal to

severity-name

The ITU severity level name. Table 89: Severity Levels lists severity names and corresponding numbers per ITU standards M.3100 X.733 & X.21 severity levels.

Table 89: Severity Levels

Severity Number	Severity Name
1	cleared
2	indeterminate (info)
3	critical
4	major
5	minor
6	warning

Values cleared, intermediate, critical, major, minor, warning

Platforms

7705 SAR Gen 2

severity

Syntax

severity {**eq** | **neq** | **lt** | **lte** | **gt** | **gte**} *severity-level*

no severity

Context

[Tree] (config>log>filter>entry>match severity)

Full Context

configure log filter entry match severity

Description

This command adds an event severity level as a match criterion. Only one severity command can be entered per event filter entry. The latest severity command overwrites the previous command.

The **no** form of this command removes the severity match criterion.

Parameters

eq | neq | lt | lte | gt | gte

Specifies the match type. Valid operators are listed in Table 90: Valid Operators.

Table 90: Valid Operators

Operator	Notes
eq	equal to
neq	not equal to
lt	less than
lte	less than or equal to
gt	greater than
gte	greater than or equal to

severity-name

Specifies the ITU severity level name. Table 91: ITU Severity Information lists severity names and corresponding numbers per ITU standards M.3100 X.733 & X.21 severity levels.

Table 91: ITU Severity Information

Severity Number	Severity Name
1	cleared
2	indeterminate (info)
3	critical
4	major
5	minor
6	warning

Values cleared, intermediate, critical, major, minor, warning

Platforms

7705 SAR Gen 2

27.104 sf-offset

sf-offset

Syntax

sf-offset *offset-value*

no sf-offset

Context

[Tree] (config>service>vprn>isis>if>level sf-offset)

Full Context

configure service vprn isis interface level sf-offset

Description

If the pre-FEC error rate of the associated DWDM port crosses the configured **sf-threshold**, this offset-value is added to the IS-IS interface metric. This parameter is only effective if the interface is associated with a DWDM port and the **sf-threshold** value is configured under that port.

The **no** form of this command reverts the offset value to 0.

Default

no sf-offset

Parameters

offset-value

Specifies the amount the interface metric is increased by if the **sf-threshold** is crossed.

Values 0 to 16777215

Platforms

7705 SAR Gen 2

sf-offset

Syntax

sf-offset *offset-value*

no sf-offset

Context

[Tree] (config>router>isis>if>level sf-offset)

Full Context

configure router isis interface level sf-offset

Description

If the pre-FEC error rate of the associated DWDM port crosses the configured **sf-threshold**, this offset-value is added to the IS-IS interface metric. This parameter is only effective if the interface is associated with a DWDM port and the **sf-threshold** value is configured under that port.

The **no** form of this command reverts the offset value to 0.

Default

no sf-offset

Parameters

offset-value

Specifies the amount the interface metric is increased by if the **sf-threshold** is crossed.

Values 0 to 16777215

Platforms

7705 SAR Gen 2

27.105 sf-threshold

sf-threshold

Syntax

sf-threshold *threshold* [**multiplier** *multiplier*]

no sf-threshold

Context

[Tree] (config>port>ethernet>crc-monitor sf-threshold)

Full Context

configure port ethernet crc-monitor sf-threshold

Description

This command specifies the error rate at which to declare the Signal Fail condition on an Ethernet interface. The value represents M*10E-N errored frames over total frames received over W seconds of the sliding window. The CRC errors on the interface are sampled once per second. A default of 10 seconds is used when there is no additional window-size configured. The multiplier keyword is optional. If the multiplier keyword is omitted or **no sf-threshold** is specified the multiplier will return to the default value of 1.

Default

no sf-threshold

Parameters

threshold

Specifies the threshold value.

Values 1 to 9

multiplier

Specifies the multiplier value.

Values 1 to 9

Platforms

7705 SAR Gen 2

27.106 sgt-qos

sgt-qos

Syntax

sgt-qos

Context

[Tree] (config>router sgt-qos)

[Tree] (config>service>vprn sgt-qos)

Full Context

configure router sgt-qos

configure service vprn sgt-qos

Description

Commands in this context configure DSCP/dot1p remarking for self-generated traffic.

Platforms

7705 SAR Gen 2

27.107 sham-link

sham-link

Syntax

sham-link *ip-int-name ip-address*

Context

[Tree] (config>service>vprn>ospf>area sham-link)

Full Context

configure service vprn ospf area sham-link

Description

This command is similar to a virtual link with the exception that metric must be included in order to distinguish the cost between the MPLS-VPRN link and the backdoor.

Parameters

ip-int-name

The local interface name used for the sham-link. This is a mandatory parameter and interface names must be unique within the group of defined IP interfaces for **config>router>if**, **config>service>ies>if** and **config>service>vprn>if** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters, the entire string must be enclosed between double quotes. If the IP interface name does not exist or does not have an IP address configured, an error message will be returned.

ip-address

The IP address of the sham-link neighbor in IP address dotted decimal notation. This parameter is the remote peer of the sham link's IP address used to set up the sham-link. This is a mandatory parameter and must be a valid IP address.

Platforms

7705 SAR Gen 2

27.108 sham-neighbor

sham-neighbor

Syntax

sham-neighbor [*ip-address*]

no sham-neighbor

Context

[Tree] (debug>router>ospf sham-neighbor)

Full Context

debug router ospf sham-neighbor

Description

This command enables debugging of the OSPFv2 sham-link neighbor.

Parameters

ip-address

Debugs the sham-link neighbor identified by this IP address.

Platforms

7705 SAR Gen 2

27.109 shell

shell

Syntax

shell -password *password*

no shell

Context

[Tree] (environment shell)

Full Context

environment shell

Description

This command allows Nokia technical support to access the **shell** commands. **shell** commands are used only by Nokia technical support for troubleshooting.

The **no** form of this command disables the **shell** commands.

Parameters

password

Specifies the password to access the **shell** commands, up to 256 characters.

Platforms

7705 SAR Gen 2

27.110 shortcut-local-ttl-propagate

shortcut-local-ttl-propagate

Syntax

[**no**] **shortcut-local-ttl-propagate**

Context

[Tree] (config>router>ldp shortcut-local-ttl-propagate)

[Tree] (config>router>mpls shortcut-local-ttl-propagate)

Full Context

configure router ldp shortcut-local-ttl-propagate

configure router mpls shortcut-local-ttl-propagate

Description

This command configures the TTL handling of locally generated packets for all LSP shortcuts originating on this ingress LER. It applies to all LDP or RSVP LSPs that are used to resolve static routes, BGP routes, and IGP routes.

The user can enable or disable the propagation of the TTL from the header of an IP packet into the header of the resulting MPLS packet independently for local and transit packets forwarded over an LSP shortcut.

Local IP packets include ICMP Ping, traceroute, and OAM packets, that are destined to a route that is resolved to the LSP shortcut. Transit IP packets are all IP packets received on any IES interface and destined to a route that is resolved to the LSP shortcut

By default, the feature propagates the TTL from the header of locally generated IP packets into the label stack of the resulting MPLS packets forwarded over the LSP shortcut. This is referred to as Uniform mode.

When the **no** form of this command is enabled, TTL propagation is disabled on all locally generated IP packets, including ICMP Ping, traceroute, and OAM packets, that are destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack. This is referred to as Pipe mode.

Default

shortcut-local-ttl-propagate

Platforms

7705 SAR Gen 2

27.111 shortcut-transit-ttl-propagate

shortcut-transit-ttl-propagate

Syntax

[**no**] **shortcut-transit-ttl-propagate**

Context

[Tree] (config>router>mpls shortcut-transit-ttl-propagate)

[Tree] (config>router>ldp shortcut-transit-ttl-propagate)

Full Context

configure router mpls shortcut-transit-ttl-propagate

configure router ldp shortcut-transit-ttl-propagate

Description

This command configures the TTL handling of transit packets for all LSP shortcuts originating on this ingress LER. It applies to all LDP or RSVP LSPs that are used to resolve static routes, BGP routes, and IGP routes.

The user can enable or disable the propagation of the TTL from the header of an IP packet into the header of the resulting MPLS packet independently for local and transit packets forwarded over an LSP shortcut.

By default, the feature propagates the TTL from the header of transit IP packets into the label stack of the resulting MPLS packets forwarded over the LSP shortcut. This is referred to as Uniform mode.

When the **no** form of the command is enabled, TTL propagation is disabled on all transit IP packets received on any IES interface and destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack. This is referred to as Pipe mode.

Default

shortcut-transit-ttl-propagate

Platforms

7705 SAR Gen 2

27.112 shortcut-tunnel

shortcut-tunnel

Syntax

shortcut-tunnel

Context

[Tree] (config>router>bgp>next-hop-resolution shortcut-tunnel)

Full Context

configure router bgp next-hop-resolution shortcut-tunnel

Description

This command creates the context to configure the tunnel types that can be used to resolve unlabeled IPv4 and IPv6 BGP routes.

The following tunnel types are supported for resolving IPv4 routes and IPv6 routes with IPv4-mapped IPv6 next-hop addresses: bgp, ldp, rsvp, sr-isis, sr-ospf, sr-policy and sr-te. In this context:

- **bgp** — refers to IPv4 tunnels created by receiving BGP label-unicast IPv4 routes for /32 IPv4 prefixes.
- **ldp** — refers to /32 and shorter length LDP FEC prefixes imported into the tunnel table. For IPv4 NLRI, BGP selects the LDP FEC that is the longest-prefix-match (LPM) of the BGP next-hop address. For IPv6 NLRI, BGP selects the /32 FEC that is an exact match of the BGP next-hop address.

- **rsvp** — refers to RSVP tunnels in the tunnel table to IPv4 destinations. This option allows BGP to use the best metric RSVP LSP to the address of the BGP next-hop. This address can correspond to the system interface or to another loopback interface of the remote BGP router. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel id.
- **sr-isis** — refers to segment routing tunnels (shortest path) to IPv4 destinations reachable by the IS-IS protocol. This option allows BGP to use the segment routing tunnel in the tunnel table submitted by the lowest preference IS-IS instance or (in case of a tie) the lowest numbered IS-IS instance.
- **sr-ospf** — refers to segment routing tunnels (shortest path) to IPv4 destinations reachable by the OSPF protocol. This option allows BGP to use the segment routing tunnel in the tunnel table submitted by the lowest preference OSPF instance or (in case of a tie) the lowest numbered OSPF instance.
- **sr-policy** — refers to segment routing policies with an IPv4 endpoint that are statically configured in the local router or learned through BGP routes (AFI 1/SAFI 73). For BGP to resolve the next hop of an unlabeled IPv4 or IPv6 route using a segment routing policy the highest numbered color extended community attached to the IPv4 or IPv6 route must match the color of the segment routing policy.
- **sr-te** — refers to traffic engineered (TE) segment routing tunnels. This option allows BGP to use the best metric SR-TE tunnel to the address of the BGP next-hop. In the case of multiple SR-TE tunnels with the same lowest metric, BGP selects the tunnel with the lowest tunnel id.
- **udp** — refers to MPLSoUDPoIPv4 tunnels set up by action of the BGP import policies.

The following tunnel types are supported for resolving IPv6 routes with IPv6 next-hops that are not IPv4-mapped IPv6 addresses: ldp, sr-isis, and sr-policy. In this context:

- **ldp** — refers to /128 LDP FEC prefixes in the tunnel table. BGP selects the /128 FEC that is an exact match of the BGP next-hop address.
- **sr-isis** — refers to segment routing tunnels (shortest path) to IPv6 destinations reachable by the IS-IS protocol. This option allows BGP to use the segment routing tunnel in the tunnel table submitted by the lowest preference IS-IS instance or (in case of a tie) the lowest numbered IS-IS instance.
- **sr-policy** — refers to segment routing policies with a null IPv4 endpoint (0.0.0.0) that are statically configured in the local router or learned through BGP routes (AFI 1/SAFI 73). For BGP to resolve the next hop of an IPv6 route using a segment routing policy the highest numbered color extended community attached to the IPv6 route must match the color of the segment routing policy and its color bits must be set to '01' or '10'.

Platforms

7705 SAR Gen 2

27.113 show-ipsec-keys

show-ipsec-keys

Syntax

[**no**] **show-ipsec-keys**

Context

[Tree] (config>ipsec show-ipsec-keys)

Full Context

configure ipsec show-ipsec-keys

Description

This command enables user to optionally include IKE-SA or CHILD-SA keys in the output of **debug ipsec** or **admin ipsec display-key**.

The **no** form of this command disallows the user from including keys in the output.

Default

no show-ipsec-keys

Platforms

7705 SAR Gen 2

27.114 show-request

show-request

Syntax

show-request [**ca** *ca-profile-name*]

Context

[Tree] (admin>certificate>cmpv2 show-request)

Full Context

admin certificate cmpv2 show-request

Description

This command displays current the CMPv2 pending request toward the specified CA. If there is no pending request, the last pending request is displayed including the status (success/fail/rejected) and the receive time of last CMPv2 message from server.

The following information is included in the output:

- Request type, original input parameter (password is not displayed), checkAfter and reason in of last PollRepContent, time of original command input.

Parameters

ca-profile-name

Specifies a ca-profile name, up to 32 characters. If not specified, the system will display pending requests of all ca-profiles.

Platforms

7705 SAR Gen 2

27.115 shutdown

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>script-control>script-policy shutdown)

[Tree] (config>router>mpls>fwd-policies>fwd-policy>egress-statistics shutdown)

[Tree] (config>system>grpc-tunnel>tunnel>handler shutdown)

[Tree] (config>router>fad>flex-algo shutdown)

[Tree] (config>router>mpls>static-lsp shutdown)

[Tree] (config>system>telemetry>notification-bundling shutdown)

[Tree] (config>system>grpc-tunnel>tunnel shutdown)

[Tree] (config>system>telemetry>persistent-subscriptions>subscription shutdown)

[Tree] (config>system>time>ntp shutdown)

[Tree] (config>system>time>sntp shutdown)

[Tree] (config>system>lldp shutdown)

[Tree] (config>system>telemetry>destination-group>tcp-keepalive shutdown)

[Tree] (config>service>vpls>sap>dhcp6>ldra shutdown)

[Tree] (config>system>grpc-tunnel>destination-group>tcp-keepalive shutdown)

[Tree] (config>router>mpls>fwd-policies>fwd-policy>ingress-statistics shutdown)

[Tree] (config>system>script-control>script shutdown)

[Tree] (config>system>cron>sched shutdown)

Full Context

configure system script-control script-policy shutdown

configure router mpls forwarding-policies forwarding-policy egress-statistics shutdown

configure system grpc-tunnel tunnel handler shutdown

configure router flexible-algorithm-definitions flex-algo shutdown

configure router mpls static-lsp shutdown

configure system telemetry notification-bundling shutdown

configure system grpc-tunnel tunnel shutdown

configure system telemetry persistent-subscriptions subscription shutdown

configure system time ntp shutdown

configure system time sntp shutdown

configure system lldp shutdown

configure system telemetry destination-group tcp-keepalive shutdown

configure service vpls sap dhcp6 ldra shutdown

configure system grpc-tunnel destination-group tcp-keepalive shutdown

configure router mpls forwarding-policies forwarding-policy ingress-statistics shutdown

configure system script-control script shutdown

configure system cron schedule shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command places the entity into an administratively enabled state.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>origin-validation>rpki-session shutdown)

[Tree] (config>router>pim>rp>ipv6>bsr-candidate shutdown)

[Tree] (config>router>igmp>if shutdown)

[Tree] (config>router>static-route-entry>indirect shutdown)

[Tree] (config>router>pim>rp>ipv6>embedded-rp shutdown)

[Tree] (config>router>if shutdown)

[Tree] (config>router>static-route-entry>next-hop shutdown)

[Tree] (config>router>igmp shutdown)

[Tree] (config>system>management-interface>cli>md-cli>environment>progress-indicator shutdown)

[Tree] (config>router>mld>if shutdown)

[Tree] (config>router>static-route-entry>black-hole shutdown)

[Tree] (config>router>pim>rp>rp-candidate shutdown)

[Tree] (config>router>pim shutdown)

[Tree] (config>router>pim>rp>ipv6>rp-candidate shutdown)

[Tree] (config>router>mld shutdown)

[Tree] (config>router>pim>interface shutdown)

[Tree] (config>router>pim>rp>bsr-candidate shutdown)

Full Context

configure router origin-validation rpki-session shutdown

configure router pim rp ipv6 bsr-candidate shutdown

configure router igmp interface shutdown

configure router static-route-entry indirect shutdown

configure router pim rp ipv6 embedded-rp shutdown

configure router interface shutdown

configure router static-route-entry next-hop shutdown

configure router igmp shutdown

configure system management-interface cli md-cli environment progress-indicator shutdown

configure router mld interface shutdown

configure router static-route-entry black-hole shutdown

configure router pim rp rp-candidate shutdown

configure router pim shutdown

configure router pim rp ipv6 rp-candidate shutdown

configure router mld shutdown

configure router pim interface shutdown

configure router pim rp bsr-candidate shutdown

Description

This command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system-generated configuration files.

The **no** form of this command places the entity into an administratively enabled state.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>service>ies>if shutdown)

[Tree] (config>service>ies>if>spoke-sdp>control-channel-status shutdown)

[Tree] (config>service>ies shutdown)

[Tree] (config>service>ies>if>dhcp shutdown)

[Tree] (config>service>ies>if>vrrp shutdown)

[Tree] (config>service>ies>if>spoke-sdp shutdown)

[Tree] (config>service>ies>if>dhcp>proxy-server shutdown)

Full Context

configure service ies interface shutdown

configure service ies interface spoke-sdp control-channel-status shutdown

configure service ies shutdown

configure service ies interface dhcp shutdown

configure service ies interface vrrp shutdown

configure service ies interface spoke-sdp shutdown

configure service ies interface dhcp proxy-server shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-lag shutdown)

[Tree] (config>redundancy>multi-chassis>peer>sync shutdown)

[Tree] (config>redundancy>multi-chassis>peer shutdown)

Full Context

configure redundancy multi-chassis peer mc-lag shutdown

configure redundancy multi-chassis peer sync shutdown

configure redundancy multi-chassis peer shutdown

Description

The **shutdown** command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they can be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

Shutting down a subscriber interface will operationally shut down all child group interfaces and SAPs. Shutting down a group interface will operationally shut down all SAPs that are part of that group-interface.

The **no** form of this command puts an entity into the administratively enabled state.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>service>vprn>rip>group shutdown)

[Tree] (config>service>vprn>if shutdown)

[Tree] (config>service>vprn>if>sap>ipsec-tunnel shutdown)
[Tree] (config>service>vprn>ospf shutdown)
[Tree] (config>service>vprn>isis>if shutdown)
[Tree] (config>service>vprn>ospf>area>if shutdown)
[Tree] (config>system>security>radius shutdown)
[Tree] (config>service>vprn>log>log-id shutdown)
[Tree] (config>service>vprn>isis shutdown)
[Tree] (config>service>vprn>ospf>area>sham-link shutdown)
[Tree] (config>service>vprn>bgp>group shutdown)
[Tree] (config>service>vprn>pim shutdown)
[Tree] (config>service>vprn>rip>group>neighbor shutdown)
[Tree] (config>service>vprn>bgp-ipvpn>mpls shutdown)
[Tree] (config>service>vprn>igmp>if shutdown)
[Tree] (config>service>vprn shutdown)
[Tree] (config>service>vprn>ospf3 shutdown)
[Tree] (config>service>vprn>ospf3>area>if shutdown)
[Tree] (config>service>vprn>ntp shutdown)
[Tree] (config>service>vprn>if>ipv6>vrrp shutdown)
[Tree] (config>service>vprn>nw-if shutdown)
[Tree] (config>service>vprn>if>sap shutdown)
[Tree] (config>service>vprn>aaa>rmt-srv>radius shutdown)
[Tree] (config>service>vprn>router-advert>if shutdown)
[Tree] (config>service>vprn>ospf3>area>virtual-link shutdown)
[Tree] (config>service>vprn>igmp shutdown)
[Tree] (config>service>vprn>rip shutdown)
[Tree] (config>service>vprn>pim>if shutdown)
[Tree] (config>service>vprn>pim>rp>bsr-candidate shutdown)
[Tree] (config>service>vprn>if>vrrp shutdown)
[Tree] (config>service>vprn>bgp-evpn>mpls shutdown)
[Tree] (config>service>vprn>pim>rp>ipv6>bsr-candidate shutdown)
[Tree] (config>service>vprn>bgp>group>neighbor shutdown)
[Tree] (config>service>vprn>bgp shutdown)
[Tree] (config>service>vprn>pim>rp>ipv6>rp-candidate shutdown)
[Tree] (config>service>vprn>pim>rp>ipv6>embedded-rp shutdown)
[Tree] (config>service>vprn>ospf>area>virtual-link shutdown)

Full Context

configure service vprn rip group shutdown
configure service vprn interface shutdown
configure service vprn interface sap ipsec-tunnel shutdown
configure service vprn ospf shutdown
configure service vprn isis interface shutdown
configure service vprn ospf area interface shutdown
configure system security radius shutdown
configure service vprn log log-id shutdown
configure service vprn isis shutdown
configure service vprn ospf area sham-link shutdown
configure service vprn bgp group shutdown
configure service vprn pim shutdown
configure service vprn rip group neighbor shutdown
configure service vprn bgp-ipvpn mpls shutdown
configure service vprn igmp interface shutdown
configure service vprn shutdown
configure service vprn ospf3 shutdown
configure service vprn ospf3 area interface shutdown
configure service vprn ntp shutdown
configure service vprn interface ipv6 vrrp shutdown
configure service vprn network-interface shutdown
configure service vprn interface sap shutdown
configure service vprn aaa remote-servers radius shutdown
configure service vprn router-advertisement interface shutdown
configure service vprn ospf3 area virtual-link shutdown
configure service vprn igmp shutdown
configure service vprn rip shutdown
configure service vprn pim interface shutdown
configure service vprn pim rp bsr-candidate shutdown
configure service vprn interface vrrp shutdown
configure service vprn bgp-evpn mpls shutdown
configure service vprn pim rp ipv6 bsr-candidate shutdown
configure service vprn bgp group neighbor shutdown
configure service vprn bgp shutdown
configure service vprn pim rp ipv6 rp-candidate shutdown

configure service vprn pim rp ipv6 embedded-rp shutdown

configure service vprn ospf area virtual-link shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

If the AS number was previously changed, the BGP AS number inherits the new value.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>service>vprn shutdown)

[Tree] (config>service>vpls>mesh-sdp shutdown)

[Tree] (config>service>vpls shutdown)

[Tree] (config>service>vpls>sap shutdown)

[Tree] (config>service>ies>if>sap shutdown)

[Tree] (config>service>vpls>spoke-sdp shutdown)

Full Context

configure service vprn shutdown

configure service vpls mesh-sdp shutdown

configure service vpls shutdown

configure service vpls sap shutdown

configure service ies interface sap shutdown

configure service vpls spoke-sdp shutdown

Description

This command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no**

shutdown command. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system-generated configuration files.

Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>service>vpls>mac-move shutdown)

[Tree] (config>service>vpls>stp shutdown)

[Tree] (config>service>vpls>spoke-sdp shutdown)

[Tree] (config>service>vpls>mld-snooping shutdown)

[Tree] (config>service>vpls>interface shutdown)

[Tree] (config>service>vpls>sap>dhcp>proxy shutdown)

[Tree] (config>service>vpls>spoke-sdp>stp shutdown)

[Tree] (config>service>vpls>bgp-ad shutdown)

[Tree] (config>service>vpls>igmp-snooping shutdown)

[Tree] (config>service>vpls>sap>stp shutdown)

Full Context

configure service vpls mac-move shutdown

configure service vpls stp shutdown

configure service vpls spoke-sdp shutdown

configure service vpls mld-snooping shutdown

configure service vpls interface shutdown

configure service vpls sap dhcp proxy-server shutdown

configure service vpls spoke-sdp stp shutdown

configure service vpls bgp-ad shutdown

configure service vpls igmp-snooping shutdown

configure service vpls sap stp shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>port>ethernet>ssm shutdown)

[Tree] (config>port-xc>pxc shutdown)

[Tree] (config>card>mda shutdown)

[Tree] (config>port>ethernet>dampening shutdown)

[Tree] (config>redundancy>multi-chassis>peer>mc-ep shutdown)

[Tree] (config>redundancy>multi-chassis>ipsec-domain shutdown)

[Tree] (config>port shutdown)

[Tree] (config>card shutdown)

[Tree] (config>redundancy>multi-chassis>peer>mc-ipsec>domain shutdown)

[Tree] (config>lag shutdown)

Full Context

configure port ethernet ssm shutdown

configure port-xc pxc shutdown

configure card mda shutdown

configure port ethernet dampening shutdown

configure redundancy multi-chassis peer mc-endpoint shutdown

configure redundancy multi-chassis ipsec-domain shutdown

configure port shutdown

configure card shutdown

configure redundancy multi-chassis peer mc-ipsec domain shutdown

configure lag shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within.

This command is supported on TDM satellite.

The **no** form of this command administratively enables an entity.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>mpls>lsp-history shutdown)

Full Context

configure router mpls lsp-history shutdown

Description

This command enables the collection of up to the last 100 significant events for each RSVP-TE and SR-TE LSP.

A shutdown of the **lsp-history** pauses the collection of events, but does not remove previously collected events from memory.

The **no** form of this command disables the collection of significant events for LSPs.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>port>ethernet>dot1x>macsec>sub-port shutdown)

Full Context

configure port ethernet dot1x macsec sub-port shutdown

Description

This command shuts down the MACsec under this sub-port specifically, including MKA negotiation. In the shutdown state, this port is not MACsec capable and all PDUs will be transmitted and expected without encryption and authentication.

The **no** form of this command puts the port in MACsec-enabled mode. A valid CA, different than any other CA configured on any other sub-port of this port and also a *max-peer* value larger than 0 must be configured. In MACsec-enabled mode, packets are sent in cleartext until the MKA session is up, and if the **rx-must-be-encrypted** is set on the port, all incoming packets with no MACsec encapsulations are dropped.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>service>sdp shutdown)

[Tree] (config>service>pw-template>stp shutdown)

[Tree] (config>service>sdp>keep-alive shutdown)

Full Context

configure service sdp shutdown

configure service pw-template stp shutdown

configure service sdp keep-alive shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Platforms

7705 SAR Gen 2

shutdown

Syntax

shutdown

[**no**] **shutdown**

Context

[Tree] (config>service>vpls>bgp-evpn>mpls shutdown)

[Tree] (config>service>epipe>bgp-evpn>mpls shutdown)

Full Context

configure service vpls bgp-evpn mpls shutdown

configure service epipe bgp-evpn mpls shutdown

Description

This command controls the administrative state of EVPN-MPLS, EVPN-VXLAN, or EVPN-SRv6 in the service.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>service>vpls>proxy-nd shutdown)

[Tree] (config>service>vpls>proxy-arp shutdown)

Full Context

configure service vpls proxy-nd shutdown

configure service vpls proxy-arp shutdown

Description

This command enables and disables the proxy-ARP and proxy-nd functionality. ARP/GARP/ND messages will be snooped and redirected to the CPM for lookup in the proxy-ARP/proxy-ND table. The proxy-ARP/proxy-ND table is populated with IP->MAC pairs received from different sources (EVPN, static, dynamic). When the **shutdown** command is issued, it flushes the dynamic/EVPN dup proxy-ARP/proxy-ND table entries and instructs the system to stop snooping ARP/ND frames. All the static entries are kept in the table as *inactive*, regardless of their previous *Status*.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>service>epipe>site shutdown)

[Tree] (config>service>epipe>spoke-sdp shutdown)

[Tree] (config>service>epipe shutdown)

[Tree] (config>service>epipe>sap shutdown)

Full Context

configure service epipe site shutdown

configure service epipe spoke-sdp shutdown

configure service epipe shutdown

configure service epipe sap shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described as follows in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>service>epipe>bgp-vpws shutdown)

Full Context

configure service epipe bgp-vpws shutdown

Description

This command administratively enables/disables the local BGP VPWS instance. On de-activation an MP-UNREACH-NLRI is sent for the local NLRI.

The **no** form of this command enables the BGP VPWS addressing and the related BGP advertisement. The associated BGP VPWS MP-REACH-NLRI will be advertised in an update message and the corresponding received NLRIs must be considered to instantiate the data plane.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>service>vpls>bgp-vpls shutdown)

Full Context

configure service vpls bgp-vpls shutdown

Description

This command administratively enables/disables the local BGP VPLS instance. On de-activation an MP-UNREACH-NLRI must be sent for the local NLRI.

The **no** form of this command enables the BGP VPLS addressing and the related BGP advertisement. The associated BGP VPLS MP-REACH-NLRI will be advertised in an update message and the corresponding received NLRIs must be considered to instantiate the data plane. RT, RD usage: same as in the BGP AD solution, if the values are not configured here, the value of the VPLS-id from under the bgp-ad node is used. If VPLS-id value is not configured either the MH site cannot be activated – i.e. no shutdown returns an error. Same applies if a pseudowire template is not specified under the BGP node.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>service>ies>if>ipv6>secure-nd shutdown)

Full Context

configure service ies interface ipv6 secure-nd shutdown

Description

This command enables or disables Secure Neighbor Discovery (SeND) on the interface.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>security>tacplus shutdown)

[Tree] (config>service>vprn>aaa>rmt-srv>tacplus shutdown)

Full Context

configure system security tacplus shutdown

configure service vprn aaa remote-servers tacplus shutdown

Description

This command administratively disables the TACACS+ protocol operation. Shutting down the protocol does not remove or change the configuration other than the administrative state.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables the protocol which is the default state.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>service>vprn>static-route-entry>next-hop shutdown)

[Tree] (config>service>vprn>static-route-entry>black-hole shutdown)

[Tree] (config>service>vprn>static-route-entry>ipsec-tunnel shutdown)

[Tree] (config>service>vprn>static-route-entry>grt shutdown)

[Tree] (config>service>vprn>static-route-entry>indirect shutdown)

Full Context

configure service vprn static-route-entry next-hop shutdown

configure service vprn static-route-entry black-hole shutdown

configure service vprn static-route-entry ipsec-tunnel shutdown

configure service vprn static-route-entry grt shutdown

configure service vprn static-route-entry indirect shutdown

Description

This command causes the static route to be placed in an administratively down state and removed from the active route-table

Default
no shutdown

Platforms
7705 SAR Gen 2

shutdown

Syntax
[**no**] **shutdown**

Context
[Tree] (config>service>vprn>if>send shutdown)

Full Context
configure service vprn interface ipv6 secure-nd shutdown

Description
This command enables or disables Secure Neighbor Discovery (SeND) on the interface.

Platforms
7705 SAR Gen 2

shutdown

Syntax
[**no**] **shutdown**

Context
[Tree] (config>router>ldp>targ-session>peer shutdown)

[Tree] (config>router>ldp>targ-session>peer-template shutdown)

[Tree] (config>router>ldp>if-params>if shutdown)

[Tree] (config>router>ldp>if-params>if>ipv6 shutdown)

[Tree] (config>router>ldp>aggregate-prefix-match shutdown)

[Tree] (config>router>ldp>if-params>if>ipv4 shutdown)

[Tree] (config>router>ldp shutdown)

Full Context
configure router ldp targeted-session peer shutdown

configure router ldp targeted-session peer-template shutdown

configure router ldp interface-parameters interface shutdown

configure router ldp interface-parameters interface ipv6 shutdown

configure router ldp aggregate-prefix-match shutdown

configure router ldp interface-parameters interface ipv4 shutdown

configure router ldp shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted. For an LDP interface, the **shutdown** command exists under the main interface context and under each of the interface IPv4 and IPv6 contexts.

- **shutdown** under the **interface** context brings down both IPv4 and IPv6 Hello adjacencies and stops Hello transmission in both contexts.
- **shutdown** under the **interface** IPv4 or IPv6 contexts brings down the Hello adjacency and stops Hello transmission in that context only.

The user can also delete the entire IPv4 or IPv6 context under the interface with the **no ipv4** or **no ipv6** command which in addition to bringing down the Hello adjacency will delete the configuration.

The **no** form of this command administratively enables an entity.

Unlike other commands and parameters where the default state is not indicated in the configuration file, the **shutdown** and **no shutdown** states are always indicated in system generated configuration files.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>ldp>targeted-session>auto-tx>ipv4 shutdown)

[Tree] (config>router>ldp>targeted-session>auto-rx>ipv4 shutdown)

Full Context

configure router ldp targeted-session auto-tx ipv4 shutdown

configure router ldp targeted-session auto-rx ipv4 shutdown

Description

This command administratively disables the capabilities associated with automatically sending targeted Hello messages through the **auto-tx** command or processing targeted Hello messages through the **auto-rx** command.

The **no** form of this command administratively enables the capabilities associated with the **auto-tx** and **auto-rx** commands.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>mpls>lsp>secondary shutdown)

[Tree] (config>router>mpls>interface shutdown)

[Tree] (config>router>mpls shutdown)

[Tree] (config>router>mpls>lsp>primary shutdown)

Full Context

configure router mpls lsp secondary shutdown

configure router mpls interface shutdown

configure router mpls shutdown

configure router mpls lsp primary shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

MPLS is not enabled by default and must be explicitly enabled (**no shutdown**).

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command places the entity into an administratively enabled state.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>mpls>pce-initiated-lsp>sr-te shutdown)

Full Context

configure router mpls pce-initiated-lsp sr-te shutdown

Description

This command administratively enables or disables the **sr-te** context for PCE initiated LSPs. A shutdown of the **sr-te** context under **pce-initiated-lsp** causes an error to be generated for new PCInitate messages, and existing PCE-initiated LSPs are taken to the **oper-down** state.

The **no** form of this command administratively enables the **sr-te** context for PCE initiated LSP.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>mpls>if>label-map shutdown)

Full Context

configure router mpls interface label-map shutdown

Description

This command disables the label map definition. This drops all packets that match the specified *in-label* specified in the **label-map** *in-label* command.

The **no** form of this command administratively enables the defined label map action.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>mpls>lsp-template shutdown)

[Tree] (config>router>mpls>lsp shutdown)

Full Context

configure router mpls lsp-template shutdown

configure router mpls lsp shutdown

Description

This command disables the existing LSP including the primary and any standby secondary paths.

To shutdown only the primary enter the **config router mpls lsp** *lsp-name* **primary** *path-name* **shutdown** command.

To shutdown a specific standby secondary enter the **config router mpls lsp** *lsp-name* **secondary** *path-name* **shutdown** command. The existing configuration of the LSP is preserved.

Use the **no** form of this command to restart the LSP. LSPs are created in a shutdown state. Use this command to administratively bring up the LSP.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>mpls>path shutdown)

Full Context

configure router mpls path shutdown

Description

This command disables the existing LSPs using this path. All services using these LSPs are affected. Binding information, however, is retained in those LSPs. Paths are created in the **shutdown** state.

The **no** form of this command administratively enables the path. All LSPs, where this path is defined as primary or defined as standby secondary, are (re)established.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>pcep>pcc shutdown)

Full Context

configure router pcep pcc shutdown

Description

This command administratively disables the PCC or PCE process.

The following PCE parameters can only be modified when the PCEP session is shut down:

- **local-address**
- **keepalive**
- **dead-timer**

The **unknown-message-rate** PCE parameter can be modified without shutting down the PCEP session.

The following PCC parameters can only be modified when the PCEP session is shut down:

- **local-address**
- **keepalive**
- **dead-timer**
- **peer**

The following PCC parameters can be modified without shutting down the PCEP session:

- **report-path-constraints**
- **unknown-message-rate**

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>rsvp shutdown)

[Tree] (config>router>rsvp>interface shutdown)

Full Context

configure router rsvp shutdown

configure router rsvp interface shutdown

Description

This command disables the RSVP protocol instance or the RSVP-related functions for the interface. The RSVP configuration information associated with this interface is retained. When RSVP is administratively disabled, all the RSVP sessions are torn down. The existing configuration is retained.

The **no** form of this command administratively enables RSVP on the interface.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>mpls>fwd-policies>fwd-policy>nh-grp shutdown)

Full Context

configure router mpls forwarding-policies forwarding-policy next-hop-group shutdown

Description

This command shuts down an NHG entry in a forwarding policy.

When an NHG is shut down, it is removed from the data path entry of the forwarding policy.

The **no** form of this command brings up an NHG entry in a forwarding policy.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>mpls>fwd-policies>fwd-policy shutdown)

Full Context

configure router mpls forwarding-policies forwarding-policy shutdown

Description

This command shuts down the forwarding policy.

The **no** form of this command enables the forwarding policy.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>mpls>fwd-policies shutdown)

Full Context

configure router mpls forwarding-policies shutdown

Description

This command shuts down the **forwarding-policies** context; causing all forwarding policies to be removed from the data path, however they remain in the MPLS forwarding database.

The **no** form of this command enables the **forwarding-policies** context.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ipsec>tunnel-group shutdown)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel shutdown)

[Tree] (config>service>ies>if>sap>ipsec-gw>dhcp6 shutdown)

[Tree] (config>ipsec>client-db shutdown)

[Tree] (config>service>vprn>if>sap>ipsec-gw>dhcp6 shutdown)

[Tree] (config>service>vprn>if>sap>ipsec-gw>lcl-addr-assign shutdown)

[Tree] (config>service>ies>if>sap>ipsec-gw shutdown)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel shutdown)

[Tree] (config>isa>tunnel-grp shutdown)

[Tree] (config>service>ies>if>sap>ipsec-gw>dhcp shutdown)

[Tree] (config>service>ies>if>sap>ipsec-gw>lcl-addr-assign shutdown)

[Tree] (config>service>vprn>if>sap>ipsec-gw shutdown)

[Tree] (config>service>vprn>if>sap>ip-tunnel shutdown)

[Tree] (config>service>ies>if>sap>ip-tunnel shutdown)

[Tree] (config>ipsec>client-db>client shutdown)

[Tree] (config>ipsec>cert-profile shutdown)

[Tree] (config>service>vprn>if>sap>ipsec-gw>dhcp shutdown)

Full Context

configure redundancy multi-chassis peer mc-ipsec tunnel-group shutdown

configure service ies interface ipsec ipsec-tunnel shutdown

configure service ies interface sap ipsec-gw dhcp6 shutdown

configure ipsec client-db shutdown

configure service vprn interface sap ipsec-gw dhcp6 shutdown

configure service vprn interface sap ipsec-gw local-address-assignment shutdown

configure service ies interface sap ipsec-gw shutdown

configure service vprn interface ipsec ipsec-tunnel shutdown

configure isa tunnel-group shutdown

configure service ies interface sap ipsec-gw dhcp shutdown

configure service ies interface sap ipsec-gw local-address-assignment shutdown

configure service vprn interface sap ipsec-gw shutdown

configure service vprn interface sap ip-tunnel shutdown

configure service ies interface sap ip-tunnel shutdown

configure ipsec client-db client shutdown

configure ipsec cert-profile shutdown

configure service vprn interface sap ipsec-gw dhcp shutdown

Description

This command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The **shutdown** command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>security>pki>ca-prof>auto-crl-update shutdown)

Full Context

configure system security pki ca-profile auto-crl-update shutdown

Description

This command disables the auto CRL update.

The **no** form of this command enables an auto CRL update. Upon **no shutdown**, if the configured CRL file does not exist, is invalid or is expired or if the schedule-type is next-update-based and current time passed (Next-Update_of_existing_CRL - pre-update-time), then system will start downloading CRL right away.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>nat>outside>pool shutdown)

[Tree] (config>service>vprn>nat>outside>pool shutdown)

[Tree] (config>router>nat>inside>deterministic>address-map shutdown)

[Tree] (config>service>vprn>nat>inside>deterministic>address-map shutdown)

[Tree] (config>isa>nat-group shutdown)

Full Context

configure router nat outside pool shutdown

configure service vprn nat outside pool shutdown

configure router nat inside deterministic address-map shutdown

configure service vprn nat inside deterministic address-map shutdown

configure isa nat-group shutdown

Description

This command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The **shutdown** command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>mirror>mirror-dest shutdown)

[Tree] (config>mirror>mirror-source shutdown)

[Tree] (config>service>vprn>ip-mirror-interface>spoke-sdp shutdown)

[Tree] (config>service>vprn>ip-mirror-interface shutdown)

Full Context

configure mirror mirror-dest shutdown

configure mirror mirror-source shutdown

configure service vprn ip-mirror-interface spoke-sdp shutdown

configure service vprn ip-mirror-interface shutdown

Description

The **shutdown** command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The **shutdown** command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

The **no** form of this command puts an entity into the administratively enabled state.

Default

See Special Cases below.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (debug>mirror-source shutdown)

Full Context

debug mirror-source shutdown

Description

This command enables mirror source debugging.

The **no** form of this command clears mirror source information.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>test-oam>twamp>server shutdown)

[Tree] (config>oam-pm>session>ip>twamp-light shutdown)

[Tree] (config>oam-pm>bin-group shutdown)

[Tree] (config>saa>test shutdown)

Full Context

configure test-oam twamp server shutdown

configure oam-pm session ip twamp-light shutdown

configure oam-pm bin-group shutdown

configure saa test shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Entities are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the entity becomes administratively up and then tries to enter the operationally up state.

The **no** form of this command administratively enables the entity.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>service>vprn>twamp-light>reflector shutdown)

[Tree] (config>router>twamp-light>reflector shutdown)

Full Context

configure service vprn twamp-light reflector shutdown

configure router twamp-light reflector shutdown

Description

This command disables or enables TWAMP Light functionality within the context where the configuration exists, either the base router instance or the service. Enabling the base router context enables the IES prefix list since the IES service uses the configuration under the base router instance.

The **no** form of this command allows the router instance or the service to accept TWAMP Light packets for processing.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>filter>redirect-policy shutdown)

[Tree] (config>filter>log>summary shutdown)

[Tree] (config>filter>redirect-policy>destination shutdown)

Full Context

configure filter redirect-policy shutdown

configure filter log summary shutdown

configure filter redirect-policy destination shutdown

Description

Administratively enables/disabled (AdminUp/AdminDown) an entity. Downing an entity does not change, reset or remove any configuration settings or statistics. Many objects must be shutdown before they may be deleted.

The **shutdown** command administratively downs an entity. Administratively downing an entity changes the operational state of the entity to down.

Unlike other commands and parameters where the default state will not be indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

The **no** form of the command puts an entity into the administratively enabled state.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>if>ipv6>secure-nd shutdown)

Full Context

configure router interface ipv6 secure-nd shutdown

Description

This command enables or disables Secure Neighbor Discovery (SeND) on the interface.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>vrrp>policy shutdown)

[Tree] (config>router>if>ipv6>vrrp shutdown)

[Tree] (config>router>if>vrrp shutdown)

Full Context

configure vrrp policy shutdown

configure router interface ipv6 vrrp shutdown

configure router interface vrrp shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables an entity.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown** [**active**] [**standby**]

[**no**] **shutdown** [*cflash-id*]

Context

[Tree] (file shutdown)

Full Context

file shutdown

Description

This command shuts down (unmounts) the specified CPM(s).

Use the **no shutdown** [**active**] [**standby**] command to enable one or both CPM.

Use the **no shutdown** [*cflash-id*] command to enable a compact flash (cf1:, cf2:, or cf3:) on the CPM/CCM. The **no shutdown** command can be issued for a specific slot when no compact flash is present. When a flash card is installed in the slot, the card will be activated upon detection.

In redundant systems, use the **no shutdown** command on cf3: on both SF/CPMs or CCMs in order to facilitate synchronization. See the **config>redundancy synchronize** command.

Note:
The **shutdown** command must be issued prior to removing a flash card. If no parameters are specified, then the drive referred to by the current working directory will be shut down.

LED Status Indicators

Table 92: LED Status Indicators lists the possible states for the compact flash and their LED status indicators.

Table 92: LED Status Indicators

State	Description
Operational	If a compact flash is present in a drive and operational (no shutdown), the respective LED is lit green. The LED flickers when the compact flash is accessed. **Note:** Do not remove the compact flash during a read/write operation.

State	Description
Flash defective	If a compact flash is defective, the respective LED blinks amber to reflect the error condition and a trap is raised.
Flash drive shut down	When the compact flash drive is shut down and a compact flash present, the LED is lit amber. In this state, the compact flash can be ejected.
No compact flash present, drive shut down	If no compact flash is present and the drive is shut down the LED is unlit.
No compact flash present, drive enabled	If no compact flash is present and the drive is not shut down the LED is unlit.
Ejecting a compact flash	The compact flash drive should be shut down before ejecting a compact flash card. The LED should turn to solid (not blinking) amber. This is the only mode to safely remove the flash card. If a compact flash drive is not shut down before a compact flash is ejected, the LED blinks amber for approximately 5 seconds before shutting off.

The **shutdown** or **no shutdown** state is not saved in the configuration file. Following a reboot all compact flash drives are in their default state.

Default

no shutdown

Parameters

cflash-id

Specifies the compact flash slot ID to be shut down or enabled. If *cflash-id* is specified, the drive is shut down or enabled. If no *cflash-id* is specified, the drive referred to by the current working directory is assumed. If a slot number is not specified, then the active CPM is assumed.

Values cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Default the current compact flash device

active

Specifies that all drives on the active CPM are shutdown or enabled.

standby

Specifies that all drives on the standby CPM are shutdown or enabled.

When both **active** and **standby** keywords are specified, then all drives on both CPM are shutdown.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>management-interface>remote-management shutdown)

[Tree] (config>system>management-interface>remote-management>manager shutdown)

Full Context

configure system management-interface remote-management shutdown

configure system management-interface remote-management manager shutdown

Description

This command administratively disables remote management.

The **no** form of this command administratively enables remote management.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>alarms shutdown)

Full Context

configure system alarms shutdown

Description

This command enables or disables the Facility Alarm functionality. When enabled, the Facility Alarm sub-system tracks active and cleared facility alarms and controls the Alarm LEDs on the CPMs. When Facility Alarm functionality is enabled, the alarms are viewed using the show system alarms command(s).

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>log>accounting-policy shutdown)

[Tree] (config>log>event-trigger>event>trigger-entry shutdown)

[Tree] (config>log>log-id shutdown)

[Tree] (config>log>event-handling>handler shutdown)

[Tree] (config>log>event-handling>handler>action-list>entry shutdown)

[Tree] (config>log>event-trigger>event shutdown)

Full Context

configure log accounting-policy shutdown

configure log event-trigger event trigger-entry shutdown

configure log log-id shutdown

configure log event-handling handler shutdown

configure log event-handling handler action-list entry shutdown

configure log event-trigger event shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables an entity.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>netconf>listen shutdown)

Full Context

configure system netconf listen shutdown

Description

This command disables the NETCONF server. The **shutdown** command is blocked if there are any active NETCONF sessions. Use the **admin disconnect** command to disconnect all NETCONF sessions before shutting down the NETCONF service.

The **no** form of this command enables the NETCONF server.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter shutdown)

[Tree] (config>system>security>keychain>direction>uni>send>entry shutdown)

[Tree] (config>system>security>keychain>direction>uni>receive>entry shutdown)

[Tree] (config>system>security>keychain>direction>bi>entry shutdown)

[Tree] (config>system>security>keychain shutdown)

[Tree] (config>system>security>mgmt-access-filter>ip-filter shutdown)

Full Context

configure system security management-access-filter ipv6-filter shutdown

configure system security keychain direction uni send entry shutdown

configure system security keychain direction uni receive entry shutdown

configure system security keychain direction bi entry shutdown

configure system security keychain shutdown

configure system security management-access-filter ip-filter shutdown

Description

This command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command puts an entity into the administratively enabled state.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>security>pki>ca-profile shutdown)

Full Context

configure system security pki ca-profile shutdown

Description

Use this command to enable or disable the ca-profile. The system verifies the configured cert-file and crl-file. If the verification fails, then the **no shutdown** command fails.

The ca-profile in a **shutdown** state cannot be used in certificate authentication.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>security>ssh>key-re-exchange>client shutdown)

[Tree] (config>system>security>ssh>key-re-exchange>server shutdown)

Full Context

configure system security ssh key-re-exchange client shutdown

configure system security ssh key-re-exchange server shutdown

Description

This command stops the key exchange. It sets the minutes and bytes to infinity so there will not be any key exchange during the PDU transmission.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>security>ldap>server shutdown)

[Tree] (config>system>security>ldap shutdown)

Full Context

configure system security ldap server shutdown

configure system security ldap shutdown

Description

In the **ldap** context, this command enables or disabled LDAP protocol operations.

In the **server** context, this command enables or disables the LDAP server. To perform **no shutdown**, an LDAP server address is required. To change the address, the user first needs to shut down the server.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>grpc>gnmi shutdown)

Full Context

configure system grpc gnmi shutdown

Description

This command stops the gNMI service.

The **no** form of this command starts the gNMI service.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>grpc shutdown)

Full Context

configure system grpc shutdown

Description

This command stops the gRPC server. This closes all of the associated TCP connections and immediately purges all RIB entries that were programmed using the RibApi Service.

The **shutdown** command is not blocked if there are active gRPC sessions. Shutting down gRPC will terminate all active gRPC sessions.

The **no** form of this command starts the gRPC server.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>grpc>tcp-keepalive shutdown)

Full Context

configure system grpc tcp-keepalive shutdown

Description

This command stops the TCP keepalives from being sent to all gRPC clients.

The **no** form of this command restarts the sending of TCP keepalives to all gRPC clients.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>snmp>streaming shutdown)

Full Context

configure system snmp streaming shutdown

Description

This command administratively disables proprietary SNMP request/response bundling and TCP-based transport mechanism for optimizing network management of the router nodes.

The **no** form of the command administratively re-enables SNMP request/response bundling and TCP-based transport mechanism.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>snmp shutdown)

Full Context

configure system snmp shutdown

Description

This command administratively disables SNMP agent operations. System management can then only be performed using the command line interface (CLI). Shutting down SNMP does not remove or change configuration parameters other than the administrative state. This command does not prevent the agent from sending SNMP notifications to any configured SNMP trap destinations. SNMP trap destinations are configured under the **config>log>snmp-trap-group** context.

This command is automatically invoked in the event of a reboot when the processing of the configuration file fails to complete or when an SNMP persistent index file fails while the **bof persist on** command is enabled.

The **no** form of the command administratively enables SNMP which is the default state.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>security>tls>cert-profile shutdown)

Full Context

configure system security tls cert-profile shutdown

Description

This command disables the certificate profile. When the certificate profile is disabled, it will not be sent to the TLS server.

The **no** form of the command enables the certificate profile and allows it to be sent to the TLS server.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>system>security>tls>client-tls-profile shutdown)

[Tree] (config>system>security>tls>server-tls-profile shutdown)

Full Context

configure system security tls client-tls-profile shutdown

configure system security tls server-tls-profile shutdown

Description

This command administratively enables or disables the TLS profile. If the TLS profile is shut down, the TLS operational status will be down. Therefore, if the TLS profile is shut down, any application using TLS should not attempt to send any PDUs.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>bgp>group shutdown)

[Tree] (config>router>bgp>segment-routing shutdown)

[Tree] (config>router>bgp shutdown)

[Tree] (config>router>bgp>group>neighbor shutdown)

Full Context

configure router bgp group shutdown

configure router bgp segment-routing shutdown

configure router bgp shutdown

configure router bgp group neighbor shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables an entity.

Unlike other commands and parameters where the default state is not indicated in the configuration file, the **shutdown** and **no shutdown** states are always indicated in system generated configuration files.

Default administrative states for services and service entities are described in Special Cases.

The **no** form of this command places an entity in an administratively enabled state.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>isis>segment-routing shutdown)

[Tree] (config>router>isis>segm-rtng>mapping-server shutdown)

[Tree] (config>router>isis>igp-shortcut shutdown)

[Tree] (config>router>isis>interface shutdown)

[Tree] (config>router>isis shutdown)

Full Context

configure router isis segment-routing shutdown

configure router isis segment-routing mapping-server shutdown

configure router isis igp-shortcut shutdown

configure router isis interface shutdown

configure router isis shutdown

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables an entity.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>isis>flex-algos shutdown)

Full Context

configure router isis flexible-algorithms shutdown

Description

This command enables IS-IS flexible algorithms. If it is enabled with the **no shutdown** command the router starts supporting the flexible algorithms IGP LSDB extensions. Flexible algorithm IGP LSDB extensions are by default not enabled.

The **no** form of this command enables the router to support flexible algorithms.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>ospf>flex-algos shutdown)

Full Context

configure router ospf flexible-algorithms shutdown

Description

This command enables OSPFv2 flexible algorithms. If **no shutdown** is configured, the router enables support for the flexible algorithms IGP LSDB extensions. Flexible algorithm IGP LSDB extensions are disabled by default.

The **no** form of this command enables the router to support flexible algorithms.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>ospf3>area>virtual-link shutdown)

[Tree] (config>router>ospf>igp-shortcut shutdown)

[Tree] (config>router>ospf shutdown)

[Tree] (config>router>ospf>area>virtual-link shutdown)

[Tree] (config>router>ospf3 shutdown)

[Tree] (config>router>ospf3>area>interface shutdown)

[Tree] (config>router>ospf>segm-rtng shutdown)

[Tree] (config>router>ospf>segm-rtng>mapping-server shutdown)

Full Context

configure router ospf3 area virtual-link shutdown

configure router ospf igp-shortcut shutdown

configure router ospf shutdown

configure router ospf area virtual-link shutdown

configure router ospf3 shutdown

configure router ospf3 area interface shutdown

configure router ospf segment-routing shutdown

configure router ospf segment-routing mapping-server shutdown

Description

The **shutdown** command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. The operational state of the entity is disabled as well as the operational state of any entities contained within.

Many objects must be shut down before they may be deleted. Many entities must be explicitly enabled using the **no shutdown** command.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

The **no** form of this command puts an entity into the administratively enabled state.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>rip>group>neighbor shutdown)

[Tree] (config>router>rip shutdown)

[Tree] (config>router>ripng>group>neighbor shutdown)

[Tree] (config>router>ripng shutdown)

[Tree] (config>router>ripng>group shutdown)

[Tree] (config>router>rip>group shutdown)

Full Context

configure router rip group neighbor shutdown

configure router rip shutdown

configure router ripng group neighbor shutdown

configure router ripng shutdown

configure router ripng group shutdown

configure router rip group shutdown

Description

This command administratively disables an entity. Downing an entity does not change, reset or remove any configuration settings or statistics. Many objects must be shutdown before they may be deleted.

The **shutdown** command administratively downs an entity. Administratively downing an entity changes the operational state of the entity to down and the operational state of any entities contained within the administratively down entity.

Unlike other commands and parameters where the default state will not be indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

The **no** form of the command puts an entity into the administratively enabled state.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>segment-routing>maintenance-policy shutdown)

Full Context

configure router segment-routing maintenance-policy shutdown

Description

This command deactivates all segment routing policies and removes the associated entries from the forwarding plane of the router.

The **no** form of this command enables all segment routing policies so that they can be revalidated and reinstalled as necessary.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>segment-routing>sr-policies shutdown)

Full Context

configure router segment-routing sr-policies shutdown

Description

This command deactivates all segment routing policies and removes the associated entries from the forwarding plane of the router.

It is necessary to execute this shutdown if you want to make a change to the reserved-label-block reference.

The **no** form of this command enables all segment routing policies so that they can be revalidated and reinstalled as necessary.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (conf>router>segment-routing>sr-policies>policy>seg-list shutdown)

Full Context

configure router segment-routing sr-policies static-policy segment-list shutdown

Description

This command deactivates a segment-list. If this is done on an active policy with more than one segment list, then traffic forwarded by the policy will be diverted to the remaining segment-lists.

The **no** form of this command enables the segment list so that it can be validated and installed as necessary.

Default

shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (conf>router>segment-routing>sr-policies>policy shutdown)

Full Context

configure router segment-routing sr-policies static-policy shutdown

Description

This command deactivates the associated static policy and causes another policy for the same (color, endpoint) combination to be promoted as the active path, assuming there is another valid policy.

It is necessary to execute this shutdown if you want to make critical configuration changes to the static policy.

The **no** form of this command enables the static policy so that it can be validated and installed as necessary.

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>port>ethernet>dot1x>per-host-authentication shutdown)

Full Context

configure port ethernet dot1x per-host-authentication shutdown

Description

This command administratively configures per-host authentication on the port.

The **no** form of this command administratively enables per-host authentication on the port.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>port>ethernet>dot1x shutdown)

Full Context

configure port ethernet dot1x shutdown

Description

This command administratively configures the 802.1x functionality (consisting of packet extraction and processing on the CPM) on the port.

The **no** form of this command administratively enables the 802.1x functionality on the port.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>bgp>egress-peer-engineering shutdown)

Full Context

configure router bgp egress-peer-engineering shutdown

Description

This command administratively enables or disables BGP-EPE. If enabled, peer node SIDs and peer adjacency SIDs are advertised in BGP-LS.

The **no** form of this command places the entity into an administratively enabled state and prevents peer node SIDs and peer adjacency SIDs from being advertised in BGP-LS.

Default

no shutdown

Platforms

7705 SAR Gen 2

shutdown

Syntax

[**no**] **shutdown**

Context

[Tree] (config>router>bgp>group>egress-engineering shutdown)

Full Context

configure router bgp group egress-engineering shutdown

Description

This command administratively enables or disable egress engineering on a BGP neighbor or group of neighbors.

If this command is enabled along with the **egress-peer-engineering** command in BGP, SIDs in the form of MPLS labels are allocated for the segments toward the neighbor and to all links (adjacencies). These adjacencies are then advertised in BGP LS.

The **no** form of this command places the entity into an administratively enabled state.

Default

no shutdown

Platforms

7705 SAR Gen 2

28 s Commands – Part II

28.1 sid

sid

Syntax

sid label *value*

Context

[Tree] (config>router>isis>segm-rtng>adjacency-set sid)

[Tree] (config>router>ospf>segm-rtng>adjacency-set sid)

Full Context

configure router isis segment-routing adjacency-set sid

configure router ospf segment-routing adjacency-set sid

Description

This command allows a static SID value to be assigned to an adjacency set in IS-IS or OSPF segment routing.

The **label** option specifies the value is assigned to an MPLS label.

The **no** form of this command removes the adjacency SID.

Parameters

label value

Specifies the value of adjacency SID label.

Values 18432 to 524287 | 1048575 (FP4 only)

Platforms

7705 SAR Gen 2

28.2 sid-map

sid-map

Syntax

sid-map node-sid {**index** *value* [**range** *value*]} **prefix** {{*ip-address*/*mask*} | {*ip-address*} {*netmask*}} [**set-flags** {**s**}] [**level** { **1** | **2** | **1/2**}] [**clear-n-flag**]

no sid-map node-sid index *value*

Context

[Tree] (config>router>isis>segm-rtng>mapping-server sid-map)

Full Context

configure router isis segment-routing mapping-server sid-map

Description

This command configures the Segment Routing mapping server database in IS-IS.

The user enters the node SID index for one or a range of prefixes by specifying the first index value and optionally a range value can be entered. The default value for the range option is 1. Only the first prefix in a consecutive range of prefixes must be entered. The user can enter the first prefix with a mask lower than 32 and the SID or label binding TLV is advertised, but the routers will not resolve these prefix SIDs and will generate a trap.

By setting the S-flag, the user can indicate to the IS-IS routers in the rest of the network that the flooding scope of the SID or label binding TLV is the entire domain. In that case, a router receiving the TLV advertisement should leak it between ISIS levels. If leaked from level 2 to level 1, the D-flag must be set and once set the TLV cannot be leaked back into level 2. Otherwise, the S-flag is clear by default and the TLV must not be leaked by routers that receive the mapping server advertisement.

Note that the SR OS does not leak this TLV between IS-IS instances and does not support the multi-topology SID/Label Binding TLV format.

In addition, the user can specify the mapping server own flooding scope for the generated SID or label binding TLV using the **level** option. This option allows the user to narrow the flooding scope configured under the router IS-IS level-capability for a one or more SID or label binding TLVs if required. The default flooding scope of the mapping server is Layer 1 or Layer 2, which can be narrowed by the value configured under the router IS-IS level-capability.

The A-flag and M-flag are not supported by the mapping server feature. The mapping client ignores the flags.

Each time a prefix or a range of prefixes is configured in the SR mapping database in any routing instance, the router issues for this prefix or range of prefixes, a prefix-SID sub-TLV within a ISIS SID or label binding TLV in that instance. The flooding scope of the TLV from the mapping server is determined as explained above. No further check of the reachability of that prefix in the mapping server route table is performed. Additionally, no check is performed if the SID index is a duplicate of an existing prefix in the local IGP instance database or if the SID index is out of range with the local SRGB.

The **no** form of this command deletes the range of node SIDs beginning with the specified index value.

Parameters

index

Specifies the node SID index for the IS-IS prefix that is advertised in a SID/Label Binding TLV.

Values 0 to 4294967295

value

Specifies the node SID range for the IS-IS prefix that is advertised in a SID/Label Binding TLV.

Values 0 to 65535

ip-address/mask

Specifies the IP address and mask.

Values *ip-address:* **a.b.c.d.** (host bits must be 0)

mask: **0 to 32**

ip-address netmask

Specifies the IP address netmask.

Values **a.b.c.d.** (network bits all 1 and host bits all 0)

set-flags

Specifies the flooding scope of the SID/Label binding TLV.

Default **S-flag clear**

The TLV is not leaked by routers receiving the mapping server advertisement

level {1 | 2| 1/2}

Configures the mapping server own flooding scope for the generated SID/Label binding TLV.

Default 1/2

clear-n-flag

Specifies whether the node-sid flag (N-flag) should be cleared in a SID Label Binding TLV.

Platforms

7705 SAR Gen 2

sid-map

Syntax

sid-map node-sid index *index-value* [**range** *range-value*] **prefix** *ip-addressImask* [*netmask*]

sid-map node-sid index *index-value* [**range** *range-value*] **prefix** *ip-addressImask* [*netmask*] **scope** {**area** *area-id* | **as**}

no sid-map node-sid index *index-value*

Context

[Tree] (config>router>ospf>segm-rtng>mapping-server sid-map)

Full Context

configure router ospf segment-routing mapping-server sid-map

Description

This command configures the Segment Routing mapping server database in OSPF.

The user enters the node SID index for one or a range of prefixes by specifying the first index value and optionally a range value. The default value for the range option is 1. Only the first prefix in a consecutive range of prefixes must be entered. If the user enters the first prefix with a mask lower than 32, the OSPF Extended Prefix Range TLV is advertised but a router which receives it will not resolve SID and instead originates a trap.

The user specifies the mapping server own flooding scope for the generated OSPF Extended Prefix Range TLV using the scope option. There is no default value. If the scope is a specific area, then the TLV is flooded only in that area.

An ABR that propagates an intra-area OSPF Extended Prefix Range TLV flooded by the mapping server in that area into other areas, sets the inter-area flag (IA-flag). The ABR also propagates the TLV if received with the inter-area flag set from other ABR nodes but only from the backbone to leaf areas and not vice-versa. However, if the exact same TLV is advertised as an intra-area TLV in a leaf area, the ABR will not flood the inter-area TLV into that leaf area.

Note:
SR OS does not leak this TLV between OSPF instances.

Each time a prefix or a range of prefixes is configured in the SR mapping database in any routing instance, the router issues for this prefix, or range of prefixes, a prefix-SID sub-TLV within a OSPF Extended Prefix Range TLV in that instance. The flooding scope of the TLV from the mapping server is determined as previously explained. No further check of the reachability of that prefix in the mapping server route table is performed and no check if the SID index is duplicate with some existing prefix in the local IGP instance database or if the SID index is out of range with the local SRGB.

The **no** form of this command deletes the range of node SIDs beginning with the specified index value.

Default

no prefix-sid-range

Parameters

index index-value

Specifies the index.

Values 0 to 4294967295

range range-value

Specifies the range.

Values 1 to 65535

prefix ip-address/mask

Specifies the IP address in dotted decimal notation.

Values ip-address/mask:

- ip-address a.b.c.d (host bits must be 0)

mask: 0 to 132

netmask

Specifies the netmask.

Values netmask — a.b.c.d (network bits all 1 and host bits all 0)

area area-id

Configures the mapping server own flooding scope for the generated OSPF Extended Prefix Range TLV.

Values ip-address | 0 to 4294967295

Platforms

7705 SAR Gen 2

28.3 sid-protection

sid-protection

Syntax

[**no**] **sid-protection**

Context

[Tree] (config>router>isis>interface sid-protection)

Full Context

configure router isis interface sid-protection

Description

This command enables or disables adjacency SID protection by LFA and remote LFA.

While LFA and remote LFA Fast-Reroute (FRR) protection is enabled for all node SIDs and local adjacency SIDs when the user enables the **loopfree-alternates** option in IS-IS or OSPF at the LER and LSR, there are applications where the user wants traffic to never divert from the strict hop computed by CSPF for

a SR-TE LSP. In that case, the user can disable protection for all adjacency SIDs formed over a given network IP interface using this command.

The protection state of an adjacency SID is advertised in the B-FLAG of the IS-IS or OSPF Adjacency SID sub-TLV.

Default

sid-protection

Platforms

7705 SAR Gen 2

sid-protection

Syntax

[**no**] **sid-protection**

Context

[Tree] (config>router>ospf>area>interface sid-protection)

Full Context

configure router ospf area interface sid-protection

Description

This command enables or disables adjacency SID protection by LFA and remote LFA.

LFA and remote LFA Fast-Reroute (FRR) protection is enabled for all node SIDs and local adjacency SIDs when the user enables the **loopfree-alternate** option in IS-IS or OSPF at the LER and LSR. However, may be applications where the user never wants traffic to divert from the strict hop computed by CSPF for an SR-TE LSP. In this case, the user can disable protection for all adjacency SIDs formed over a particular network IP interface using this command.

The protection state of an adjacency SID is advertised in the B-FLAG of the IS-IS or OSPF Adjacency SID sub-TLV.

Default

sid-protection

Platforms

7705 SAR Gen 2

28.4 signaling

signaling

Syntax

signaling *signaling*

Context

[Tree] (config>service>epipe>spoke-sdp-fec signaling)

Full Context

configure service epipe spoke-sdp-fec signaling

Description

This command enables a user to configure this router as the active or passive T-PE for signaling this MS-PW, or to automatically select whether this T-PE is active or passive based on the prefix. In an active role, this endpoint initiates MS-PW signaling without waiting for a T-LDP label mapping message to arrive from the far end T-PE. In a passive role, it will wait for the initial label mapping message from the far end before sending a label mapping for this end of the PW. In auto mode, if the SAII has the greater prefix value, then the router will initiate MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAII has the greater value prefix, then the router will assume that the far end T-PE will initiate MS-PW signaling and will wait for that label mapping message before responding with a T-LDP label mapping message for the MS-PW in the reverse direction.

The **no** form of this command means that the router T-PE automatically selects the which router will initiate MS-PW signaling based on the prefix values configured in the SAII and TAII of the spoke SDP, as previously described.

Default

signaling auto

Parameters

signaling

Configures this router as the active T-PE for signaling this MS-PW.

Values auto, master

Platforms

7705 SAR Gen 2

signaling

Syntax

signaling {**off** | **tldp** | **bgp**}

Context

[Tree] (config>service>sdp signaling)

Full Context

configure service sdp signaling

Description

This command specifies the signaling protocol used to obtain the ingress and egress pseudowire labels in frames transmitted and received on the SDP. When signaling is *off* then labels are manually configured when the SDP is bound to a service. The signaling value can only be changed while the administrative status of the SDP is down. Additionally, the signaling can only be changed on an SDP if that SDP is not in use by BGP-AD or BGP-VPLS. BGP signaling can only be enabled if that SDP does not already have pseudowires signaled over it.

Note:
If the **tldp** option is selected as the mechanism for exchanging service labels over an MPLS or GRE SDP and the T-LDP session is automatically established, an explicit T-LDP session that is subsequently configured takes precedence over the automatic T-LDP session. However, if the explicit, manually-configured session is then removed, the system does not revert to the automatic session and the automatic session is also deleted. To address this, recreate the T-LDP session by disabling and re-enabling the SDP using the **shutdown** and **no shutdown** commands.

The **no** form of this command is not applicable. To modify the signaling configuration, the SDP must be administratively shut down and then the signaling parameter can be modified and re-enabled.

Default

signaling tldp

Parameters

off

Ingress and egress signal auto-labeling is not enabled. If this parameter is selected, then each service using the specified SDP must manually configure VPN labels. This configuration is independent of the SDP's transport type, GRE, MPLS (RSVP or LDP).

tldp

Ingress and egress pseudowire signaling using T-LDP is enabled. Default value used when BGP AD automatically instantiates the SDP.

bgp

Ingress and egress pseudowire signaling using BGP is enabled. Default value used when BGP VPLS automatically instantiates the SDP.

Platforms
7705 SAR Gen 2

28.5 signature-list

signature-list

Syntax
signature-list *name*
no signature-list

Context
[Tree] (config>system>security>tls>client-tls-profile signature-list)

Full Context
configure system security tls client-tls-profile signature-list

Description
This command assigns an existing TLS 1.3 signature list to the TLS client profile.

The **no** form of this command removes the signature list from the client profile.

Default
no signature-list

Parameters
name
Specifies the name of the signature list, up to 32 characters.

Platforms
7705 SAR Gen 2

signature-list

Syntax
signature-list *name*
no signature-list

Context
[Tree] (config>system>security>tls>server-tls-profile signature-list)

Full Context

configure system security tls server-tls-profile signature-list

Description

This command assigns an existing TLS 1.3 signature list to the TLS server profile.

The **no** form of this command removes the signature list from the server profile.

Default

no signature-list

Parameters

name

Specifies the name of the signature list, up to 32 characters.

Platforms

7705 SAR Gen 2

28.6 significant-change

significant-change

Syntax

significant-change *delta*

no significant-change

Context

[Tree] (config>log>acct-policy>cr significant-change)

Full Context

configure log accounting-policy custom-record significant-change

Description

This command configures the significant change required to generate the record. The custom record is only generated when the change in the reference counters equals or exceeds the configured (non-zero) significant change value. Only the reference counters for which there are corresponding counters configured under the related queues and policers are used for the significant change comparison. For reference queues and policers, the change applies to the sum of all configured reference queue and policer counters. When no reference counters are configured or **significant-change** is zero, the significant change reporting is not active.

Default

significant-change 0

Parameters

delta

Specifies the delta change (significant change) that is required for the custom record to be written to the XML file.

Values 0 to 4294967295 (For custom-record-aa-sub only values 0 or 1 are supported.)

Platforms

7705 SAR Gen 2

28.7 single-sfm-overload

single-sfm-overload

Syntax

single-sfm-overload [**holdoff-time** *holdoff-time*]

no single-sfm-overload

Context

[Tree] (config>service>vprn single-sfm-overload)

Full Context

configure service vprn single-sfm-overload

Description

This command configures OSPF, OSPFv3 and IS-IS to set overload when the router has fewer than the full set of SFMs functioning, which reduces forwarding capacity. Setting overload enables a router to still participate in exchanging routing information, but routes all traffic away from it.

The **no** form of this command configures the router to not set overload if an SFM fails.

Default

no single-sfm-overload

Parameters

holdoff-time

Specifies the delay between detecting SFM failures and setting overload.

Values 1 to 600 seconds

Default 0 seconds

Platforms

7705 SAR Gen 2

single-sfm-overload

Syntax

single-sfm-overload [**holdoff-time** *holdoff-time*]

no single-sfm-overload

Context

[Tree] (config>router single-sfm-overload)

Full Context

configure router single-sfm-overload

Description

This command configures OSPF, OSPFv3 and IS-IS to set overload when the router has fewer than the full set of SFMs functioning, which reduces forwarding capacity. Setting overload enables a router to still participate in exchanging routing information, but routes all traffic away from it.

The **no** form of this command configures the router to not set overload if an SFM fails.

Default

no single-sfm-overload

Parameters

holdoff-time

Specifies the delay between detecting SFM failures and setting overload.

Values 1 to 600 seconds

Default 0 seconds

Platforms

7705 SAR Gen 2

28.8 site

site

Syntax

site *name* [**create**]

no site *name*

Context

[Tree] (config>service>vpls site)

Full Context

configure service vpls site

Description

This command configures a VPLS site.

The no form of this command removes the name from the configuration.

Parameters

name

Specifies a site name up to 32 characters in length.

create

This keyword is mandatory while creating a VPLS site.

Platforms

7705 SAR Gen 2

site

Syntax

site *name* [**create**]

no site name

Context

[Tree] (config>service>epipe site)

Full Context

configure service epipe site

Description

This command configures a Epipe site.

The **no** form of this command removes the name from the configuration.

Parameters

name

Specifies a site name up to 32 characters in length.

create

This keyword is mandatory while creating a Epipe service.

Platforms

7705 SAR Gen 2

28.9 site-activation-timer

site-activation-timer

Syntax

site-activation-timer *seconds*

no site-activation-timer

Context

[Tree] (config>redundancy>bgp-multi-homing site-activation-timer)

Full Context

configure redundancy bgp-multi-homing site-activation-timer

Description

This command defines the amount of time the service manager will keep the local sites in standby status, waiting for BGP updates from remote PEs before running the DF election algorithm to decide whether the site should be unblocked. The timer is started when one of the following event occurs only if the site is operationally up:

- Manual site activation using "no shutdown” at site-id level or at member object(s) level (for example, SAP(s) or PW(s)
- Site activation after a failure

The **no** form of this command sets the value to 2.

Default

no site-activation-timer

Parameters

seconds

Specifies the timer, in seconds.

Values 1 to 100

Platforms

7705 SAR Gen 2

site-activation-timer

Syntax

site-activation-timer *seconds*

no site-activation-timer

Context

[Tree] (config>redundancy>bgp-multi-homing site-activation-timer)

Full Context

configure redundancy bgp-multi-homing site-activation-timer

Description

This command defines the amount of time the service manager will keep the local sites in standby status, waiting for BGP updates from remote PEs before running the DF election algorithm to decide whether the site should be unblocked. The timer is started when one of the following events occurs if the site is operationally up:

- Manual site activation using the **no shutdown** command at site-id level or at member object(s) level (SAP(s) or PW(s))
- Site activation after a failure

Default

no site-activation-timer

Parameters

seconds

Specifies the standby status in seconds.

Values 0 to 100

Default 2

Platforms

7705 SAR Gen 2

site-activation-timer

Syntax

site-activation-timer *seconds*

no site-activation-timer

Context

[Tree] (config>service>vpls>site site-activation-timer)

Full Context

configure service vpls site site-activation-timer

Description

This command configures the time-period the system keeps the local sites in standby status, waiting for BGP updates from remote PEs before running the DF (designated-forwarder) election algorithm to decide whether the site should be unblocked. This timer if terminated if an update is received for which the remote PE has transitioned from DF to non-DF.

The no form of this command removes the value from the configuration.

Default

site-activation-timer 2

Parameters

seconds

Specifies the site activation timer in seconds.

Values 0 to 100

Platforms

7705 SAR Gen 2

site-activation-timer

Syntax

site-activation-timer *seconds*

no site-activation-timer

Context

[Tree] (config>service>epipe>site site-activation-timer)

Full Context

configure service epipe site site-activation-timer

Description

This command configures the time-period the system keeps the local sites in standby status, waiting for BGP updates from remote PEs before running the DF (designated-forwarder) election algorithm to decide whether the site should be unblocked. This timer is terminated if an update is received for which the remote PE has transitioned from DF to non-DF.

The **no** form of this command removes the value from the configuration.

Default

site-activation-timer 2

Parameters

seconds

Specifies the site activation timer in seconds.

Values 0 to 100

Platforms

7705 SAR Gen 2

28.10 site-id

site-id

Syntax

site-id *value*

no site-id

Context

[Tree] (config>service>vpls>site site-id)

Full Context

configure service vpls site site-id

Description

This command configures the identifier for the site in this service.

Parameters

value

Specifies the site identifier.

Values 1 to 65535

Platforms

7705 SAR Gen 2

site-id

Syntax

site-id *value*

no site-id

Context

[Tree] (config>service>epipe>site site-id)

Full Context

configure service epipe site site-id

Description

This command configures the identifier for the site in this service. It must match between services but it is local to the service.

Parameters

value

Specifies the site identifier.

Values 1 to 65535

Platforms

7705 SAR Gen 2

28.11 site-min-down-timer

site-min-down-timer

Syntax

site-min-down-timer *seconds*

no site-min-down-timer

Context

[Tree] (config>redundancy>bgp-multi-homing site-min-down-timer)

Full Context

configure redundancy bgp-multi-homing site-min-down-timer

Description

This command configures the BGP multi-homing site minimum down time. When this value is set and the site goes operationally down, it remains operationally down for at least the length of time configured by this timer, regardless of whether other state changes might cause the site to go operationally up. This timer is restarted every time the site transitions from operationally up to down.

This timer is optimized in the following circumstances:

- If the site goes down on the DF but there are no BGP multi-homing peers with the same site in an up state, this timer is not used.
- If the site goes down on the DF but there are no active BGP multi-homing peers, this timer is not used.
- If this timer is active and a BGP multihoming update is received from the DF indicating its site is down, this timer is immediately terminated and the BGP multihoming algorithm is triggered to determine whether this PE should become the DF.

The **no** form of this command removes the value from the configuration.

Default

no site-min-down-timer

Parameters

seconds

Specifies the time, in seconds, that a BGP multi-homing site remains operationally down after a transition from up to down.

Values 1 to 100

Platforms

7705 SAR Gen 2

site-min-down-timer

Syntax

site-min-down-timer *min-down-time*

no site-min-down-timer

Context

[Tree] (config>service>vpls>site site-min-down-timer)

Full Context

configure service vpls site site-min-down-timer

Description

This command configures the BGP multi-homing site minimum down time. When set to a non-zero value, if the site goes operationally down it will remain operationally down for at least the length of time configured for the **site-min-down-timer**, regardless of whether other state changes would have caused it to go

operationally up. This timer is restarted every time that the site transitions from up to down. Setting this parameter to zero allows the minimum down timer to be disabled for this service.

The above operation is optimized in the following circumstances:

- If the site goes down on the designated forwarder but there are no BGP multi-homing peers with the same site in an operationally up state, then the **site-min-down-timer** is not started and is not used.
- If the site goes down on the designated forwarder but there are no active BGP multi-homing peers, then the **site-min-down-timer** is not started and is not used.
- If the s**ite-min-down-timer** is active and a BGP multi-homing update is received from the designated forwarder indicating its site has gone down, the **site-min-down-timer** is immediately terminated and this PE becomes the designated forwarder if the BGP multi-homing algorithm determines it should be the designated forwarder.

The **no** form of this command reverts to the default value.

Default

Taken from the value of **site-min-down-timer** configured for Multi-Chassis BGP multi-homing under the **config>redundancy>bgp-multi-homing** context.

Parameters

min-down-time

Specifies the time, in seconds, that a BGP multi-homing site remains operationally down after a transition from up to down.

Values 0 to 100 seconds

Platforms

7705 SAR Gen 2

site-min-down-timer

Syntax

site-min-down-timer *min-down-time*

no site-min-down-timer

Context

[Tree] (config>service>epipe>site site-min-down-timer)

Full Context

configure service epipe site site-min-down-timer

Description

This command configures the BGP multi-homing site minimum down time. When set to a non-zero value, if the site goes operationally down it will remain operationally down for at least the length of time configured for the **site-min-down-timer**, regardless of whether other state changes would have caused it to go

operationally up. This timer is restarted every time that the site transitions from up to down. Setting this parameter to zero allows the minimum down timer to be disabled for this service.

The preceding operation is optimized in the following circumstances:

- If the site goes down on the designated forwarder but there are no BGP multi-homing peers with the same site in an operationally up state, then the **site-min-down-timer** is not started and is not used.
- If the site goes down on the designated forwarder but there are no active BGP multi-homing peers, then the s**ite-min-down-timer** is not started and is not used.
- If the **site-min-down-timer** is active and a BGP multi-homing update is received from the designated forwarder indicating its site has gone down, the **site-min-down-timer** is immediately terminated and this PE becomes the designated forwarder if the BGP multi-homing algorithm determines it should be the designated forwarder.

The **no** form of this command reverts to default value.

Default

Taken from the value of **site-min-down-timer** configured for Multi-Chassis BGP multi-homing under the **config>redundancy>bgp-multi-homing** context.

Parameters

min-down-time

Specifies the time, in seconds, that a BGP multi-homing site remains operationally down after a transition from up to down.

Values 0 to 100

Platforms

7705 SAR Gen 2

28.12 site-preference

site-preference

Syntax

site-preference *preference-value*

no site-preference

Context

[Tree] (config>service>epipe>site site-preference)

Full Context

configure service epipe site site-preference

Description

This command defines the value to advertise in the VPLS preference field of the BGP VPWS and BGP Multi-homing NLRI extended community. This value can be changed without having to shutdown the site itself. The site-preference is only applicable to VPWS services.

When not configured, the default is zero, indicating that the VPLS preference is not in use.

Default

no site-preference, value=0

Parameters

preference-value

Specifies the preference value to advertise in the NLRI L2 extended community for this site.

Values 1 to 65535

primary

Sets the site-preference to 65535.

backup

Sets the site-preference to 1.

Platforms

7705 SAR Gen 2

28.13 size

size

Syntax

size *octets*

no size

Context

[Tree] (config>saa>test>type-multi-line>lsp-ping size)

[Tree] (config>saa>test>type-multi-line>lsp-ping>sr-policy size)

Full Context

configure saa test type-multi-line lsp-ping size

configure saa test type-multi-line lsp-ping sr-policy size

Description

This command configures the MPLS echo request packet size.

The **no** form of this command reverts to the default value.

Default

size 1

Parameters

octets

Specifies the size in octets. The request payload is padded with zeros to the specified size.

Values 1 to 9786

Default 1

Platforms

7705 SAR Gen 2

size

Syntax

size *number*

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>history size)

Full Context

configure system management-interface cli md-cli environment history size

Description

This command configures the maximum size of the command history.

Parameters

number

Specifies the maximum size of the command history. A value of 0 disables the command history.

Values 0 to 1000

Platforms

7705 SAR Gen 2

28.14 sleep

sleep

Syntax

sleep [*seconds*]

Context

[Tree] (sleep)

Full Context

sleep

Description

This command causes the console session to pause operation (sleep) for 1 second (default) or for the specified number of seconds.

Default

sleep 1

Parameters

seconds

Specifies the number of seconds for the console session to sleep, expressed as a decimal integer.

Values 1 to 100

Default 1

Platforms

7705 SAR Gen 2

28.15 slice-size

slice-size

Syntax

slice-size *slice-size*

no slice-size

Context

[Tree] (config>mirror>mirror-dest slice-size)

Full Context

configure mirror mirror-dest slice-size

Description

This command enables mirrored frame truncation and specifies the maximum size, in bytes, of a mirrored frame that can be transmitted to the mirror destination.

This command enables mirroring larger frames than the destination packet decode equipment can handle. It also allows conservation of mirroring resources by limiting the size of the packet stream through the router and the core network.

When defined, the mirror **slice-size** creates a threshold that truncates a mirrored frame to a specific size. For example, if the value of 256 bytes is defined, a frame larger than 256 bytes will only have the first 256 bytes transmitted to the mirror destination. The original frame is not affected by the truncation. The mirrored frame size may increase if encapsulation information is added during transmission through the network core or out the mirror destination SAP to the packet/protocol decode equipment.

The actual capability of the router to transmit a sliced or non-sliced frame is also dictated by the mirror destination SDP **path-mtu** or the mirror destination SAP physical MTU. Packets that require a larger MTU than the mirroring destination supports are discarded if the defined **slice-size** does not truncate the packet to an acceptable size.

Notes:

- When configuring IP mirroring, packet slice is rejected as an incorrect option as it will cause IP packets to be rejected by the next hop with an IP header verification error.
- Slice-size is not supported by CEM encap-types or IP-mirroring.

The **no** form of this command disables mirrored packet truncation.

Parameters

slice-size

Specifies the number of bytes to which mirrored frames are truncated, expressed as a decimal integer.

Values 128 to 9216

Platforms

7705 SAR Gen 2

28.16 snap-oui

snap-oui

Syntax

snap-oui {**zero** | **non-zero**}

no snap-oui

Context

[Tree] (config>qos>sap-ingress>mac-criteria>entry>match snap-oui)

Full Context

configure qos sap-ingress mac-criteria entry match snap-oui

Description

Configures an IEEE 802.3 LLC SNAP Ethernet frame OUI zero or non-zero value to be used as a service ingress QoS policy match criterion.

The **no** form of this command removes the criterion from the match criteria.

Default

no snap-oui

Parameters

zero

Specifies to match packets with the 3-byte OUI field in the SNAP-ID set to zero.

non-zero

Specifies to match packets with the 3-byte OUI field in the SNAP-ID not set to zero.

Platforms

7705 SAR Gen 2

snap-oui

Syntax

snap-oui {**zero** | **non-zero**}

no snap-oui

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry>match snap-oui)

Full Context

configure system security management-access-filter mac-filter entry match snap-oui

Description

This command configures an IEEE 802.3 LLC SNAP Ethernet Frame OUI zero or non-zero value to be used as a MAC filter match criterion.

The **no** form of this command removes the criterion from the match criteria.

Default

no snap-oui

Parameters

zero

Specifies to match packets with the three-byte OUI field in the SNAP-ID set to zero.

non-zero

Specifies to match packets with the three-byte OUI field in the SNAP-ID not set to zero.

Platforms

7705 SAR Gen 2

28.17 snap-pid

snap-pid

Syntax

snap-pid *snap-pid*

no snap-pid

Context

[Tree] (config>qos>sap-ingress>mac-criteria>entry>match snap-pid)

Full Context

configure qos sap-ingress mac-criteria entry match snap-pid

Description

Configures an IEEE 802.3 LLC SNAP Ethernet frame PID value to be used as a service ingress QoS policy match criterion.

This is a 2-byte protocol id that is part of the IEEE 802.3 LLC SNAP Ethernet Frame that follows the 3-byte OUI field.

The snap-pid field, etype field, ssap, and dsap fields are mutually exclusive and cannot be part of the same match criteria.

The **snap-pid** match criteria is independent of the OUI field within the SNAP header. Two packets with different 3-byte OUI fields, but the same PID field, will both match the same policy entry based on a snap-pid match criteria.

The **no** form of this command removes the snap-pid value as the match criteria.

Default

no snap-pid

Parameters

snap-pid

The 2-byte snap-pid value to be used as a match criterion in hexadecimal.

Values 0x0000 to 0xFFFF

Platforms

7705 SAR Gen 2

snap-pid

Syntax

snap-pid *snap-pid*

no snap-pid

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry>match snap-pid)

Full Context

configure system security management-access-filter mac-filter entry match snap-pid

Description

This command configures an IEEE 802.3 LLC SNAP Ethernet Frame PID value to be used as a MAC filter match criterion.

This is a two-byte protocol id that is part of the IEEE 802.3 LLC SNAP Ethernet Frame that follows the three-byte OUI field.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and may not be part of the same match criteria. Refer to the *7705 SAR Gen 2 Router Configuration Guide* for information about MAC Match Criteria Exclusivity Rules fields that are exclusive based on the frame format.

Note:
The snap-pid match criterion is independent of the OUI field within the SNAP header. Two packets with different three-byte OUI fields but the same PID field will both match the same filter entry based on a snap-pid match criteria.

The **no** form of this command removes the snap-pid value as the match criteria.

Default
no snap-pid

Parameters
pid-value

Specifies the two-byte snap-pid value to be used as a match criterion in hexadecimal.

Values 0x0000 to 0xFFFF

Platforms
7705 SAR Gen 2

28.18 snmp

snmp

Syntax
snmp

Context
[Tree] (config>service>vprn snmp)

Full Context
configure service vprn snmp

Description
Commands in this context configure SNMP parameters for this VPRN.

Platforms
7705 SAR Gen 2

snmp

Syntax
snmp

Context
[Tree] (config>system>security>user snmp)

Full Context

configure system security user snmp

Description

This command creates the context to configure SNMP group membership for a specific user and defines encryption and authentication parameters.

All SNMPv3 users must be configured with the commands available in this CLI node.

The OS always uses the configured SNMPv3 user name as the security user name.

Platforms

7705 SAR Gen 2

snmp

Syntax

snmp

Context

[Tree] (config>system snmp)

[Tree] (config>system>security snmp)

Full Context

configure system snmp

configure system security snmp

Description

This command creates the context to configure SNMPv1, SNMPv2, and SNMPv3 parameters**.**

Platforms

7705 SAR Gen 2

28.19 snmp-trap-group

snmp-trap-group

Syntax

snmp-trap-group *log-id* | *log-name* [**name** *log-name*]

no snmp-trap-group *log-id* | *log-name*

Context

[Tree] (config>service>vprn>log snmp-trap-group)

Full Context

configure service vprn log snmp-trap-group

Description

This command creates the context to configure a group of SNMP trap receivers and their operational parameters for a specific *log-id*.

A group specifies the types of SNMP traps and specifies the log ID that will receive the group of SNMP traps. The user must configure a trap group before SNMP traps can be sent.

To suppress the generation of all alarms and traps, see the **event-control** command. To suppress alarms and traps that are sent to this log-id, see the **filter** command. After alarms and traps are generated, they can be directed to one or more SNMP trap groups. Log events that can be forwarded as SNMP traps are always defined on the main event source.

The **no** form of this command deletes the SNMP trap group.

Parameters

log-id* | *log-name

Specifies the log ID or name (up to 32 characters).

Values *log-id*: 1 to 100

name *log-name*

Specifies an optional log name of a log configured in the **log-id** context, up to 32 characters, that can be used to refer to the log after it is created. Alarms and traps cannot be sent to the trap receivers until a valid *log-id* exists.

Platforms

7705 SAR Gen 2

snmp-trap-group

Syntax

snmp-trap-group *log-id* | *log-name* [**name** *log-name*]

no snmp-trap-group *log-id* | *log-name*

Context

[Tree] (config>log snmp-trap-group)

Full Context

configure log snmp-trap-group

Description

This command creates the context to configure a group of SNMP trap receivers and their operational parameters for a specified *log-id*.

A group specifies the types of SNMP traps and the log ID which that will receive the SNMP trap group. The user must configure a trap to send SNMP traps.

To suppress the generation of all alarms and traps, see the **event-control** command. To suppress alarms and traps that are sent to this log ID, see the filter command. When alarms and traps are generated, they can be directed to one or more SNMP trap groups. Log events that can be forwarded as SNMP traps are always defined at the main event source.

The **no** form of this command deletes the SNMP trap group.

Parameters

log-id* | *log-name

Specifies the log ID or log name (up to 32 characters).

Values *log-id*: 1 to 100

name *log-name*

Specifies an optional log name of a log configured in the **log-id** context, up to 32 characters, that can be used to refer to the log after it is created. Alarms and traps cannot be sent to the trap receivers until a valid *log-id* exists.

Platforms

7705 SAR Gen 2

28.20 snoop

snoop

Syntax

[**no**] **snoop**

Context

[Tree] (config>service>vpls>sap>dhcp snoop)

[Tree] (config>service>vpls>spoke-sdp>dhcp snoop)

[Tree] (config>service>vpls>mesh-sdp>dhcp snoop)

Full Context

configure service vpls sap dhcp snoop

configure service vpls spoke-sdp dhcp snoop

configure service vpls mesh-sdp dhcp snoop

Description

This command enables snooping of DHCP or DHCP6 messages on the SAP or SDP. Enabling DHCP or DHCP6 snooping on interfaces (SAPs and SDP bindings) is required where DHCP or DHCP6 messages important to lease state table population are received, or where Option 82 information is to be inserted. This includes interfaces that are in the path to receive messages from either DHCP or DHCP6 servers or from subscribers.

The **no** form of this command disables DHCP or DHCP6 snooping on the specified SAP or SDP binding.

Default

no snoop

Platforms

7705 SAR Gen 2

28.21 sntp

sntp

Syntax

[**no**] **sntp**

Context

[Tree] (config>system>time sntp)

Full Context

configure system time sntp

Description

This command creates the context to edit the Simple Network Time Protocol (SNTP).

SNTP can be configured in either broadcast or unicast client mode. SNTP is a compact, client-only version of the NTP. SNTP can only receive the time from SNTP/NTP servers. It cannot be used to provide time services to other systems.

The system clock is automatically adjusted at system initialization time or when the protocol first starts up.

When the time differential between the SNTP/NTP server and the system is more than 2.5 seconds, the time on the system is gradually adjusted.

SNTP is created in an administratively enabled state (**no shutdown**).

The **no** form of the command removes the SNTP instance and configuration. SNTP does not need to be administratively disabled when removing the SNTP instance and configuration.

Default

sntp

Platforms

7705 SAR Gen 2

28.22 socket

socket

Syntax

socket [**neighbor** *ip-address* | **group** *name*]

no socket

Context

[Tree] (debug>router>bgp socket)

Full Context

debug router bgp socket

Description

This command logs all TCP socket events to the debug log.

The **no** form of this command disables debugging.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x [-interface] (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H
- d: [0 to 255]D
- interface: up to 32 characters for link local addresses

group *name*

Debugs only events affecting the specified peer group name, up to 64 characters, and associated neighbors.

Platforms

7705 SAR Gen 2

28.23 solicited-release

solicited-release

Syntax

[**no**] **solicited-release**

Context

[Tree] (config>router>dhcp6>server>lease-hold-time-for solicited-release)

[Tree] (config>router>dhcp>server>lease-hold-time-for solicited-release)

[Tree] (config>service>vprn>dhcp6>server>lease-hold-time-for solicited-release)

[Tree] (config>service>vprn>dhcp>server>lease-hold-time-for solicited-release)

Full Context

configure router dhcp6 local-dhcp-server lease-hold-time-for solicited-release

configure router dhcp local-dhcp-server lease-hold-time-for solicited-release

configure service vprn dhcp6 local-dhcp-server lease-hold-time-for solicited-release

configure service vprn dhcp local-dhcp-server lease-hold-time-for solicited-release

Description

This command enables the server to hold up a lease even in case of solicited release; for example, when the server receives a normal DHCP release message.

The **no** form of this command disables the ability of the server to hold up a lease when a solicited release is received.

Platforms

7705 SAR Gen 2

28.24 source

source

Syntax

[**no**] **source** *ip-address*

[**no**] **source** *src-ipv6-address*

Context

[Tree] (config>service>vpls>mesh-sdp>mld-snooping>static>group source)

[Tree] (config>service>vpls>spoke-sdp>mld-snooping>static>group source)

[Tree] (config>service>vpls>sap>igmp-snooping>static>group source)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping>static>group source)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping>static>group source)

[Tree] (config>service>vpls>sap>mld-snooping>static>group source)

Full Context

configure service vpls mesh-sdp mld-snooping static group source

configure service vpls spoke-sdp mld-snooping static group source

configure service vpls sap igmp-snooping static group source

configure service vpls spoke-sdp igmp-snooping static group source

configure service vpls mesh-sdp igmp-snooping static group source

configure service vpls sap mld-snooping static group source

Description

This command specifies a IPv4 or IPv6 unicast address that sends data on an interface. This enables a multicast receiver host to signal a router the group to receive multicast traffic from, and from the sources that the traffic is expected.

The **source** command is mutually exclusive with the specification of individual sources for the same group.

The **source** command in combination with the group is used to create a specific (S,G) static group entry.

Static (s,g) entries cannot be entered when a starg is already created.

Use the **no** form of this command to remove the source from the configuration.

Parameters

ip-address

Specifies the IPv4 unicast address

src-ipv6-address

Specifies the IPv6 unicast address.

Platforms

7705 SAR Gen 2

source

Syntax

source *ip-address*

no source

Context

[Tree] (config>service>ies>if>sap>ip-tunnel source)

[Tree] (config>service>vprn>if>sap>ip-tunnel source)

Full Context

configure service ies interface sap ip-tunnel source

configure service vprn interface sap ip-tunnel source

Description

This command configures the source IPv4 or IPv6 address to use for an IP tunnel. This configuration applies to the outer IP header of the encapsulated packets. The IPv4 or IPv6 address must belong to the one of the IP subnets associated with the public SAP interface of the tunnel-group. The **source** address, **remote-ip** address and **backup-remote-ip** address of a tunnel must all belong to the same address family (IPv4 or IPv6). When the source address contains an IPv6 address it must be a global unicast address.

The **no** form of this command deletes the source address from the tunnel configuration. The tunnel must be administratively shutdown before issuing the **no source** command.

Default

no source

Parameters

ip-address

Specifies an IPv4 address or an IPv6 address.

Values

ipv4-address	a.b.c.d	
ipv6-address		x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
	x	[0..FFFF]H
	d	[0..255]D

Platforms

7705 SAR Gen 2

source

Syntax

[**no**] **source** *ip-address*

Context

[Tree] (config>service>vprn>igmp>ssm-translate>grp-range source)

Full Context

configure service vprn igmp ssm-translate grp-range source

Description

This command specifies the source IP address for the group range. Whenever a (*,G) report is received in the range specified by **grp-range** *start* and *end* parameters, it is translated to an (S,G) report with the value of this object as the source address.

Parameters

ip-address

Specifies the IP address that will be sending data.

Platforms

7705 SAR Gen 2

source

Syntax

source *ip-address*

Context

[Tree] (config>service>vprn>igmp>if>static>group source)

Full Context

configure service vprn igmp interface static group source

Description

This command specifies an IPv4 unicast address that sends data on an interface. This enables a multicast receiver host to signal a router the group is to receive multicast traffic from, and from the sources that the traffic is expected.

The **source** command is mutually exclusive with the specification of individual sources for the same group.

The **source** command in combination with the group is used to create a specific (S,G) static group entry.

Use the **no** form of this command to remove the source from the configuration.

Parameters

ip-address

Specifies the IPv4 unicast address.

Platforms

7705 SAR Gen 2

source

Syntax

[**no**] **source** *src-ipv6-address*

Context

[Tree] (config>service>vprn>mld>if>static>group source)

Full Context

configure service vprn mld interface static group source

Description

This command specifies an IPv6 unicast address that sends data on an interface. This enables a multicast receiver host to signal a router the group to receive multicast traffic from, and from the sources that the traffic is expected.

The **source** command is mutually exclusive with the specification of individual sources for the same group.

The **source** command, in combination with the group, is used to create a specific (S,G) static group entry.

The **no** form of this command removes the source from the configuration.

Parameters

src-ipv6-address

Specifies the IPv6 unicast address.

Platforms

7705 SAR Gen 2

source

Syntax

[**no**] **source** *ip-address*

Context

[Tree] (config>service>vprn>mld>ssm-translate>grp-range source)

Full Context

configure service vprn mld ssm-translate grp-range source

Description

This command specifies the source IP address for the group range. Whenever a (*,G) report is received in the range specified by **grp-range** *start* and *end* parameters, it is translated to an (S,G) report with the value of this object as the source address.

Parameters

ip-address

Specifies the IP address that will be sending data.

Platforms

7705 SAR Gen 2

source

Syntax

[**no**] **source** *ip-address*

Context

[Tree] (config>router>igmp>if>ssm-translate>grp-range source)

[Tree] (config>router>igmp>ssm-translate>grp-range source)

Full Context

configure router igmp interface ssm-translate grp-range source

configure router igmp ssm-translate grp-range source

Description

This command specifies the source IP address for the group range. Whenever a (*,G) report is received in the range specified by **grp-range** *start* and *end* parameters, it is translated to an (S,G) report with the value of this object as the source address.

Parameters

ip-address

Specifies the IP address that will be sending data.

Platforms

7705 SAR Gen 2

source

Syntax

[**no**] **source** *ip-address*

Context

[Tree] (config>router>igmp>if>static>group source)

Full Context

configure router igmp interface static group source

Description

This command specifies a IPv4 unicast address that sends data on an interface. This enables a multicast receiver host to signal a router the group to receive multicast traffic from, and from the source(s) that the traffic is expected.

The **source** command is mutually exclusive with the specification of individual sources for the same group.

The **source** command in combination with the group is used to create a specific (S,G) static group entry.

The **no** form of the command removes the source from the configuration.

Parameters

ip-address

Specifies the IPv4 unicast address.

Platforms

7705 SAR Gen 2

source

Syntax

[**no**] **source** *src-ipv6-address*

Context

[Tree] (config>router>mld>if>static>group source)

Full Context

configure router mld interface static group source

Description

This command specifies an IPv6 unicast address that sends data on an interface. This enables a multicast receiver host to signal a router the group to receive multicast traffic from, and from the source(s) that the traffic is expected.

The **source** command is mutually exclusive with the specification of individual sources for the same group.

The **source** command, in combination with the group, is used to create a specific (S,G) static group entry.

The **no** form of this command removes the source from the configuration.

Parameters

src-ipv6-address

Specifies the IPv6 unicast address.

Platforms

7705 SAR Gen 2

source

Syntax

[**no**] **source** *ipv6-address*

Context

[Tree] (config>router>mld>if>ssm-translate>grp-range source)

[Tree] (config>router>mld>ssm-translate>grp-range source)

Full Context

configure router mld interface ssm-translate grp-range source

configure router mld ssm-translate grp-range source

Description

This command specifies the source IP address for the group range. Whenever a (*,G) report is received in the range specified by **grp-range** *start* and *end* parameters, it is translated to an (S,G) report with the value of this object as the source address.

The **no** form of this command removes the IPv6 address form the group range configuration.

Parameters

ipv6-address

Specifies the IPv6 address that will be sending data.

Platforms

7705 SAR Gen 2

source

Syntax

source *ip-address*

no source

Context

[Tree] (config>oam-pm>session>ip source)

Full Context

configure oam-pm session ip source

Description

This command defines the source IP address that the session controller (launch point) uses for the test. The source address must be a local resident IP address in the context; otherwise, the response packets

are processed by the TWAMP Light application. Only source addresses configured as part of TWAMP tests can process the reflected TWAMP packets from the session reflector.

The **no** form of this command removes the source address parameters.

Parameters

source

Indicates the launch point.

ip-address

Specifies the source IP address that the session controller (launch point) uses for the test.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

Platforms

7705 SAR Gen 2

28.25 source-address

source-address

Syntax

source-address *ipv6-address*

no source-address

Context

[Tree] (config>service>vprn>if>ipv6>dhcp6-relay source-address)

[Tree] (config>service>ies>if>ipv6>dhcp6-relay source-address)

Full Context

configure service vprn interface ipv6 dhcp6-relay source-address

configure service ies interface ipv6 dhcp6-relay source-address

Description

This command configures the source IPv6 address of the DHCPv6 relay messages.

The **no** form of this command reverts to the default.

Parameters

ipv6-address

Specifies the source IPv6 address of the DHCPv6 relay messages.

Values

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x: [0 to FFFF]H
d: [0 to 255]D

Platforms

7705 SAR Gen 2

source-address

Syntax

source-address *ip-address*

no source-address

Context

[Tree] (config>system>management-interface>remote-management source-address)

Full Context

configure system management-interface remote-management source-address

Description

This command configures the address local to this device that NISH uses to connect to this node.

If this command is also configured for a specific manager in the **config**>**system**> **management-interface**>**remote-management**>**manager** context, that configuration takes precedence.

The **no** form of this command causes the system to select the source address based on the selected routing instance of the manager.

Parameters

ip-address

Specifies the IP address that NISH managers use to connect to the node.

Values

ipv4-address: a.b.c.d
ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x - [0 to FFFF]H

d - [0 to 255]D

Platforms

7705 SAR Gen 2

source-address

Syntax

source-address *ip-address*

no source-address

Context

[Tree] (config>system>management-interface>remote-management>manager source-address)

Full Context

configure system management-interface remote-management manager source-address

Description

This command configures the address local to this device that this NISH manager uses to connect to this node.

This command takes precedence over the command configured in the global context (**config**>**system**>**management-interface**>**remote-management**).

The **no** form of this command causes the source address to be inherited from the global context (**config**>**system**>**management-interface**>**remote-management**).

Parameters

ip-address

Specifies the IP address that NISH managers use to connect to the node.

Values

ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

Platforms

7705 SAR Gen 2

source-address

Syntax

source-address *ip-address*

no source-address

Context

[Tree] (config>redundancy>multi-chassis>peer source-address)

Full Context

configure redundancy multi-chassis peer source-address

Description

This command specifies the source address used to communicate with the multi-chassis peer.

The **no** form of this command reverts to the default.

Parameters

ip-address

Specifies the source address used to communicate with the multi-chassis peer.

Values

ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

Platforms

7705 SAR Gen 2

source-address

Syntax

source-address *ip-address*

no source-address

Context

[Tree] (config>aaa>radius-srv-plcy>servers source-address)

Full Context

configure aaa radius-server-policy servers source-address

Description

This command configures the source address of the RADIUS packet. The system IP address must be configured in order for the RADIUS client to work. See "Configuring a System Interface" in the *7705 SAR Gen 2 Router Configuration Guide*.

Note:
The system IP address must only be configured if the source-address is not specified. When the no source-address command is executed, the source address is determined at the moment the request is sent. This address is also used in the *nas-ip-address* attribute: over there it is set to the system IP address if no source-address was given.

The **no** form of this command reverts to the default value.

Parameters

ip-address

Specifies the source address of RADIUS packet.

Platforms

7705 SAR Gen 2

source-address

Syntax

source-address

Context

[Tree] (config>service>vprn source-address)

Full Context

configure service vprn source-address

Description

Commands in this context specify the source address and application that should be used in all unsolicited packets.

Platforms

7705 SAR Gen 2

source-address

Syntax

source-address [*ip-address*]

no source-address

Context

[Tree] (config>filter>redirect-policy>dest>ping-test source-address)

Full Context

configure filter redirect-policy destination ping-test source-address

Description

This command configures the source address to use in the IP packet of the ping test for this destination.

Default

no source-address

Parameters

ip-address

The source address of the IP packet. This can be IPv4 only for an IPv4 destination and IPv6 only for an IPv6 destination.

Values

ipv4-address:	a.b.c.d.
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

Platforms

7705 SAR Gen 2

source-address

Syntax

source-address

Context

[Tree] (config>system>security source-address)

Full Context

configure system security source-address

Description

This command configures the IP source address that is used in all unsolicited packets sent by the application.

The configured source address applies only to packets transmitted in-band (for example, a network port on an IOM). Packets transmitted out-of-band on the management interface on the CPM Ethernet port use the address of the CPM Ethernet port as the IP source address in the packet.

When a source address is specified for the **ptp** application, the port-based 1588 hardware timestamping assist function will be applied to PTP packets matching the IPv4 address of the router interface used to ingress the 7705 SAR Gen 2 or IP address specified in this command. If the IP address is removed, then the port-based 1588 hardware timestamping assist function will only be applied to PTP packets matching the IPv4 address of the router interface.

Platforms

7705 SAR Gen 2

source-address

Syntax

source-address *ip-address*

source-address prefix-list *prefix-list-name*

no source-address

Context

[Tree] (config>router>policy-options>policy-statement>entry>from source-address)

Full Context

configure router policy-options policy-statement entry from source-address

Description

This command specifies the source address that is embedded in the join or prune packet as a filter criterion.

The **no** form of this command removes the criterion from the configuration.

This command specifies a multicast data source address as a match criterion for this entry.

Default

no source-address

Parameters

ip-address

Specifies the IP prefix for the IP match criterion in dotted decimal notation.

Values ipv4-address:

- a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x:x
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

prefix-list-name

The prefix list name. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

source-address

Syntax

source-address

Context

[Tree] (config>router>pim source-address)

[Tree] (config>service>vprn>pim source-address)

Full Context

configure router pim source-address

configure service vprn pim source-address

Description

Commands in this context configure the source IP address for PIM messages.

Platforms

7705 SAR Gen 2

28.26 source-port

source-port

Syntax

source-port *port*

source-port grpc

no source-port

Context

[Tree] (config>system>management-interface>remote-management source-port)

Full Context

configure system management-interface remote-management source-port

Description

This command configures the TCP port local to this device that NISH uses to send packets to this node.

If this command is also configured for a specific manager in the **config**>**system**> **management-interface**>**remote-management**>**manager** context, that configuration takes precedence.

The **no** form of this command causes the system to select the default gRPC port, 57400.

Default

source-port grpc

Parameters

port

Specifies the TCP source port.

Values 1 to 65535

grpc

Keyword that specifies the default gRPC protocol port as the source port.

Platforms

7705 SAR Gen 2

source-port

Syntax

source-port *port*

source-port grpc

no source-port

Context

[Tree] (config>system>management-interface>remote-management>manager source-port)

Full Context

configure system management-interface remote-management manager source-port

Description

This command configures the TCP port local to this device that this NISH manager uses to send packets to this node.

This command takes precedence over the same command configured in the global context (**config**>**system**>**management-interface**>**remote-management**).

The **no** form of this command causes the source port to be inherited from the global context (**config**>**system**>**management-interface**>**remote-management**).

Parameters

port

Specifies the TCP source port.

Values 1 to 65535

Default 57400

grpc

Keyword that specifies the default gRPC protocol port as the source port.

Platforms

7705 SAR Gen 2

28.27 source-udp-port

source-udp-port

Syntax

source-udp-port *udp-port-number*

no source-udp-port

Context

[Tree] (config>oam-pm>session>ip source-udp-port)

Full Context

configure oam-pm session ip source-udp-port

Description

This command should only be used when the source UDP port for the session-sender twamp-test packet must be specified.

The **no** form of this command means the session-sender automatically assigns the source UDP port from the available dynamic (private) UDP range.

Parameters

udp-port-number

Specifies the UDP source port.

Values 64374 to 64383

Platforms

7705 SAR Gen 2

28.28 sp-reverse-route

sp-reverse-route

Syntax

sp-reverse-route [**ignore-default-route**]

no sp-reverse-route

Context

[Tree] (config>ipsec>tnl-temp sp-reverse-route)

Full Context

configure ipsec tunnel-template sp-reverse-route

Description

This command enables the system to automatically create a reverse route based on dynamic LAN-to-LAN tunnel's TSi in private service.

If **ignore-default-route** is specified, the system ignores any full range traffic selector when creating a reverse route. Otherwise, the system refuses to create a CHILD_SA if any full range traffic selector is included in TSi.

The **no** form of this command disables sp-reverse-route.

Default

no sp-reverse-route

Parameters

ignore-default-route

Specifies to ignore any full range traffic selector in TSi.

Platforms

7705 SAR Gen 2

28.29 space

space

Syntax

[**no**] **space**

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>command-completion space)

Full Context

configure system management-interface cli md-cli environment command-completion space

Description

This command enables completion on the space character.

The **no** form of this command reverts to the default value.

Default

space

Platforms

7705 SAR Gen 2

28.30 spe-address

spe-address

Syntax

spe-address *global-id:prefix*

no spe-address

Context

[Tree] (config>service>pw-routing spe-address)

Full Context

configure service pw-routing spe-address

Description

This command configures a single S-PE Address for the node to be used for dynamic MS-PWs. This value is used for the pseudowire switching point TLV used in LDP signaling, and is the value used by pseudowire status signaling to indicate the PE that originates a pseudowire status message. Configuration of this parameter is mandatory to enable dynamic MS-PW support on a node.

If the S-PE Address is not configured, spoke-sdps that use dynamic MS-PWs and pw-routing local-prefixes cannot be configured on a T-PE. Furthermore, the node will send a label release for any label mappings received for FEC129 AII type 2.

The S-PE Address cannot be changed unless the dynamic ms-pw configuration is removed. Furthermore, changing the S-PE Address will also result in all dynamic MS-PWs for which this node is an S-PE being released. It is recommended that the S-PE Address should be configured for the life of an MS-PW configuration after reboot of the router.

The **no** form of this command removes the configured S-PE Address.

Default

no spe-address

Parameters

global-id

Specifies a 4-octet value that is unique to the service provider. For example, the global ID can contain the 2-octet or 4-octet value of the provider's Autonomous System Number (ASN).

Values

<global-id:prefix>:	<global-id>:{<prefix>| <ipaddress>}
global-id	1 to 4294967295
prefix	1 to 4294967295

	ipaddress	a.b.c.d

Platforms

7705 SAR Gen 2

28.31 speed

speed

Syntax

speed {**10** | **100** | **1000** | **10000** | **25000** | **40000** | **50000** | **100000**}

Context

[Tree] (config>port>ethernet speed)

Full Context

configure port ethernet speed

Description

For ports that support multiple speeds, this command configures the port speed to be used. This applies to the following:

- fast Ethernet when autonegotiate is disabled
- 10/100/1000 Ethernet when autonegotiate is disabled
- 10/1G ports supporting 10G SFP+ or 1G SFP
- 40/100G ports supporting QSFP28s on non-connector-based MDAs

If the port is configured to autonegotiate this parameter is ignored. Speed cannot be configured for ports that are part of a Link Aggregation Group (LAG).

Default

dependent on port type

Parameters

10

Sets the link to 10 Mb/s speed.

100

Sets the link to 100 Mb/s speed.

1000

Sets the link to 1000 Mb/s speed.

10000

Sets the link to 10000 Mb/s speed.

25000

Sets the link to 25000 Mb/s speed.

40000

Sets the link to 40000 Mb/s speed.

50000

Sets the link to 50000 Mb/s speed.

100000

Sets the link to 100000 Mb/s speed.

Platforms

7705 SAR Gen 2

speed

Syntax

speed *speed*

Context

[Tree] (bof speed)

Full Context

bof speed

Description

This command configures the speed for the CPM management Ethernet port when autonegotiation is disabled in the running configuration and the Boot Option File (BOF).

If the port is configured to autonegotiate, this parameter is ignored.

Available speed options are dependent on the specific CPM variant in the system.

Default

speed 100

Parameters

speed

Sets the link speed, in Mb/s.

Values 10, 100, 1000

Platforms

7705 SAR Gen 2

28.32 spf

spf

Syntax

[**no**] **spf** [*level-number*] [*system-id*]

Context

[Tree] (debug>router>isis spf)

Full Context

debug router isis spf

Description

This command enables debugging for IS-IS SFP.

The **no** form of the command disables debugging.

Parameters

system-id

When specified, only the specified system-id is debugged. A 6-octet system identifier (xxxx.xxxx.xxxx).

level-number

Specifies the interface level (1, 2, or 1 and 2).

Platforms

7705 SAR Gen 2

spf

Syntax

spf [**type**] [*dest-addr*]

no spf

Context

[Tree] (debug>router>ospf3 spf)

[Tree] (debug>router>ospf spf)

Full Context

debug router ospf3 spf

debug router ospf spf

Description

This command enables debugging for OSPF SPF. Information regarding overall SPF start and stop times will be shown. To see detailed information regarding the SPF calculation of a given route, the route must be specified as an optional argument.

Parameters

type

Specifies the area to debug.

Values intra-area, inter-area, external

dest-addr

Specifies the destination IP address to debug.

Platforms

7705 SAR Gen 2

28.33 spf-wait

spf-wait

Syntax

spf-wait *spf-wait* [**spf-initial-wait** *initial-wait*] [**spf-second-wait** *second-wait*]

no spf-wait

Context

[Tree] (config>service>vprn>isis>timers spf-wait)

Full Context

configure service vprn isis timers spf-wait

Description

This command configures the maximum interval, in milliseconds, between two consecutive SPF calculations. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs are controlled with this command.

Subsequent SPF runs (if required) occur at exponentially increasing intervals of the **spf-second-wait** interval. For example, if the **spf-second-wait** interval is 1000, the next SPF will run after 2000 milliseconds, and the next SPF after 4000 milliseconds, and so on, until it reaches the **spf-wait** value. The SPF interval remains at the **spf-wait** value until no more SPF runs are scheduled in that interval. After a full interval without any SPF runs, the SPF interval drops back to the **spf-initial-wait** value.

Note:
The timer granularity is 100 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

The **no** form of this command reverts to the default value.

Default

spf-wait 10000 spf-initial-wait 1000 spf-second-wait 1000

Parameters

spf-wait

Specifies the maximum interval, in milliseconds, between two consecutive SPF calculations.

Values 10 to 120000

initial-wait

Specifies the initial SPF calculation delay, in milliseconds, after a topology change.

Values 10 to 100000

second-wait

Specifies the hold time, in milliseconds, between the first and second SPF calculation.

Values 10 to 100000

Platforms

7705 SAR Gen 2

spf-wait

Syntax

spf-wait *max-spf-wait* [**spf-initial-wait** *spf-initial-wait*] [**spf-second-wait** *spf-second-wait*]

no spf-wait

Context

[Tree] (config>service>vprn>ospf3>timers spf-wait)

[Tree] (config>service>vprn>ospf>timers spf-wait)

Full Context

configure service vprn ospf3 timers spf-wait

configure service vprn ospf timers spf-wait

Description

This command configures the maximum interval between two consecutive SPF calculations, in milliseconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs are controlled with this command.

Subsequent SPF runs (if required) occur at exponentially increasing intervals of the **spf-second-wait** interval. For example, if the **spf-second-wait** interval is 1000, the next SPF will run after 2000 milliseconds, and the next SPF after 4000 milliseconds, and so on, until it reaches the **spf-wait** value. The SPF interval stays at the **spf-wait** value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval drops back to **spf-initial-wait** value.

Note: The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity; for example, a configured value of 550 ms is internally rounded down to 500 ms.

The **no** form of this command reverts to the default.

Default

spf-wait 10000 spf-initial-wait 1000 spf-second-wait 1000

Parameters

max-spf-wait

Specifies the maximum interval, in milliseconds, between two consecutive SPF calculations.

Values 10 to 120000

spf-initial-wait

Specifies the initial SPF calculation delay, in milliseconds, after a topology change.

Values 10 to 100000

spf-second-wait

Specifies the hold time, in milliseconds, between the first and second SPF calculation.

Values 10 to 100000

Platforms

7705 SAR Gen 2

spf-wait

Syntax

spf-wait *max-wait* [**initial-wait** *initial-wait*] [**second-wait** *second-wait*]

no spf-wait

Context

[Tree] (config>router>bgp>optimal-route-reflection spf-wait)

Full Context

configure router bgp optimal-route-reflection spf-wait

Description

This command controls the interval between consecutive SPF calculations performed by the TE DB in support of BGP optimal route reflection. The time parameters of this command implement an exponential back-off algorithm.

The **no** form of this command causes a return to default values.

Default

no spf-wait

Parameters

max-wait

Specifies the maximum interval in seconds between two consecutive SPF calculations.

Values 1 to 600

Default 60

initial-wait initial-wait

Specifies the initial SPF calculation delay in seconds after a topology change.

Values 1 to 300

Default 5

second-wait second-wait

Specifies the delay in seconds between the first and second SPF calculation.

Values 1 to 300

Default 15

Platforms

7705 SAR Gen 2

spf-wait

Syntax

spf-wait *spf-wait* [**spf-initial-wait** *initial-wait*] [**spf-second-wait** *second-wait*]]

no spf-wait

Context

[Tree] (config>router>isis>timers spf-wait)

Full Context

configure router isis timers spf-wait

Description

This command configures the maximum interval between two consecutive SPF calculations, in milliseconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs are controlled with this command.

Subsequent SPF runs (if required) occur at exponentially increasing intervals of the **spf-second-wait** interval. For example, if the **spf-second-wait** interval is 1000, the next SPF will run after 2000 milliseconds, and the next SPF after 4000 milliseconds, and so on, until it reaches the **spf-wait** value. The SPF interval stays at the **spf-wait** value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval drops back to the **spf-initial-wait** value.

Note:
The timer granularity is 100 ms. Timer values are rounded down to the nearest granularity; for example, a configured value of 550 ms is internally rounded down to 500 ms.

The **no** form of this command reverts to the default value.

Default

spf-wait 10000 spf-initial-wait 1000 spf-second-wait 1000

Parameters

spf-wait

Specifies the maximum interval, in milliseconds, between two consecutive SPF calculations.

Values 10 to 120000

initial-wait

Specifies the initial SPF calculation delay, in milliseconds, after a topology change.

Values 10 to 100000

second-wait

Specifies the hold time, in milliseconds, between the first and second SPF calculation.

Values 10 to 100000

Platforms

7705 SAR Gen 2

spf-wait

Syntax

spf-wait *max-spf-wait* [**spf-initial-wait** *spf-initial-wait* [**spf-second-wait** *spf-second-wait*]]

no spf-wait

Context

[Tree] (config>router>ospf3>timers spf-wait)

[Tree] (config>router>ospf>timers spf-wait)

Full Context

configure router ospf3 timers spf-wait

configure router ospf timers spf-wait

Description

This command configures the maximum interval between two consecutive SPF calculations, in milliseconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs are controlled with this command.

Subsequent SPF runs (if required) occurs at exponentially increasing intervals of the **spf-second-wait** interval. For example, if the **spf-second-wait** interval is 1000, the next SPF will run after 2000 milliseconds, and the next SPF after 4000 milliseconds, and so on, until it reaches the **spf-wait** value. The SPF interval stays at the **spf-wait** value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval drops back to the **spf-initial-wait** value.

The timer must be entered in increments of 100 milliseconds. Values entered that do not match this requirement are rejected.

Note:
The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

The **no** form of this command reverts to the default value.

Default

spf-wait 10000 spf-initial-wait 1000 spf-second-wait 1000

Parameters

max-spf-wait

Specifies the maximum interval, in milliseconds, between two consecutive SPF calculations.

Values 10 to 120000

spf-initial-wait

Specifies the initial SPF calculation delay, in milliseconds, after a topology change.

Values 10 to 100000

spf-second-wait

Specifies the hold time, in milliseconds, between the first and second SPF calculation.

Values 10 to 100000

Platforms

7705 SAR Gen 2

28.34 spi

spi

Syntax

spi *spi*

no spi

Context

[Tree] (config>ipsec>static-sa spi)

Full Context

configure ipsec static-sa spi

Description

This command configures the SPI key value for an IPsec manual SA.

This command specifies the SPI (Security Parameter Index) used to lookup the instruction to verify and decrypt the incoming IPsec packets when the value of the **direction** command is **inbound**.

The SPI value specifies the SPI that will be used in the encoding of the outgoing packets when the when the value of the **direction** command is **outbound**. The remote node can use this SPI to lookup the instruction to verify and decrypt the packet.

If **no spi** is selected, then this static SA cannot be used.

The **no** form of this command reverts to the default value.

Default

no spi

Parameters

spi

Specifies the security parameter index for this SA.

Values 256 to 16383

Platforms

7705 SAR Gen 2

28.35 split-horizon

split-horizon

Syntax

split-horizon

no split-horizon

Context

[Tree] (config>service>vprn>bgp>group>neighbor split-horizon)

[Tree] (config>service>vprn>bgp>group split-horizon)

[Tree] (config>service>vprn>bgp split-horizon)

Full Context

configure service vprn bgp group neighbor split-horizon

configure service vprn bgp group split-horizon

configure service vprn bgp split-horizon

Description

This command enables the use of split-horizon. When applied globally, to a group, or a specific peer, split-horizon prevents routes from being reflected back to a peer that sends the best route. It applies to routes of all address families and to any type of sending peer; confed-EBGP, EBGP and IBGP.

The configuration default is **no split-horizon**, meaning that no effort is taken to prevent a best route from being reflected back to the sending peer.

Caution:
Use of the **split-horizon** command may have a detrimental impact on peer and route scaling and therefore operators are encouraged to use it only when absolutely needed.

The **no** form of this command disables split horizon command which allows the lower level to inherit the setting from an upper level.

Default

no split-horizon

Platforms

7705 SAR Gen 2

split-horizon

Syntax

split-horizon {**enable** | **disable**}

no split-horizon

Context

[Tree] (config>service>vprn>rip split-horizon)

[Tree] (config>service>vprn>rip>group split-horizon)

[Tree] (config>service>vprn>rip>group>neighbor split-horizon)

[Tree] (config>service>vprn>ripng>group>neighbor split-horizon)

[Tree] (config>service>vprn>ripng split-horizon)

[Tree] (config>service>vprn>ripng>group split-horizon)

Full Context

configure service vprn rip split-horizon

configure service vprn rip group split-horizon

configure service vprn rip group neighbor split-horizon

configure service vprn ripng group neighbor split-horizon

configure service vprn ripng split-horizon

configure service vprn ripng group split-horizon

Description

This command enables the use of split-horizon. RIP uses split horizon with poison reverse to protect from such problems as "counting to infinity". Split horizon with poison reverse means that routes learned from a neighbor through a given interface are advertised in updates out of the same interface but with a metric of 16 (infinity).

The **no** form of this command disables the **split-horizon** command, which allows the lower level to inherit the setting from an upper level.

Default

split-horizon enable

Parameters

enable

Enables split horizon and poison reverse.

disable

Enables split horizon without poison reverse. This allows the routes to be readvertised on interfaces other than the interface that learned the route, with the advertised metric equaling an increment of the metric-in value.This configuration parameter can be set at three levels: global level (applies to all groups and neighbor interfaces), group level

(applies to all neighbor interfaces in the group) or neighbor level (only applies to the specified neighbor interface). The most specific value is used. In particular, if no value is set (**no split-horizon**), the lower level inherits the setting from the less-specific level.

Platforms

7705 SAR Gen 2

split-horizon

Syntax

[**no**] **split-horizon**

Context

[Tree] (config>router>bgp split-horizon)

[Tree] (config>router>bgp>group split-horizon)

[Tree] (config>router>bgp>group>neighbor split-horizon)

Full Context

configure router bgp split-horizon

configure router bgp group split-horizon

configure router bgp group neighbor split-horizon

Description

This command enables the use of split-horizon. Split-horizon prevents routes from being reflected back to a peer that sends the best route. It applies to routes of all address families and to any type of sending peer; confed-EBGP, EBGP and IBGP.

The configuration default is **no split-horizon**, meaning that no effort is taken to prevent a best route from being reflected back to the sending peer.

Default

no split-horizon

Platforms

7705 SAR Gen 2

split-horizon

Syntax

split-horizon {**enable** | **disable**}

no split-horizon

Context

[Tree] (config>router>ripng>group>neighbor split-horizon)

[Tree] (config>router>rip split-horizon)

[Tree] (config>router>ripng>group split-horizon)

[Tree] (config>router>ripng split-horizon)

[Tree] (config>router>rip>group>neighbor split-horizon)

[Tree] (config>router>rip>group split-horizon)

Full Context

configure router ripng group neighbor split-horizon

configure router rip split-horizon

configure router ripng group split-horizon

configure router ripng split-horizon

configure router rip group neighbor split-horizon

configure router rip group split-horizon

Description

This command enables the use of split-horizon.

RIP uses split-horizon with poison-reverse to protect from such problems as "counting to infinity". Split-horizon with poison reverse means that routes learned from a neighbor through a given interface are advertised in updates out of the same interface but with a metric of 16 (infinity).

The **split-horizon disable** command enables split horizon without poison reverse. This allows the routes to be re-advertised on interfaces other than the interface that learned the route, with the advertised metric equaling an increment of the metric-in value.

This configuration parameter can be set at three levels: global level (applies to all groups and neighbor interfaces), group level (applies to all neighbor interfaces in the group) or neighbor level (only applies to the specified neighbor interface). The most specific value is used. In particular if no value is set (**no split-horizon**), the setting from the less specific level is inherited by the lower level.

The **no** form of the command disables split horizon command which allows the lower level to inherit the setting from an upper level.

Default

enabled

Parameters

enable

Specifies enable split horizon and poison reverse.

disable

Specifies disable split horizon allowing routes to be re-advertised on the same interface on which they were learned with the advertised metric incremented by the **metric-in** value.

Platforms

7705 SAR Gen 2

28.36 split-horizon-group

split-horizon-group

Syntax

split-horizon-group [*group-name*] [**residential-group**] [**create**]

Context

[Tree] (config>service>vpls split-horizon-group)

Full Context

configure service vpls split-horizon-group

Description

This command creates a new split horizon group for the VPLS instance. Traffic arriving on a SAP or spoke-SDP within this split horizon group will not be copied to other SAPs or spoke-SDPs in the same split horizon group.

A split horizon group must be created before SAPs and spoke-SDPs can be assigned to the group.

The split horizon group is defined within the context of a single VPLS. The same group-name can be re-used in different VPLS instances.

Up to 30 split horizon groups can be defined per VPLS instance. Half are supported in i-VPLS.

The **no** form of this command removes the group name from the configuration.

Default

A split horizon group is by default not created as a residential-group.

Parameters

group-name

Specifies the name of the split horizon group to which the SDP belongs

residential-group

Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that:

a) SAPs which are members of this Residential Split Horizon Group will have:

- Double-pass queuing at ingress as default setting (can be disabled)
- STP disabled (cannot be enabled)
- ARP reply agent enabled per default (can be disabled)
- MAC pinning enabled per default (can be disabled)

- Downstream broadcast packets are discarded thus also blocking the unknown, flooded traffic
- Downstream multicast packets are allowed when IGMP snooping is enabled

b) Spoke SDPs which are members of this Residential Split Horizon Group will have:

- Downstream multicast traffic supported
- Double-pass queuing is not applicable
- STP is disabled (can be enabled)
- ARP reply agent is not applicable (dhcp-lease-states are not supported on spoke-SDPs)
- MAC pinning enabled per default (can be disabled)

Platforms

7705 SAR Gen 2

split-horizon-group

Syntax

split-horizon-group *name*

no split-horizon-group

Context

[Tree] (config>service>vpls>bgp-evpn>mpls split-horizon-group)

Full Context

configure service vpls bgp-evpn mpls split-horizon-group

Description

This command allows the user to configure an explicit split-horizon-group for all BGP-EVPN MPLS or SRv6 destinations that can be shared by other SAPs and/or spoke SDPs. The use of explicit split-horizon-groups for EVPN-MPLS or SRv6 and spoke SDPs allows the integration of VPLS and EVPN-MPLS or SRv6 networks.

If the **split-horizon-group** command for **bgp-evpn>mpls/srv6** contexts is not used, the default split-horizon-group (that contains all the EVPN destinations) is still used, but it is not possible to refer to it on SAPs/spoke SDPs. User-configured split-horizon-groups can be configured within the service context. The same group-name can be associated to SAPs, spoke SDPs, pw-templates, pw-template-bindings and EVPN-MPLS or SRv6 destinations. The configuration of **bgp-evpn>mpls/srv6> split-horizon-group** is only allowed if **bgp-evpn>mpls/srv6** is shutdown; no changes are allowed when **bgp-evpn>mpls/srv6** is **no shutdown**.

When the SAPs and/or spoke SDPs (manual or BGP-AD-discovered) are configured within the same **split-horizon-group** as the EVPN-MPLS or SRv6 endpoints, MAC addresses are still learned on them but they are not advertised in BGP-EVPN. If provider-tunnel is enabled in the bgp-evpn service, the SAPs and SDP bindings that share the same split-horizon-group of the EVPN-MPLS provider-tunnel are brought operationally down if the point-to-multipoint tunnel is operationally up.

Default

no split-horizon-group

Parameters

name

Specifies the split-horizon-group name.

Platforms

7705 SAR Gen 2

split-horizon-group

Syntax

split-horizon-group *group-name*

no split-horizon-group

Context

[Tree] (config>service>vpls>site split-horizon-group)

Full Context

configure service vpls site split-horizon-group

Description

This command configures the value of split-horizon group associated with this site.

The **no** form of this command reverts the default.

Default

no split-horizon-group

Parameters

group-name

Specifies a split-horizon group name

Platforms

7705 SAR Gen 2

split-horizon-group

Syntax

split-horizon-group *group-name*

no split-horizon-group

Context

[Tree] (config>service>pw-template split-horizon-group)

Full Context

configure service pw-template split-horizon-group

Description

This command creates a new split horizon group (SGH).

Comparing a "residential" SGH and a "regular" SHG is that a residential SHG:

- Has different defaults for the SAP or SDP that belong to this group (ARP reply agent enabled (SAP only), MAC pinning enabled). These can be disabled in the configuration.
- Does not allow enabling spanning tree (STP) on a SAP. It is allowed on an SDP.
- Does not allow for downstream broadcast (broadcast/unknown unicast) on a SAP. It is allowed on an SDP.
- On a SAP, downstream multicast is only allowed when IGMP is enabled (for which an MFIB state exists; only IP multicast); on a SDP, downstream mcast is allowed.

When the feature was initially introduced, residential SHGs were also using ingress shared queuing by default to increase SAP scaling.

A residential SAP (SAP that belongs to a RSHG) is used to scale the number of SAPs in a single VPLS instance. The limit depends on the hardware used and is higher for residential SAPs (where there is no need for egress multicast replication on residential SAPs) than for regular SAPs. Therefore, residential SAPs are useful in residential aggregation environments (for example, triple play networks) with a VLAN/subscriber model.

The **no** form of the command removes the group name from the configuration.

Parameters

group-name

Specifies the name of the split horizon group to which the SDP belongs.

residential-group

Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that:

- SAPs which are members of this Residential Split Horizon Group will have:
  - Double-pass queuing at ingress as default setting (can be disabled)
  - STP disabled (cannot be enabled)
  - ARP reply agent enabled per default (can be disabled)
  - MAC pinning enabled per default (can be disabled)
  - Downstream Broadcast packets are discarded thus also blocking the unknown, flooded traffic
  - Downstream Multicast packets are allowed when IGMP snooping is enabled
- Spoke SDPs which are members of this Residential Split Horizon Group will have:
  - Downstream multicast traffic supported

- Double-pass queuing is not applicable
- STP is disabled (can be enabled)
- ARP reply agent is not applicable on the 7705 SAR Gen 2 (dhcp-lease-states are not supported on spoke SDPs)
- MAC pinning enabled per default (can be disabled)

Platforms

7705 SAR Gen 2

28.37 spoke-sdp

spoke-sdp

Syntax

spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] **endpoint** [**no-endpoint**] [**root-leaf-tag** | **leaf-ac**]

no spoke-sdp *sdp-id*[**:***vc-id*]

Context

[Tree] (config>service>vpls spoke-sdp)

Full Context

configure service vpls spoke-sdp

Description

This command binds a service to an existing service destination point (SDP). A spoke-SDP is treated like the equivalent of a traditional bridge "port" where flooded traffic received on the spoke-SDP is replicated on all other "ports" (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a VPLS service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Default

No *sdp-id* is bound to a service.

Parameters

sdp-id

Specifies the SDP identifier

Values 1 to 17407

vc-id

Specifies the virtual circuit identifier

Values 1 to 4294967295

vc-type

This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the bindings VC type causes the binding to signal the new VC type to the far end when signaling is enabled. VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

The VC type value for Ethernet is 0x0005.

The VC type value for an Ethernet VLAN is 0x0004.

Values ether, vlan

ether

Defines the VC type as Ethernet. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke-SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke-SDP binding. (hex 5)

vlan

Defines the VC type as VLAN. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke-SDP bindings. The VLAN VC-type inserts one dot1q tag within each encapsulated Ethernet packet transmitted to the far end and strips one dotQ tag, if a tag is present, from traffic received on the pseudowire.

Note: The system expects a symmetrical configuration with its peer, specifically it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. As some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. Consequently, protocol extractions will not necessarily function for asymmetrical configurations as they would with a symmetrical configurations resulting in an unexpected operation.

group-name

Specifies the name of the split horizon group to which the SDP belongs

endpoint

Specifies the service endpoint to which this SDP bind is attached. The service ID of the SDP binding must match the service ID of the service endpoint.

no endpoint

Removes the association of a spoke-SDP with an explicit endpoint name

root-leaf-tag

Specifies a tagging spoke-SDP under an E-Tree VPLS. When a tag SDP binding is required, it is created with a root-leaf-tag flag. Only VLAN tag SDP bindings are supported. The VLAN type must be set to VC VLAN type. The root-leaf-tag parameter indicates this SDP binding is a tag SDP that will use a default VID tag of 1 for root and 2 for leaf. The SDP binding tags egress E-Tree traffic with root and leaf VIDs as appropriate. Root and leaf VIDs are only significant between peering VPLS but the values must be consistent on each end. On ingress a tag SDP binding removes the VID tag on the interface between VPLS in the same E-Tree service. The tag SDP receives root tagged traffic and marks the traffic with a root indication internally. This option is not available on BGP EVPN-enabled E-Tree services.

leaf-ac

Specifies an access (AC) spoke-SDP binding under a E-Tree VPLS as a leaf access (AC) SDP. The default E-Tree SDP binding type is a root-ac if *leaf-ac* or *root-leaf-tag* is not specified at SDP creation. This option is only available when the VPLS is designated as an E-Tree VPLS. BGP EVPN-enabled E-Tree VPLS services support the **leaf-ac** option.

Platforms

7705 SAR Gen 2

spoke-sdp

Syntax

spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **ipipe**}] [**create**]

no spoke-sdp *sdp-id*[**:***vc-id*]

Context

[Tree] (config>service>ies>if spoke-sdp)

[Tree] (config>service>vprn>if spoke-sdp)

Full Context

configure service ies interface spoke-sdp

configure service vprn interface spoke-sdp

Description

This command binds a service to an existing service destination point (SDP).

A spoke SDP is treated like the equivalent of a traditional bridge "port" where flooded traffic received on the spoke SDP is replicated on all other "ports" (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service is down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with an IES service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

Class-based forwarding is not supported on a spoke SDP used for termination on an IES or VPRN services. All packets are forwarded over the default LSP.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router. The spoke SDP must be shut down first before it can be deleted from the configuration.

Default

no spoke-sdp

Parameters

sdp-id

Specifies the SDP identifier. Allowed values are integers in the range of 1 and 17407 for existing SDPs.

vc-id

Specifies the virtual circuit identifier.

Values 1 to 4294967295

vc-type

Specifies the encapsulation and pseudowire type for the spoke SDP.

Values ether: specifies Ethernet pseudowire as the type of virtual circuit (VC) associated with the SDP binding

ipipe: specifies Ipipe pseudowire as the type of virtual circuit (VC) associated with the SDP binding

Default ether

create

Keyword used to create the spoke SDP. The create keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

spoke-sdp

Syntax

spoke-sdp *sdp-id:vc-id*

no spoke-sdp

Context

[Tree] (config>service>vpls>site spoke-sdp)

Full Context

configure service vpls site spoke-sdp

Description

This command binds a service to an existing service destination point (SDP).

The **no** form of this command removes the parameter from the configuration.

Parameters

sdp-id

Specifies the SDP identifier

Values 1 to 17407

vc-id

Specifies the virtual circuit identifier.

Values 1 to 429496729

Platforms

7705 SAR Gen 2

spoke-sdp

Syntax

spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] [**no-endpoint**]

spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **endpoint** *endpoint-name* [**icb**]

no spoke-sdp *sdp-id*[**:***vc-id*]

Context

[Tree] (config>service>epipe spoke-sdp)

Full Context

configure service epipe spoke-sdp

Description

This command binds a service to an existing service destination point (SDP). A spoke SDP is treated like the equivalent of a traditional bridge "port" where flooded traffic received on the spoke SDP is replicated on all other "ports" (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with an Epipe, VPLS, VPRN, VPRN service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

This command can also be used to associate a GRE tunnel carrying Ethernet payload with an Epipe and terminate it on a PW port referenced within the same Epipe service. The spoke SDP represents a L2oGRE tunnel with SDP delivery type set to **eth-gre-bridged**. With this configuration, the **vc-id** is unused since there is no multiplexing of Ethernet payload within the same tunnel. The **vc-id** value is included only to maintain the expected spoke SDP structure within an EPIPE service. For L2oGRE tunnels, the **vc-id** can be set to any arbitrary value within its configurable range.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Default

No *sdp-id* is bound to a service.

Parameters

sdp-id

The SDP identifier.

Values 1 to 17407

vc-id

The virtual circuit identifier. The VC-ID is not used with L2TPv3 SDPs or L2oGRE tunnels, however it must be configured.

Values 1 to 4294967295

vc-type

This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the bindings VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

The VC type value for Ethernet is 0x0005.

The VC type value for an Ethernet VLAN is 0x0004.

Values ethernet

ether

Defines the VC type as Ethernet. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings.

Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke SDP binding.

vlan

Defines the VC type as VLAN. The top VLAN tag, if a VLAN tag is present, is stripped from traffic received on the pseudowire, and a VLAN tag is inserted when forwarding into the pseudowire. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings.

The VLAN VC-type requires at least one dot1q tag within each encapsulated Ethernet packet transmitted to the far end.

Note: The system expects a symmetrical configuration with its peer, specifically it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. As some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. Consequently, protocol extractions will not necessarily function for asymmetrical configurations as they would with a symmetrical configurations resulting in an unexpected operation.

no-endpoint

Removes the association of a spoke SDP with an explicit endpoint name.

endpoint-name

Specifies the name of the service endpoint.

icb

Specifies the spoke SDP as an inter-chassis backup SDP binding.

Platforms

7705 SAR Gen 2

spoke-sdp

Syntax

[**no**] **spoke-sdp** *spoke-id*

Context

[Tree] (config>service>vpls>mac-move>secondary-ports spoke-sdp)

[Tree] (config>service>vpls>mac-move>primary-ports spoke-sdp)

Full Context

configure service vpls mac-move secondary-ports spoke-sdp

configure service vpls mac-move primary-ports spoke-sdp

Description

This command declares a specified spoke-SDP as a primary (or secondary) VPLS port.

Parameters

spoke-id

Specifies the SDP ID to configure as the primary VPLS port

Values 1 to 17407

vc-id

Specifies the virtual circuit identifier

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

spoke-sdp

Syntax

spoke-sdp *sdp-id:vc-id* [**create**] [**no-endpoint**]

spoke-sdp *sdp-id:vc-id* [**create**] **endpoint** *name* [**icb**]

no sdp *sdp-id:vc-id*

Context

[Tree] (config>mirror>mirror-dest spoke-sdp)

[Tree] (config>mirror>mirror-dest>remote-source spoke-sdp)

Full Context

configure mirror mirror-dest spoke-sdp

configure mirror mirror-dest remote-source spoke-sdp

Description

This command binds an existing (mirror) service distribution path (SDP) to the mirror destination service ID.

Spoke SDPs are used to send and receive mirrored traffic between mirror source and destination routers in a remote mirroring solution. A spoke SDP configured in the remote-source context (**remote-src>spoke-sdp**) is used on the destination router. A spoke SDP configured in the mirror service context (**mirror-dest>spoke-sdp**) is used on the source router.

The destination node should be configured with **remote-src>spoke-sdp** entries when using L2TPv3, MPLS-TP or LDP IPv6 LSP SDPs in the remote mirroring solution. For all other types of SDPs, r**emote-source>far-end** entries should be used.

Spoke SDPs are not applicable when routable LI encapsulation is employed (mirror-dest>encap).

A mirror destination service that is configured for a destination router must not be configured as for a source router.

The **no** form of this command removes the SDP binding from the mirror destination service.

Default

An SDP ID is bound to a mirror destination service ID. If no SDP is bound to the service, the mirror destination will be local and cannot be sent to another router over the core network.

Parameters

sdp-id:vc-id

Specifies a locally unique SDP identification (ID) number. The SDP ID must exist. If the SDP ID does not exist, an error will occur and the command will not execute.

For mirror services, the *vc-id* defaults to the *service-id*. However, there are scenarios where the *vc-id* is being used by another service. In this case, the SDP binding cannot be created. So, to avoid this, the mirror service SDP bindings now accepts *vc-ids*.

Values 1 to 17407

no-endpoint

Removes the association of a SAP or a SDP with an explicit endpoint name.

name

Specifies the name of the endpoint associated with the SAP.

icb

Indicates that the SDP is of type Inter-Chassis Backup (ICB). This is a special pseudowire used for MC-LAG and pseudowire redundancy application.

An explicitly named endpoint can have a maximum of one SAP and one ICB. Once a SAP is added to the endpoint, only one more object of type ICB SDP is allowed. The ICB SDP cannot be added to the endpoint if the SAP is not part of a MC-LAG instance. This means that all other SAP types cannot exist on the same endpoint as an ICB SDP since non Ethernet SAP cannot be part of a MC-LAG instance. Conversely, a SAP which is not part of a MC-LAG instance cannot be added to an endpoint which already has an ICB SDP.

An explicitly named endpoint, which does not have a SAP object, can have a maximum of four SDPs, which can include any of the following: a single primary SDP, one or many secondary SDPs with precedence, and a single ICB SDP.

Default Null. The user should explicitly configure this option at create time. The user can remove the ICB type simply by retyping the SDP configuration without the **icb** keyword.

Platforms

7705 SAR Gen 2

spoke-sdp

Syntax

spoke-sdp *sdp-id:vc-id* [**create**]

no spoke-sdp *sdp-id:vc-id*

Context

[Tree] (config>service>vprn>ip-mirror-interface spoke-sdp)

Full Context

configure service vprn ip-mirror-interface spoke-sdp

Description

This command binds a service to an existing SDP.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with the VPRN service. SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router. The spoke SDP must be shut down before it can be deleted from the configuration.

Parameters

sdp-id

Specifies SDP identifier.

Values 1 to 32767

vc-id

Specifies the virtual circuit identifier.

Values 1 to 4294967295

create

Keyword used to create an IP mirror interface.

Platforms

7705 SAR Gen 2

28.38 spoke-sdp-fec

spoke-sdp-fec

Syntax

spoke-sdp-fec

spoke-sdp-fec *spoke-sdp-fec-id* [**fec** *fec-type*] [**aii-type** *aii-type*] [**create**]

spoke-sdp-fec *spoke-sdp-fec-id* **no-endpoint**

spoke-sdp-fec *spoke-sdp-fec-id* [**fec** *fec-type*] [**aii-type** *aii-type*] [**create**] **endpoint** *name* [**icb**]

Context

[Tree] (config>service>epipe spoke-sdp-fec)

Full Context

configure service epipe spoke-sdp-fec

Description

This command binds a service to an existing service destination point (SDP), using a dynamic MS-PW.

A spoke SDP is treated like the equivalent of a traditional bridge "port" where flooded traffic received on the spoke SDP is replicated on all other "ports" (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

When using dynamic MS-PWs, the particular SDP to bind-to is automatically selected based on the Target Attachment Individual Identifier (TAII) and the path to use, specified under spoke SDP FEC. The selected SDP will terminate on the first hop S-PE of the MS-PW. Therefore, an SDP must already be defined in the config>service>sdp context that reaches the first hop router of the MS-PW. The router will in order to associate an SDP with a service. If an SDP to that is not already configured, an error message is generated. If the sdp-id does exist, a binding between that sdp-id and the service is created.

It differs from the spoke-sdp command in that the spoke-sdp command creates a spoke SDP binding that uses a pseudowire with the PW ID FEC. However, the spoke-sdp-fec command enables pseudowires with other FEC types to be used. Only the Generalized ID FEC (FEC129) may be specified using this command.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Parameters

spoke-sdp-fec-id

An unsigned integer value identifying the spoke SDP.

Values 1 to 4294967295

fec-type

An unsigned integer value for the type of the FEC used by the MS-PW.

Values 129 to 130

aii-type

An unsigned integer value for the Attachment Individual Identifier (AII) type used to identify the MS-PW endpoints.

Values 1 to 2

endpoint-name

Specifies the name of the service endpoint.

no endpoint

Adds or removes a spoke SDP association.

icb

Configures the spoke SDP as an inter-chassis backup SDP binding.

Platforms

7705 SAR Gen 2

28.39 spt-switchover-threshold

spt-switchover-threshold

Syntax

spt-switchover-threshold {*grp-ip-address/mask* | *grp-ip-address netmask*} *spt-threshold*

spt-switchover-threshold *grp-ipv6-addr/prefix-length spt-threshold*

no spt-switchover-threshold {*grp-ip-address/mask* | *grp-ip-address netmask*}

no spt-switchover-threshold *grp-ipv6-addr/prefix-length*

Context

[Tree] (config>service>vprn>pim spt-switchover-threshold)

Full Context

configure service vprn pim spt-switchover-threshold

Description

This command configures a shortest path tree (SPT tree) switchover threshold for a group prefix.

Parameters

grp-ip-address

Specifies the multicast group address.

grp-ipv6-address

Specifies the multicast group address.

prefix-length

Specifies the address prefix length.

Values

grp-ipv6-address	: x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x [0 to FFFF]H

d [0 to 255]D

prefix-length [1 to 128]

mask

Defines the mask of the multicast-ip-address.

Values 4 to 32

netmask

The subnet mask in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

spt-threshold

Specifies the configured threshold in kilobits per second (kb/s) for the group to which this (S,G) belongs. For a group G configured with a threshold, switchover to SPT for an (S,G) is attempted only if the (S,G)'s rate exceeds this configured threshold.

Platforms

7705 SAR Gen 2

spt-switchover-threshold

Syntax

spt-switchover-threshold {*grp-ipv4-prefixIipv4-prefix-length* | *grp-ipv4-prefix netmask* | *grp-ipv6-prefixIipv6-prefix-length*} *spt-threshold*

no spt-switchover-threshold {*grp-ipv4-prefixIipv4-prefix-length* | *grp-ipv4-prefix netmask* | *grp-ipv6-prefixIipv6-prefix-length*}

Context

[Tree] (config>router>pim spt-switchover-threshold)

Full Context

configure router pim spt-switchover-threshold

Description

This command configures shortest path (SPT) tree switchover thresholds for group prefixes.

PIM-SM routers with directly connected routers receive multicast traffic initially on a shared tree rooted at the Rendezvous Point (RP). Once the traffic arrives on the shared tree and the source of the traffic is known, a switchover to the SPT tree rooted at the source is attempted.

For a group that falls in the range of a prefix configured in the table, the corresponding threshold value determines when the router should switch over from the shared tree to the source specific tree. The switchover is attempted only if the traffic rate on the shared tree for the group exceeds the configured threshold.

In the absence of any matching prefix in the table, the default behavior is to switchover when the first packet is seen. In the presence of multiple prefixes matching a given group, the most specific entry is used.

The **no** form of this command removes the parameters from the PIM configuration.

Parameters

grp-ipv4-prefix

Specifies the group IPv4 multicast address in dotted decimal notation.

Values a.b.c.d

ipv4-prefix-length

Specifies the length of the IPv4 prefix.

Values 4 to 32

netmask

Specifies the netmask associated with the IPv4 prefix, expressed in dotted decimal notation. Network bits must be 1, and host bits must be 0.

Values a.b.c.d

grp-ipv6-prefix

Specifies the group IPv6 multicast address in hexadecimal notation.

Values xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

xx — 0 to FF (hex)

ipv6-prefix-length

Specifies the length of the IPv6 prefix.

Values 8 to 128

spt-threshold

Specifies the configured threshold in kilobits per second (kb/s) for a group prefix. A switchover is attempted only if the traffic rate on the shared tree for the group exceeds this configured threshold. When the **infinity** keyword is specified, no switchover will occur at any time, regardless of the traffic level is detected.

Values 1 to 4294967294, infinity

Platforms

7705 SAR Gen 2

28.40 sr-isis

sr-isis

Syntax

[**no**] **sr-isis**

Context

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-isis)

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-isis)

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel>resolution-filter sr-isis)

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-isis)

Full Context

configure service epipe bgp-evpn mpls auto-bind-tunnel resolution-filter sr-isis

configure service vpls bgp-evpn mpls auto-bind-tunnel resolution-filter sr-isis

configure service vprn bgp-ipvpn mpls auto-bind-tunnel resolution-filter sr-isis

configure service vprn bgp-evpn mpls auto-bind-tunnel resolution-filter sr-isis

Description

This command selects the Segment Routing (SR) tunnel type programed by an IS-IS instance in TTM.

When the **sr-isis** value (or **sr-ospf**) is enabled, an SR tunnel to the BGP next hop is selected in the TTM from the lowest-numbered IS-IS (OSPF) instance.

The **no** form of this command disables the SR-ISIS setting for the auto-bind tunnel.

Default

no sr-isis

Platforms

7705 SAR Gen 2

sr-isis

Syntax

[**no**] **sr-isis**

Context

[Tree] (config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter sr-isis)

Full Context

configure router static-route-entry indirect tunnel-next-hop resolution-filter sr-isis

Description

This command enables the use of SR-ISIS sourced tunnel entries in the TTM to resolve the associated static route next hop.

Default

no sr-isis

Platforms

7705 SAR Gen 2

sr-isis

Syntax

[**no**] **sr-isis**

Context

[Tree] (config>service>sdp sr-isis)

Full Context

configure service sdp sr-isis

Description

This command configures an MPLS SDP of LSP type ISIS Segment Routing. The SDP of LSP type sr-isis can be used with the far-end option. The signaling protocol for the service labels for an SDP using an SR tunnel can be configured to static (off), T-LDP (tldp), or BGP (bgp).

Platforms

7705 SAR Gen 2

sr-isis

Syntax

[**no**] **sr-isis**

Context

[Tree] (config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter sr-isis)

[Tree] (config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter sr-isis)

Full Context

configure router bgp next-hop-resolution shortcut-tunnel family resolution-filter sr-isis

configure router bgp next-hop-resolution labeled-routes transport-tunnel family resolution-filter sr-isis

Description

This command selects the Segment Routing (SR) tunnel type programmed by an IS-IS instance in TTM for next-hop resolution of BGP routes and labeled routes. This option allows BGP to use the segment- routing tunnel in the tunnel table submitted by the lowest preference IS-IS instance or, in case of a tie, the lowest numbered IS-IS instance.

Platforms

7705 SAR Gen 2

sr-isis

Syntax

[**no**] **sr-isis**

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls sr-isis)

Full Context

configure oam-pm session ip tunnel mpls sr-isis

Description

This command configures the specification of **sr-isis** specific tunnel information that is used to transport the test packets. Entering this context removes all other tunnel type options configured under the **configure oam-pm session ip tunnel mpls** context. Only a single **mpls** type can be configured for an OAM-PM session.

The **no** form of this command deletes the context and all configurations under it.

Parameters

ipv4-address

Specifies IPv4 address.

Values ipv4-address: a.b.c.d (host bits must be 0)

Platforms

7705 SAR Gen 2

sr-isis

Syntax

sr-isis

Context

[Tree] (config>service>vprn>auto-bind-tunnel>res-filter sr-isis)

Full Context

configure service vprn auto-bind-tunnel resolution-filter sr-isis

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

28.41 sr-label-index

sr-label-index

Syntax

sr-label-index {*value* | *param-name*} [**prefer-igp**]

no sr-label-index

Context

[Tree] (config>router>policy-options>policy-statement>default-action sr-label-index)

[Tree] (config>router>policy-options>policy-statement>entry>action sr-label-index)

Full Context

configure router policy-options policy-statement default-action sr-label-index

configure router policy-options policy-statement entry action sr-label-index

Description

This command associates a BGP segment-routing label index value with all /32 BGP labeled IPv4 routes matching the entry or policy **default-action**.

Note:
Avoid using this action in a policy entry that matches more than one /32 label-ipv4 route, otherwise SID conflicts are created.

The **sr-label-index** action only takes effect in BGP peer import policies (and only on received /32 label-ipv4 routes) and in route-table-import policies associated with the label-ipv4 RIB.

The **prefer-igp** applies only in a route-table-import policy. If **prefer-igp** is specified and BGP segment-routing uses **prefix-sid-range global**, then BGP tries, as a first priority, to use the IGP segment routing label index for the IGP route matched by the **route-table-import** policy. If the IGP route does not have an SID index, or **prefer-igp** is not configured or **prefix-sid-range** is not **global**, BGP tries to use the label index value specified by this command.

When this action occurs in a policy applied as a peer-import policy, it can add a prefix SID attribute to a received /32 label-ipv4 route that was not sent with this attribute, or it can replace the received prefix SID attribute with a new one.

If this command specifies an index value that causes a SID conflict with another BGP route, then all conflicting BGP routes are re-advertised with label values based on dynamic allocation rather than SID-based allocation.

If this command specifies an index value that causes a SID conflict with an IGP route, the BGP route is re-advertised with a label value based on dynamic allocation rather than an SID-based allocation.

The **no** form of this command causes matched BGP routes to be advertised without any new or changed prefix SID attributes.

Default

no sr-label-index

Parameters

value

Specifies the BGP segment routing label index to associate with the matched route or routes.

Values 0 to 52487

param-name

Specifies the **type** parameter variable name, up to 32 characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@”.

prefer-igp

A keyword that is applicable only in **route-table-import** policies, to instruct BGP to borrow the SID index from the IGP route if it has an SID index and the **prefix-sid-range** is **global**.

Platforms

7705 SAR Gen 2

28.42 sr-labels

sr-labels

Syntax

sr-labels start *start-value* **end** *end-value*

no sr-labels

Context

[Tree] (config>router>mpls-labels sr-labels)

Full Context

configure router mpls-labels sr-labels

Description

This command configures the range of the Segment Routing Global Block (SRGB). It is a label block which is used for assigning labels to segment routing prefix SIDs originated by this router. This range is carved from the system dynamic label range and is not instantiated by default.

This is a reserved label and once configured it cannot be used by other protocols such as RSVP, LDP, and BGP to assign a label dynamically.

Default

no sr-labels

Parameters

start-value

Specifies the start label value in the SRGB

Values 18432 to 524287 within dynamic label range | 1048575 (FP4 or FP5 only)

end-value

Specifies the end label value in the SRGB

Values 18432 to 524287 within dynamic label range | 1048575 (FP4 or FP5 only)

Platforms

7705 SAR Gen 2

28.43 sr-maintenance-policy

sr-maintenance-policy

Syntax

sr-maintenance-policy *maintenance-policy-name*

no sr-maintenance-policy

Context

[Tree] (config>router>policy-options>policy-statement>entry>action sr-maintenance-policy)

[Tree] (config>router>policy-options>policy-statement>default-action sr-maintenance-policy)

Full Context

configure router policy-options policy-statement entry action sr-maintenance-policy

configure router policy-options policy-statement default-action sr-maintenance-policy

Description

This command applies a named segment routing maintenance policy to the matching routes. It is only used for SR policy routes. The named policy must exist under the **config**>**router**>**segment-routing** context.

The **no** form of this command removes the specified maintenance policy.

Parameters

maintenance-policy-name

Specifies the name of the maintenance policy, up to 32 characters and cannot start with a space or underscore.

Platforms

7705 SAR Gen 2

28.44 sr-mpls

sr-mpls

Syntax

sr-mpls

Context

[Tree] (config>router>segment-routing sr-mpls)

Full Context

configure router segment-routing sr-mpls

Description

Commands in this context configure the SR MPLS properties.

Platforms

7705 SAR Gen 2

28.45 sr-mpls-local

sr-mpls-local

Syntax

sr-mpls-local {**none** | **all**}

Context

[Tree] (config>router>ttl-propagate sr-mpls-local)

Full Context

configure router ttl-propagate sr-mpls-local

Description

This command configures TTL or hop-limit propagation for all segment routing MPLS tunnels carrying IPv4 or IPv6 packets. This applies to IPv4 and IPv6 packets of IGP, BGP unlabelled (except 6PE), and static routes in the base router whose next hop is resolved to a Segment Routing MPLS (SR-MPLS) tunnel of any of the following types: SR-ISIS, SR-OSPF, SR-OSPF3, SR-TE LSP, and SR policy.

This command configures TTL or hop-limit propagation for CPM originated IP packets. Use the **sr-mpls-transit** command to configure TTL or hop-limit propagation for transit IP packets.

Default

sr-mpls-local all

Parameters

none

Keyword to specify that the IP TTL or hop limit is not propagated into the segment routing transport label stack.

all

Keyword to specify that the IP TTL or hop limit is propagated to all labels in the segment routing transport label stack.

Platforms

7705 SAR Gen 2

28.46 sr-mpls-transit

sr-mpls-transit

Syntax

sr-mpls-transit {**none** | **all**}

Context

[Tree] (config>router>ttl-propagate sr-mpls-transit)

Full Context

configure router ttl-propagate sr-mpls-transit

Description

This command configures TTL or hop-limit propagation for all segment routing MPLS tunnels carrying IPv4 or IPv6 packets. This applies to IPv4 and IPv6 packets of IGP, BGP unlabelled (except 6PE), and static routes in the base router whose next hop is resolved to a Segment Routing MPLS (SR-MPLS) tunnel of any of the following types: SR-ISIS, SR-OSPF, SR-OSPF3, SR-TE LSP, and SR policy.

This command configures TTL or hop-limit propagation for transit IP packets. Transit IP packets are packets of base router prefixes received on an access interface or a network interface (with or without tunnel encapsulation) and whose FIB lookup results in forwarding them over an SR-MPLS tunnel. Use the **sr-mpls-local** command to configure TTL or hop-limit propagation for CPM originated IP packets.

Default

sr-mpls-transit all

Parameters

none

Keyword to specify that the IP TTL or hop limit is not propagated into the segment routing transport label stack.

all

Keyword to specify that the IP TTL or hop limit is propagated to all labels in the segment routing transport label stack.

Platforms

7705 SAR Gen 2

28.47 sr-ospf

sr-ospf

Syntax

[**no**] **sr-ospf**

Context

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel>resolution-filter sr-ospf)

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-ospf)

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-ospf)

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-ospf)

Full Context

configure service vprn bgp-ipvpn mpls auto-bind-tunnel resolution-filter sr-ospf

configure service epipe bgp-evpn mpls auto-bind-tunnel resolution-filter sr-ospf

configure service vprn bgp-evpn mpls auto-bind-tunnel resolution-filter sr-ospf

configure service vpls bgp-evpn mpls auto-bind-tunnel resolution-filter sr-ospf

Description

This command selects the Segment Routing (SR) tunnel type programed by an OSPF instance in TTM.

When the **sr-ospf** (or **sr-isis**) value is enabled, an SR tunnel to the BGP next hop is selected in the TTM from the lowest-numbered IS-IS (OSPF) instance.

The **no** form of this command disables the SR-OSPF setting for the auto-bind tunnel.

Default

no sr-ospf

Platforms

7705 SAR Gen 2

sr-ospf

Syntax

[**no**] **sr-ospf**

Context

[Tree] (config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter sr-ospf)

Full Context

configure router static-route-entry indirect tunnel-next-hop resolution-filter sr-ospf

Description

This command enables the use of SR-OSPF sourced tunnel entries in the TTM to resolve the associated static route next hop.

Default

no sr-ospf

Platforms

7705 SAR Gen 2

sr-ospf

Syntax

[**no**] **sr-ospf**

Context

[Tree] (config>service>sdp sr-ospf)

Full Context

configure service sdp sr-ospf

Description

This command configures an MPLS SDP of LSP type OSPF Segment Routing. The SDP of LSP type sr-ospf can be used with the far-end option. The signaling protocol for the service labels for an SDP using an SR tunnel can be configured to static (off), T-LDP (tldp), or BGP (bgp).

Platforms

7705 SAR Gen 2

sr-ospf

Syntax

[**no**] **sr-ospf**

Context

[Tree] (config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter sr-ospf)

[Tree] (config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter sr-ospf)

Full Context

configure router bgp next-hop-resolution shortcut-tunnel family resolution-filter sr-ospf

configure router bgp next-hop-resolution labeled-routes transport-tunnel family resolution-filter sr-ospf

Description

This command selects the Segment Routing (SR) tunnel type programmed by an OSPF instance in TTM for next-hop resolution of BGP routes and labeled routes. This option allows BGP to use the segment routing tunnel in the tunnel table submitted by the lowest preference OSPF instance or, in case of a tie, the lowest numbered OSPF instance.

The **no** form of this command disables the use of SR-OSPF tunneling for next-hop resolution.

Platforms

7705 SAR Gen 2

sr-ospf

Syntax

[**no**] **sr-ospf**

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls sr-ospf)

Full Context

configure oam-pm session ip tunnel mpls sr-ospf

Description

This commmand configures the specification of **sr-ospfv3** specific tunnel information that is used to transport the test packets. Entering this context removes all other tunnel type options configured under the **configure oam-pm session ip tunnel mpls** context. Only a single **mpls** type can be configured for an OAM-PM session.

The **no** form of this command deletes the context and all configurations under it.

Platforms

7705 SAR Gen 2

sr-ospf

Syntax

sr-ospf

Context

[Tree] (config>service>vprn>auto-bind-tunnel>res-filter sr-ospf)

Full Context

configure service vprn auto-bind-tunnel resolution-filter sr-ospf

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

28.48 sr-ospf3

sr-ospf3

Syntax

[**no**] **sr-ospf3**

Context

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-ospf3)

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-ospf3)

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-ospf3)

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel>resolution-filter sr-ospf3)

Full Context

configure service vpls bgp-evpn mpls auto-bind-tunnel resolution-filter sr-ospf3

configure service vprn bgp-evpn mpls auto-bind-tunnel resolution-filter sr-ospf3

configure service epipe bgp-evpn mpls auto-bind-tunnel resolution-filter sr-ospf3

configure service vprn bgp-ipvpn mpls auto-bind-tunnel resolution-filter sr-ospf3

Description

This command selects the Segment Routing (SR) tunnel type programed by an OSPFv3 instance in TTM.

When the **sr-ospf3** (or **sr-isis**) command is enabled, an SR tunnel to the BGP next hop is selected in the TTM from the lowest-numbered IS-IS (OSPFv3) instance.

The **no** form of this command disables the OSPFv3 setting for the auto-bind tunnel.

Default

no sr-ospf3

Platforms

7705 SAR Gen 2

sr-ospf3

Syntax

[**no**] **sr-ospf3**

Context

[Tree] (config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter sr-ospf3)

[Tree] (config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter sr-ospf3)

Full Context

configure router bgp next-hop-resolution labeled-routes transport-tunnel family resolution-filter sr-ospf3

configure router bgp next-hop-resolution shortcut-tunnel family resolution-filter sr-ospf3

Description

This command selects the IPv6 segment routing tunnel type programmed by an OSPFv3 instance in the TTMv6 for next-hop resolution of BGP routes and labeled routes. This option allows BGP to use the segment routing tunnel in the tunnel table submitted by the lowest preference OSPFv3 instance or, in case of a tie, the lowest-numbered OSPFv3 instance.

The **no** form of this command disables the use of SR-OSPF3 for next-hop resolution.

Default

no sr-ospf3

Platforms

7705 SAR Gen 2

sr-ospf3

Syntax

sr-ospf3

Context

[Tree] (config>service>vprn>auto-bind-tunnel>res-filter sr-ospf3)

Full Context

configure service vprn auto-bind-tunnel resolution-filter sr-ospf3

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

28.49 sr-policies

sr-policies

Syntax

sr-policies

Context

[Tree] (config>router>segment-routing sr-policies)

Full Context

configure router segment-routing sr-policies

Description

This command creates the context to configure segment routing policies. A segment routing policy specifies traffic to be matched by the policy and actions to take on the matched traffic by applying the instructions encoded in one or more segment lists.

Platforms

7705 SAR Gen 2

28.50 sr-policy

sr-policy

Syntax

[**no**] **sr-policy**

Context

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel>resolution-filter sr-policy)

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-policy)

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-policy)

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-policy)

Full Context

configure service vprn bgp-ipvpn mpls auto-bind-tunnel resolution-filter sr-policy

configure service vpls bgp-evpn mpls auto-bind-tunnel resolution-filter sr-policy

configure service epipe bgp-evpn mpls auto-bind-tunnel resolution-filter sr-policy

configure service vprn bgp-evpn mpls auto-bind-tunnel resolution-filter sr-policy

Description

This command selects the tunnel type for the SR policy.

The **sr-policy** value instructs BGP to search for an SR policy with a non-null endpoint and color value that matches the BGP next hop and color extended community value of the EVPN route.

The **no** form of this command disables the SR policy setting for the auto-bind tunnel.

Default

no sr-policy

Platforms

7705 SAR Gen 2

sr-policy

Syntax

sr-policy

sr-policy color *color-id* **endpoint** *ip-address*

Context

[Tree] (config>saa>test>type-multi-line>lsp-ping sr-policy)

Full Context

configure saa test type-multi-line lsp-ping sr-policy

Description

This command configures the SR policy target FEC.

Note:
The **sr-policy** target FEC type is supported under the OAM context and under **type-multi-line node** in the SAA context.

Parameters

color color

Specifies the color ID.

Values 0 to 4294967295

endpoint *ip-address*

Specifies the endpoint address.

Values

ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

Platforms

7705 SAR Gen 2

sr-policy

Syntax

sr-policy

Context

[Tree] (config>service>vprn>auto-bind-tunnel>res-filter sr-policy)

Full Context

configure service vprn auto-bind-tunnel resolution-filter sr-policy

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

sr-policy

Syntax

[**no**] **sr-policy**

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls sr-policy)

Full Context

configure oam-pm session ip tunnel mpls sr-policy

Description

Commands in this context identify the SR policy used to tunnel IP packets for session tests.

The **no** form of this command disables the SR policy used to tunnel IP packets.

Default

no sr-policy

Platforms

7705 SAR Gen 2

28.51 sr-policy-import

sr-policy-import

Syntax

[**no**] **sr-policy-import**

Context

[Tree] (config>router>bgp sr-policy-import)

Full Context

configure router bgp sr-policy-import

Description

This command instructs BGP to import all statically-configured non-local segment routing policies from the segment routing DB into the BGP RIB so that they can be advertised, as originated routes, towards BGP peers supporting the **sr-policy-ipv4** address family.

The **no** form of this command instructs BGP to not import any statically defined segment routing policies into BGP.

Default

no sr-policy-import

Platforms

7705 SAR Gen 2

28.52 sr-te

sr-te

Syntax

[**no**] **sr-te**

Context

[Tree] (config>service>vprn>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-te)

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-te)

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter sr-te)

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel>resolution-filter sr-te)

Full Context

configure service vprn bgp-evpn mpls auto-bind-tunnel resolution-filter sr-te

configure service epipe bgp-evpn mpls auto-bind-tunnel resolution-filter sr-te

configure service vpls bgp-evpn mpls auto-bind-tunnel resolution-filter sr-te

configure service vprn bgp-ipvpn mpls auto-bind-tunnel resolution-filter sr-te

Description

This command selects the Segment Routing (SR) Traffic Engineered (SR-TE) LSP programmed in TTM.

The **sr-te** value instructs the system to search for the best metric SR-TE LSP to the address of the BGP next hop. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple SR-TE LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel ID.

The **no** form of this command disables the SR-TE LSP setting for the auto-bind tunnel.

Default

no sr-te

Platforms

7705 SAR Gen 2

sr-te

Syntax

[**no**] **sr-te**

Context

[Tree] (config>router>mpls>pce-initiated-lsp sr-te)

Full Context

configure router mpls pce-initiated-lsp sr-te

Description

This command enables support for SR-TE PCE-initiated LSPs.

The **no** form of this command removes SR-TE PCE-initiated LSP support. All PCE-initiated SR-TE LSPs are deleted.

Platforms

7705 SAR Gen 2

sr-te

Syntax

sr-te *value*

no sr-te

Context

[Tree] (config>router>mpls>tunnel-table-pref sr-te)

Full Context

configure router mpls tunnel-table-pref sr-te

Description

This command configures the tunnel table preference for an SR-TE LSP tunnel type away from its default value.

The tunnel table preference applies to the next-hop resolution of BGP routes of the following families: EVPN, IPv4, IPv6, VPN-IPv4, VPN-IPv6, label-IPv4, and label-IPv6 in the tunnel table.

This feature does not apply to a VPRN, VPLS, or VLL service with explicit binding to an SDP that enabled the **mixed-lsp-mode** option. The tunnel preference in such an SDP is fixed and is controlled by the service manager. The configuration of the tunnel table preference parameter does not modify the behavior of such an SDP and the services that bind to it.

It is recommended to not set two or more tunnel types to the same preference value. In such a situation, the tunnel table prefers the tunnel type which was first introduced in SR OS implementation historically.

The **no** form of this command reverts to the default value.

Default

sr-te 8

Parameters

value

Specifies the tunnel table preference value for SR-TE LSP.

Values	1 to 255
Default	8

Platforms

7705 SAR Gen 2

sr-te

Syntax

[**no**] **sr-te**

Context

[Tree] (config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter sr-te)

Full Context

configure router static-route-entry indirect tunnel-next-hop resolution-filter sr-te

Description

The sr-te value instructs the code to search for the set of lowest metric SR-TE LSPs to the address of the indirect next-hop. The LSP metric is provided by MPLS in the tunnel table. The static route treats a set of SR-TE LSPs with the same lowest metric as an ECMP set. The user has the option of configuring a list of SR-TE LSP names to be used exclusively instead of searching in the tunnel table. In that case, all LSPs must have the same LSP metric in order for the static route to use them as an ECMP set. Otherwise, only the LSPs with the lowest common metric value are selected.

Default

no sr-te

Platforms

7705 SAR Gen 2

sr-te

Syntax

[**no**] **sr-te**

Context

[Tree] (config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter sr-te)

[Tree] (config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter sr-te)

Full Context

configure router bgp next-hop-resolution shortcut-tunnel family resolution-filter sr-te

configure router bgp next-hop-resolution labeled-routes transport-tunnel family resolution-filter sr-te

Description

This command selects the Segment Routing (SR) tunnel type programmed by a traffic engineered (TE) instance in TTM for next-hop resolution. In the case of multiple SR-TE tunnels with the same lowest metric, BGP selects the tunnel with the lowest tunnel ID.

Platforms

7705 SAR Gen 2

sr-te

Syntax

sr-te {**legacy** | **application-specific-link-attributes**}

no sr-te

Context

[Tree] (config>router>ospf>traffic-engineering-options sr-te)

Full Context

configure router ospf traffic-engineering-options sr-te

Description

This command configures the advertisement of TE attributes of each link on a per-application basis. Two applications are supported in SR OS: RSVP-TE and SR-TE. Although the **legacy** mode of advertising TE attributes is supported, additional configurations are possible.

The **no** form of this command deletes the context.

Default

no sr-te

Parameters

legacy

Advertises the TE attributes for MPLS-enabled SR links using TE Opaque LSAs.

Note:
Do not configure the **legacy** mode if the network has both RSVP-TE and SR-TE attributes and the links are not congruent.

application-specific-link-attributes

Advertises TE information for MPLS-enabled SR links using the new Application Specific Link Attributes (ASLA) TLVs.

Platforms

7705 SAR Gen 2

sr-te

Syntax

[**no**] **sr-te**

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls sr-te)

Full Context

configure oam-pm session ip tunnel mpls sr-te

Description

This command configures specification of SR-TE specific tunnel information that is used to transport the test packets. Entering this context removes all other tunnel type options configured under the **configure oam-pm session ip tunnel mpls** context. Only a single **mpls** type can be configured for an OAM-PM session.

The **no** form of this command removes the SR-TE LSP name from the configuration.

Default

no override

Parameters

tcp-port

Specifies the source TCP port to be used in the test TCP header.

Values 0 to 65535

Platforms

7705 SAR Gen 2

sr-te

Syntax

sr-te

Context

[Tree] (config>service>vprn>auto-bind-tunnel>res-filter sr-te)

Full Context

configure service vprn auto-bind-tunnel resolution-filter sr-te

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

28.53 sr-te-lsp

sr-te-lsp

Syntax

[**no**] **sr-te-lsp** *lsp-name*

Context

[Tree] (config>service>sdp sr-te-lsp)

Full Context

configure service sdp sr-te-lsp

Description

This command configures an MPLS SDP of LSP type SR-TE.

The user can specify up to 16 SR-TE LSP names. The destination address of all LSPs must match that of the SDP far-end option. Service data packets are sprayed over the set of LSPs in the SDP using the same procedures as for tunnel selection in ECMP. Each SR-TE LSP can, however, have up to 32 next-hops at the ingress LER when the first segment is a node SID-based SR tunnel. Thus, the service data packet is forwarded over one of a maximum of 16x32 next-hops.

The **tunnel-far-end** option is not supported. In addition, the **mixed-lsp-mode** option does not support the **sr-te** tunnel type.

The signaling protocol for the service labels for an SDP using a SR-TE LSP can be configured to static (**off**), T-LDP (**tldp**), or BGP (**bgp**).

Platforms

7705 SAR Gen 2

28.54 sr-te-resignal

sr-te-resignal

Syntax

sr-te-resignal

Context

[Tree] (config>router>mpls sr-te-resignal)

Full Context

configure router mpls sr-te-resignal

Description

Commands in this context configure the re-optimization parameters of SR-TE LSPs.

Platforms

7705 SAR Gen 2

28.55 src-access-list

src-access-list

Syntax

src-access-list *list-name*

no src-access-list *list-name*

Context

[Tree] (config>system>security>snmp src-access-list)

Full Context

configure system security snmp src-access-list

Description

This command configures a list of source IP addresses used to validate SNMPv1 and SNMPv2c requests after the list is associated with one or more SNMPv1 and SNMPv2c communities.

A source access list referenced by one or more **community** instances is used to verify the source IP addresses of an SNMP request using the community, regardless of the VPRN/VRF interface (or "Base" interface) on which the request arrived. For example, if an SNMP request arrives on an interface in VPRN

"100" but the request is referencing a community, the source IP address in the packet is validated against the source address list configured for the community. This occurs regardless of whether the request is destined to a VPRN interface address and the VPRN has SNMP access enabled, or the request is destined to the base system address via GRT leaking. If the source IP address of the request message does not match the IP address of any of the **src-host** entries contained in the list, the request is discarded and logged as an SNMP authentication failure.

Caution:
Using source access list validation can impact the time it takes for an SR OS node to reply to an SNMP request. Nokia recommends keeping the lists short by including only the addresses that are needed, and to place SNMP managers that send the highest volume of requests, such as the NSP NFM-P, at the top of the list.

A maximum of 16 source access lists can be configured. Each source access lists can contain a maximum of 16 source hosts.

The **no** form of this command removes the named source access list. Users cannot remove a source access list that is referenced by one or more **community** instances.

Parameters

list-name

Specifies the name or key of the source access list. This parameter must begin with a letter (a-z or A-Z).

Platforms

7705 SAR Gen 2

28.56 src-host

src-host

Syntax

src-host *host-name* **address** *ip-address*

no src-host *host-name*

Context

[Tree] (config>system>security>snmp>src-access-list src-host)

Full Context

configure system security snmp src-access-list src-host

Description

This command configures a source IP address entry used to validate SNMPv1 and SNMPv2c requests.

The **no** form of this command removes the specified entry.

Parameters

host-name

Specifies a name for the entry, up to 32 characters.

ip-address

Specifies an allowed IPv4 or IPv6 source address for SNMP requests.

Values

ipv4-address — a.b.c.d (host bits must be 0)

ipv6-address — x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x — 0 to FFFF (hexadecimal)

d — 0 to 255 (decimal)

Platforms

7705 SAR Gen 2

28.57 src-ip

src-ip

Syntax

src-ip *ip-address*

no src-ip

Context

[Tree] (config>redundancy>mc>peer>mcr>l3ring>node>cv src-ip)

Full Context

configure redundancy multi-chassis peer mc-ring l3-ring ring-node connectivity-verify src-ip

Description

This command specifies the source IP address used in ring-node connectivity verification of this ring node.

The **no** form of this command reverts to the default.

Parameters

ip-address

Specifies the source IP address used in ring-node connectivity verification of this ring node.

Platforms

7705 SAR Gen 2

src-ip

Syntax

src-ip {*ip-address/mask* | *ip-address* [*ipv4-address-mask*] | **ip-prefix-list** *prefix-list-name*}

no src-ip

Context

[Tree] (config>qos>sap-ingress>ip-criteria>entry>match src-ip)

[Tree] (config>qos>sap-egress>ip-criteria>entry>match src-ip)

Full Context

configure qos sap-ingress ip-criteria entry match src-ip

configure qos sap-egress ip-criteria entry match src-ip

Description

This command configures a source IPv4 address range to be used as an SAP QoS policy match criterion.

To match on the source IPv4 or IPv6 address, specify the address and its associated mask; for example, 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used for IPv4.

The **no** form of this command removes the source IPv4 or IPv6 address match criterion.

Default

no src-ip

Parameters

ip-address

Specifies the source IPv4 address specified in dotted decimal notation.

Values ip-address: a.b.c.d

mask

Specifies the length in bits of the subnet mask.

Values 1 to 32

ipv4-address-mask

Specifies the subnet mask in dotted decimal notation.

Values a.b.c.d (dotted quad equivalent of mask length)

prefix-list-name

Specifies the IPv4 prefix list name, a string of up to 32 printable ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

src-ip

Syntax

src-ip {*ipv6-address/prefix-length* | *ipv6-address ipv6-address-mask* | **ipv6-prefix-list** *ipv6-prefix-list-name*}

no src-ip

Context

[Tree] (config>qos>sap-ingress>ipv6-criteria>entry>match src-ip)

[Tree] (config>qos>sap-egress>ipv6-criteria>entry>match src-ip)

Full Context

configure qos sap-ingress ipv6-criteria entry match src-ip

configure qos sap-egress ipv6-criteria entry match src-ip

Description

This command configures a source IPv6 address range to be used as an SAP QoS policy match criterion.

To match on the source IPv6 address, specify the address and its associated mask, for example, 2001:db8:1000::/64.

The **no** form of this command removes the source IPv6 address match criterion.

Default

no src-ip

Parameters

ipv6-address

Specifies the IPv6 address for the IP match criterion in hexadecimal digits.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D

prefix-length

Specifies the IPv6 prefix length for the IPv6 address expressed as a decimal integer.

Values 1 to 128

ipv6-address-mask

Specifies the IPv6 address mask.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x: [0 to FFFF]H
d: [0 to 255]D

ipv6-prefix-list-name

Specifies the IPv6 prefix list name, a string of up to 32 printable ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

src-ip

Syntax

src-ip {*ip-address/mask* | *ip-address ipv4-address-mask* | **ip-prefix-list** *ip-prefix-list-name*}

src-ip {*ipv6-address/mask* | *ipv6-address ipv6-address-mask* | **ipv6-prefix-list** *ipv6-prefix-list-name*}

no src-ip

Context

[Tree] (config>qos>network>ingress>ipv6-criteria>entry>match src-ip)

[Tree] (config>qos>network>ingress>ip-criteria>entry>match src-ip)

[Tree] (config>qos>network>egress>ip-criteria>entry>match src-ip)

[Tree] (config>qos>network>egress>ipv6-criteria>entry>match src-ip)

Full Context

configure qos network ingress ipv6-criteria entry match src-ip

configure qos network ingress ip-criteria entry match src-ip

configure qos network egress ip-criteria entry match src-ip

configure qos network egress ipv6-criteria entry match src-ip

Description

This command configures a source IPv4 or IPv6 address range to be used as a network QoS policy match criterion.

To match on the source IPv4 or IPv6 address, specify the address and its associated mask, for example, when specifying an IPv4 address, 10.1.0.0/16 or 10.1.0.0 255.255.0.0 can be used.

The **no** form of this command removes the source IPv4 or IPv6 address match criterion.

Parameters

ip-address

Specifies the source IPv4 address specified in dotted decimal notation.

Values ip-address: a.b.c.d

mask

Specifies the length in bits of the subnet mask.

Values 1 to 32

ipv4-address-mask

Specifies the subnet mask in dotted decimal notation.

Values a.b.c.d (dotted quad equivalent of mask length)

ip-prefix-list-name

Specifies an IPv4 prefix list which contains IPv4 address prefixes to be matched.

Values A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

ipv6-address

Specifies the IPv6 prefix for the IP match criterion in hex digits.

Values

ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

mask

Specifies the length of the ipv6-address expressed as a decimal integer.

Values 1 to 128

ipv6-address-mask

Specifies the eight 16-bit hexadecimal pieces representing bit match criteria.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)

ipv6-prefix-list-name

Specifies an IPv6 prefix list which contains IPv6 address prefixes to be matched.

Values A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

src-ip

Syntax

IPv4:

src-ip {*ip-address/mask* | *ip-address ipv4-address-mask* | **ip-prefix-list** *prefix-list-name*}

IPv6:

src-ip {*ipv6-address/prefix-length* | *ipv6-address ipv6-address-mask* | **ipv6-prefix-list** *prefix-list-name*}

no src-ip

Context

[Tree] (config>filter>ipv6-exception>entry>match src-ip)

[Tree] (config>filter>ipv6-filter>entry>match src-ip)

[Tree] (config>filter>ip-exception>entry>match src-ip)

Full Context

configure filter ipv6-exception entry match src-ip

configure filter ipv6-filter entry match src-ip

configure filter ip-exception entry match src-ip

Description

This command configures a source IPv4 or IPv6 address range to be used as an IP filter or IP exception match criterion.

To match on the source IPv4 or IPv6 address, specify the address and its associated mask, for example, 10.1.0.0/16 for IPv4. The conventional notation of 10.1.0.0 255.255.0.0 may also be used for IPv4.

The **no** form of the command removes the source IP address match criterion.

Default

no src-ip

Parameters

ip-address

Specifies the destination IPv4 address specified in dotted decimal notation.

Values a.b.c.d

mask

Specifies the length in bits of the subnet mask.

Values 1 to 32

ipv4-address-mask

Specifies the subnet mask in dotted decimal notation.

Values a.b.c.d (dotted quad equivalent of mask length)

ip-prefix-listoripv6-prefix-list prefix-list-name

Specifies to use a list of IP prefixes, which is a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

ipv6-address

Specifies an IPv6 prefix for the IP match criterion in hex digits.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0..FFFF]H

d: [0..255]D

prefix-length

Specifies whether a the IPv6 prefix length for the specified *ipv6-address* expressed as a decimal integer.

Values 1 to 128

ipv6-address-mask

Specifies eight 16-bit hexadecimal pieces representing bit match criteria.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x: [0..FFFF]H

d: [0..255]D

Platforms

7705 SAR Gen 2

src-ip

Syntax

src-ip *ip-prefix*[/*mask*] [*netmask*]

src-ip ip-prefix-list *ip-prefix-list-name*

no src-ip

Context

[Tree] (config>system>security>mgmt-access-filter>ip-filter>entry src-ip)

Full Context

configure system security management-access-filter ip-filter entry src-ip

Description

This command configures a source IP address range or an IP prefix list to be used as a management access filter match criterion.

The **no** form of this command removes the source IP address match criterion.

Default

no src-ip

Parameters

ip-prefix

Specifies the IP prefix for the IP match criterion in dotted decimal notation.

mask

Specifies the subnet mask length expressed as a decimal integer.

Values 1 **to** 32 (mask length), 0.0.0.0 **to** 255.255.255.255 (dotted decimal)

netmask

Specifies the dotted quad equivalent of the mask length.

Values 0.0.0.0 to 255.255.255.255

ip-prefix-list-name

Specifies the IP prefix list used as a match criterion for the source IP address. It is a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes .

Platforms

7705 SAR Gen 2

src-ip

Syntax

src-ip *ipv6-address/prefix-length*

src-ip ipv6-prefix-list *ipv6-prefix-list-name*

no src-ip

Context

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter>entry src-ip)

Full Context

configure system security management-access-filter ipv6-filter entry src-ip

Description

This command configures a source IPv6 address range or an IPv6 prefix list to be used as a management access filter match criterion.

The **no** form of this command removes the source IPv6 address match criterion.

Default

no src-ip

Parameters

ipv6-address/prefix-length

Specifies the IPv6 address for the IPv6 match criterion in dotted decimal notation. An IPv6 IP address is written as eight 4-digit (16-bit) hexadecimal numbers separated by colons. One string of zeros per address can be left out, so that 2001:db8::0:217A is the same as 2001:db8:0:0:0:0:0:217A.

Values		
	ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x: [0..FFFF]H
		d: [0..255]D
	prefix-length	1 to 128

ipv6-prefix-list-name

Specifies the IPv6 prefix list used a match criterion for the source IP address. It is a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes .

Platforms

7705 SAR Gen 2

28.58 src-ip-address

src-ip-address

Syntax

src-ip-address *ip-address*

no src-ip-address

Context

[Tree] (config>saa>test>type-multi-line>lsp-ping>sr-policy src-ip-address)

[Tree] (config>saa>test>type-multi-line>lsp-ping src-ip-address)

Full Context

configure saa test type-multi-line lsp-ping sr-policy src-ip-address

configure saa test type-multi-line lsp-ping src-ip-address

Description

This command configures the source IP address. This option is used when an OAM packet must be generated from a different address than the node's system interface address. For example, when the OAM packet is sent over an LDP LSP and the LDP LSR-ID of the corresponding LDP session to the next hop is set to an address other than the system interface address.

The **no** form of this command removes the configuration.

Parameters

ip-address

Specifies the source IP address.

Values ipv4-address: a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]D

Platforms

7705 SAR Gen 2

28.59 src-mac

src-mac

Syntax

src-mac *ieee-address*

no src-mac

Context

[Tree] (config>redundancy>mc>peer>mcr>l3ring>node>cv src-mac)

Full Context

configure redundancy multi-chassis peer mc-ring l3-ring ring-node connectivity-verify src-mac

Description

This command specifies the source MAC address used for the ring-node connectivity verification of this ring node.

If all zeros are specified, the MAC address of the system management processor (CPM) is used.

The **no** form of this command reverts to the default.

Parameters

ieee-address

Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

Platforms

7705 SAR Gen 2

src-mac

Syntax

src-mac *ieee-address* [*ieee-address-mask*]

no src-mac

Context

[Tree] (config>qos>sap-ingress>mac-criteria>entry>match src-mac)

Full Context

configure qos sap-ingress mac-criteria entry match src-mac

Description

This command configures a source MAC address or range to be used as a service ingress QoS policy match criterion.

The **no** form of this command removes the source mac as the match criteria.

Default

no src-mac

Parameters

ieee-address

Enter the 48-bit IEEE MAC address to be used as a match criterion.

Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

ieee-address-mask

This 48-bit mask can be configured using the following formats:

Table 93: Format Styles to Configure Mask

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHH	0x0FFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

To configure all packets with a source MAC OUI value of 00-03-FA to be subject to a match condition, the entry should be specified as: 003FA000000 0xFFFFFF000000

Values 0x00000000000000 to 0xFFFFFFFFFFFF (hex)

Default 0xFFFFFFFFFFFF (hex) (exact match)

Platforms

7705 SAR Gen 2

src-mac

Syntax

src-mac *ieee-address* [*ieee-address-mask*]

no src-mac

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry>match src-mac)

Full Context

configure system security management-access-filter mac-filter entry match src-mac

Description

This command configures a source MAC address or range to be used as a MAC filter match criterion.

The **no** form of this command removes the source mac as the match criteria.

Default

no src-mac

Parameters

ieee-address

Specifies the 48-bit IEEE mac address to be used as a match criterion.

Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

ieee-address-mask

Specifies a 48-bit mask that can be configured using the formats listed in Table 94: ieee-address-mask Formats:

Table 94: ieee-address-mask Formats

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHH	0x0FFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

To configure so that all packets with a source MAC OUI value of 00-03-FA are subject to a match condition then the entry should be specified as: 003FA000000 0xFFFFFF000000

Default 0xFFFFFFFFFFFF (exact match)

Values 0x00000000000000 to 0xFFFFFFFFFFFF

Platforms

7705 SAR Gen 2

28.60 src-port

src-port

Syntax

src-port {**lt** | **gt** | **eq**} *src-port-number*

src-port range *start end*

no src-port

Context

[Tree] (config>qos>sap-ingress>ipv6-criteria>entry>match src-port)

[Tree] (config>qos>sap-egress>ipv6-criteria>entry>match src-port)

[Tree] (config>qos>sap-egress>ip-criteria>entry>match src-port)

[Tree] (config>qos>sap-ingress>ip-criteria>entry>match src-port)

Full Context

configure qos sap-ingress ipv6-criteria entry match src-port

configure qos sap-egress ipv6-criteria entry match src-port

configure qos sap-egress ip-criteria entry match src-port

configure qos sap-ingress ip-criteria entry match src-port

Description

This command configures a source TCP or UDP port number or port range for a SAP QoS policy match criterion.

The **no** form of this command removes the source port match criterion.

Default

no src-port

Parameters

{lt | gt | eq} *src-port-number*

The TCP or UDP port numbers to match, specified as less than (**lt**), greater than (**gt**), or equal to (**eq**) to the source port value, specified as a decimal integer.

Values 1 to 65535 (decimal)

range *startend*

The range of TCP or UDP port values to match, specified as between the *start* and *end* source port values inclusive.

Values 1 to 65535 (decimal)

Platforms

7705 SAR Gen 2

src-port

Syntax

src-port {**lt** | **gt** | **eq**} *src-port-number*

src-port port-list *port-list-name*

src-port range *start end*

no src-port

Context

[Tree] (config>qos>network>ingress>ip-criteria>entry>match src-port)

[Tree] (config>qos>network>egress>ipv6-criteria>entry>match src-port)

[Tree] (config>qos>network>egress>ip-criteria>entry>match src-port)

[Tree] (config>qos>network>ingress>ipv6-criteria>entry>match src-port)

Full Context

configure qos network ingress ip-criteria entry match src-port

configure qos network egress ipv6-criteria entry match src-port

configure qos network egress ip-criteria entry match src-port

configure qos network ingress ipv6-criteria entry match src-port

Description

This command configures a source TCP or UDP port number, port range, or a port list for a network QoS policy match criterion.

The **no** form of this command removes the source port match criterion.

Default

no src-port

Parameters

lt

Keyword used to specify TCP or UDP port numbers to match that are less than the source port value.

gt

Keyword used to specify TCP or UDP port numbers to match that are greater than the source port value.

eq

Keyword used to specify TCP or UDP port numbers to match that are equal to the source port value.

src-port-number

The source port value, specified as a decimal integer.

Values 1 to 65535

port-list-name

Specifies a port list name, up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

start

Specifies the starting range of TCP or UDP source port values to match.

Values 1 to 65535

end

Specifies the end range of TCP or UDP source port values to match.

Values 1 to 65535

Platforms

7705 SAR Gen 2

src-port

Syntax

src-port {**lt** | **gt** | **eq**} *src-port-number*

src-port port-list *port-list-name*

src-port range *src-port-number src-port-number*

no src-port

Context

[Tree] (config>filter>ip-exception>entry>match src-port)

[Tree] (config>filter>ip-filter>entry>match src-port)

[Tree] (config>filter>ipv6-exception>entry>match src-port)

[Tree] (config>filter>ipv6-filter>entry>match src-port)

Full Context

configure filter ip-exception entry match src-port

configure filter ip-filter entry match src-port

configure filter ipv6-exception entry match src-port

configure filter ipv6-filter entry match src-port

Description

This command configures a source TCP, UDP, or SCTP port number, port range, or port match list for an IP filter or IP exception match criterion. An entry containing Layer 4 non-zero match criteria will not match non-initial (2nd, 3rd, and so on) fragments of a fragmented packet since only the first fragment contains the Layer 4 information. Similarly an entry containing "**src-port eq** 0" match criterion, may match non-initial fragments when the source port value is not present in a packet fragment and other match criteria are also met.

The **no** form of the command removes the source port match criterion.

Default

no src-port

Parameters

lt | gt | eq

Specifies the operator to use relative to *src-port-number* for specifying the port number match criteria.

lt specifies that all port numbers less than *src-port-number* match.

gt specifies that all port numbers greater than *src-port-number* match.

eq specifies that *src-port-number* must be an exact match.

src-port-number

Specifies the source port number to be used as a match criteria expressed as a decimal integer, and in hexadecimal or binary format. Below shows decimal integer only.

Values 0 to 65535

port-list-name

Specifies to use a list of ports referred to by *port-list-name*, which is a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.

src-port-number src-port-number

Specifies inclusive port range between two src-port-number values.

Platforms

7705 SAR Gen 2

src-port

Syntax

src-port {*port-id* | **cpm** | **lag** *lag-id*}

no src-port

Context

[Tree] (config>system>security>mgmt-access-filter>ipv6-filter>entry src-port)

[Tree] (config>system>security>mgmt-access-filter>ip-filter>entry src-port)

Full Context

configure system security management-access-filter ipv6-filter entry src-port

configure system security management-access-filter ip-filter entry src-port

Description

This command restricts ingress management traffic to either the CPM/CCM Ethernet port or any other logical port (for example LAG) on the device.

When the source interface is configured, only management traffic arriving on those ports satisfy the match criteria.

The **no** form of this command reverts to the default value.

Default

no src-port

Parameters

port-id

Specifies the port ID in formats shown below.

Values

slot/*mda*/*port*[.*channel*]		
	aps	keyword
	group-id	1 to 128
ccag-id	**ccag**-*id*. *path-id*[*cc-type*]	
	ccag	keyword
	id	1 to 8
	path-id	**a**, **b**
	cc-type	**.sap-net**, **.net-sap**

cpm

Matches any traffic received on any Ethernet port.

lag-id

Specifies the LAG identifier.

Values 1 to 800

Platforms

7705 SAR Gen 2

28.61 src-route-option

src-route-option

Syntax

src-route-option {**true** | **false**}

no source-route-option

Context

[Tree] (config>filter>ip-filter>entry>match src-route-option)

Full Context

configure filter ip-filter entry match src-route-option

Description

This command enables source route option match conditions. When enabled, this filter should match if a (strict or loose) source route option is present/not present at any location within the IP header, as per the value of this object. The **no** form of the command removes the criterion from the match entry.

Default

no src-route-option

Parameters

true

Enables source route option match conditions.

false

Disables source route option match conditions.

Platforms

7705 SAR Gen 2

28.62 srefresh

srefresh

Syntax

srefresh [**detail**]

no srefresh

Context

[Tree] (debug>router>rsvp>packet srefresh)

Full Context

debug router rsvp packet srefresh

Description

This command debugs srefresh packets.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about srefresh packets.

Platforms

7705 SAR Gen 2

28.63 srlb

srlb

Syntax

srlb *reserved-label-block-name*

no srlb

Context

[Tree] (config>router>ospf>segm-rtng srlb)

[Tree] (config>router>isis>segm-rtng srlb)

Full Context

configure router ospf segment-routing srlb

configure router isis segment-routing srlb

Description

This command specifies the reserved label block to use for the Segment Routing Local Block (SRLB) for the specified IS-IS or OSPF instance. The named reserved label block must already have been configured under **config**>**router**>**mpls**>**mpls-labels**.

The **no** form of this command removes an SRLB.

Parameters

reserved-label-block-name

Specifies the name of the reserved label block, up to 64 characters.

Platforms

7705 SAR Gen 2

28.64 srlg

srlg

Syntax

[**no**] **srlg**

Context

[Tree] (config>router>mpls>lsp>secondary srlg)

Full Context

configure router mpls lsp secondary srlg

Description

This command enables the use of the SRLG constraint in the computation of a secondary path for an LSP at the head-end LER. The command is configurable for both RSVP-TE and SR-TE LSPs.

When SRLG is enabled, CSPF includes the SRLG constraint in the computation of the secondary LSP path if **path-computation-method local-cspf** is configured on the LSP. CSPF returns the list of SRLG groups along with the ERO during primary path CSPF computation. At a subsequent establishment of a secondary path with the SRLG constraint, the MPLS task again queries CSPF by providing the list of SRLG group numbers to be avoided. CSPF prunes all links with interfaces that belong to the same SRLGs as the interfaces included in the ERO of the primary path. If CSPF finds a path, the secondary path is set up. If a path is not found, MPLS keeps retrying the requests to CSPF.

An SRLG enabled secondary or standby path of the LSP configured with a value of the **path-computation-method** command other than **local-cspf** remains operationally down with a failure code of srlgPrimaryCspfDisabled(25).

When an LSP is administratively enabled, the SRLG-enabled secondary path is not tried if the first attempt to bring up the primary path is in progress. The SRLG enabled secondary path is kept down temporarily with failure code srlgPrimaryPathDown(26). After this first attempt, MPLS begins setting up the SRLG-enabled standby paths. If primary path computation fails or primary path was not configured, MPLS requests CSPF to compute the secondary path using an empty primary SRLG list. The SRLG *disjoint* state field shows *True* in this scenario.

If the primary path is re-optimized, has undergone MBB, or has come back up after being down, the MPLS task check determines if any SRLG secondary paths should be re-signaled. If MPLS finds that a secondary path is no longer SRLG disjointed, and therefore becomes ineligible, MPLS puts it on a delayed MBB immediately after the expiry of the retry timer. If MBB fails at the first try, the secondary path is torn down and the path is put on retry if not active. If the secondary path is active, then it is only torn down and resignaled when the primary path is activated. The secondary path can remain active even when ineligible while the revert timer to activate the primary path is still running.

If the primary goes down while active, the LSP uses the path of an eligible SRLG secondary path if it is up. If all secondary eligible SRLG paths are down, MPLS uses a non-SRLG secondary path, if configured and up. While the LSP is using a non-SRLG secondary path, if an eligible SRLG secondary path comes back up, MPLS switches the path of the LSP to the eligible SRLG secondary path. As soon as a path for the primary is successfully computed by CSPF, MPLS schedules the delay retry MBB for the secondary path using the new SRLG list.

If the primary path goes down while inactive, for example it is waiting for the revert timer to expire, MPLS resets the SRLG list of the primary to empty and changes the state of all secondary paths, including the currently active one, to the Disjointed state. A delay retry MBB is still performed but results in no change to the active secondary path.

A secondary path that becomes ineligible as a result of an update to the SRLG membership list of the primary path has the ineligibility status removed on any of the following events:

- a successful delay retry MBB of the secondary SRLG path that makes it eligible again
- the secondary path goes down. MPLS puts the standby on retry at the expiry of the retry timer. If successful, it becomes eligible. If not successful after the retry-timer expires or the number of retries reached the number configured under the **retry-limit** parameter, it is left down.

Once the primary path of the LSP is set up and is operationally up, any subsequent changes to the SRLG group membership of an interface that the primary path is using is not considered until the next opportunity the primary path is re-signaled. The primary path may be re-signaled due to a failure or to a make-before-break operation. Make-before-break occurs as a result of a global revertive operation, a timer based or manual re-optimization of the LSP path, or an operator change to any of the path constraints.

Once an SRLG secondary path is set up and is operationally up, any subsequent changes to the SRLG group membership of an interface the secondary path is using is not considered until the next opportunity when the secondary path is re-signaled. The secondary path is re-signaled due to a failure, to a re-signaling of the primary path, or to a make before break operation. Make-before-break occurs as a result of a timer based or manual re-optimization of the secondary path, or an operator change to any of the path constraints of the secondary path, except for enabling or disabling the **srlg** command itself. Enabling or disabling the **srlg** command on an active secondary or on an active or inactive secondary standby path causes the path to be torn down and re-signaled.

In addition, the user-configured **include** or **exclude** admin group statements for a secondary path are also checked together with the SRLG constraints by CSPF.

The following behavior of the feature is specific to the SR-TE LSP.

- An SRLG-enabled SR-TE LSP secondary path with SID label hops remains operational with failure code srlgPathWithSidHops(59).
- An SR-TE LSP uses IGP advertised link SRLG information in the TE database. It does not support the use of SRLG information in the static user SRLG database (**configure router mpls srlg-database**).
- Delay Retry MBB for making a non-disjointed path a disjointed one is not supported with an SR-TE LSP. Instead, the system performs a break-before-make (that is, teardown and retry) operation. If a non-disjointed path is the active path of the LSP, that path is torn down and retried after the router switches to another path (for example, after **revert-timer** expires). If the non-disjointed path is not an active path, it is torn down and retried immediately.

The **no** form of this command reverts to the default value.

Default

no srlg

Platforms

7705 SAR Gen 2

28.65 srlg-database

srlg-database

Syntax

[**no**] **srlg-database**

Context

[Tree] (config>router>mpls srlg-database)

Full Context

configure router mpls srlg-database

Description

Commands in this context configure the link members of SRLG groups for the entire network at any node that needs to signal LSP paths (for example, a head-end node).

The **no** form of this command deletes the entire SRLG database. CSPF assumes all interfaces have no SRLG membership association if the database was not disabled with the command **config>router>mpls>user-srlg-db disable**.

Platforms

7705 SAR Gen 2

28.66 srlg-enable

srlg-enable

Syntax

[**no**] **srlg-enable**

Context

[Tree] (config>router>route-next-hop-policy>template srlg-enable)

Full Context

configure router route-next-hop-policy template srlg-enable

Description

This command configures the SRLG constraint into the route next-hop policy template.

When this command is applied to a prefix, the LFA SPF will attempt to select an LFA next-hop, among the computed ones, which uses an outgoing interface that does not participate in any of the SLRGs of the outgoing interface used by the primary next-hop.

The SRLG criterion is applied before running the LFA next-hop selection algorithm.

The **no** form deletes the SRLG constraint from the route next-hop policy template.

Default

no srlg-enable

Platforms

7705 SAR Gen 2

28.67 srlg-frr

srlg-frr

Syntax

srlg-frr [**strict**]

no srlg-frr

Context

[Tree] (config>router>mpls srlg-frr)

Full Context

configure router mpls srlg-frr

Description

This command enables the use of the SRLG constraint in the computation of FRR bypass or detour to be associated with any primary LSP path on this system.

When this option is enabled, CSPF includes the SRLG constraint in the computation of a FRR detour or bypass for protecting the primary LSP path.

CSPF prunes all links with interfaces that belong to the same SRLG as the interface that is being protected, that is, the outgoing interface at the PLR the primary path is using. If one or more paths are found, the MPLS task will select one based on best cost and will signal the bypass/detour. If not found and the user has included the **strict** option, the bypass/detour is not setup and the MPLS task will keep retrying the request to CSPF. Otherwise, if a path exists that meets the other TE constraints, other than the SRLG one, the bypass/detour is setup.

A bypass or a detour LSP path is not intended to be SRLG disjoint from the entire primary path. Only the SRLGs of the outgoing interface at the PLR that the primary path is using are avoided.

When the MPLS task is searching for an SRLG bypass tunnel to associate with the primary path of the protected LSP, it will first check if any configured manual bypass LSP with CSPF enabled satisfies the SRLG constraints. The search skips any non-CSPF manual bypass LSP because there is no ERO returned to check the SRLG constraint. If no path is found, the task will check if an existing dynamic bypass LSP satisfies the SRLG and other primary path constraints. If not found, it will make a request to CSPF.

Once the primary path of the LSP is configured and is operationally up, subsequent changes to the SRLG group membership of an interface the primary path is using are not considered by the MPLS task at the PLR for bypass/detour association until the next opportunity the bypass LSP path or the primary path is resignaled. The path may be resignaled due to a failure or a Make-Before-Break (MBB) operation. MBB occurs as a result of a global revertive operation, a timer based or manual re-optimization of the bypass LSP or LSP primary path, or a user update of the primary path constraints.

Once the bypass or detour path is set up and is operationally up, subsequent changes to the SRLG group membership of an interface the bypass/detour path is using are not considered by the MPLS task at the PLR until the next opportunity when the association with the primary LSP path is rechecked. The

association is rechecked if the bypass path is re-optimized using the timer or manual resignal MBB. Detour paths cannot be re-optimized separately from the primary path.

Enabling or disabling **srlg-frr** command only takes effect when the LSP primary path or the bypass path is resignaled. The user can either wait for the resignal timer to expire or cause the paths to be resignaled immediately by executing, at the ingress LER, the manual resignal command for the LSP primary path or for the bypass LSP path.

A MPLS interface can belong to a maximum of 64 SRLG groups. The SRLG groups are configured using the **config>router>if-attribute>srlg-group** command. The SRLG groups that an RSVP interface belong to are configured using the **srlg-group** command in the **config>router>mpls>interface** context.

The **no** form of this command reverts to the default value.

Default

no srlg-frr

Parameters

strict

Specifies the name of the SRLG group within a virtual router instance.

Values no srlg-frr (default) srlg-frr (non-strict) srlg-frr **strict** (strict)

Platforms

7705 SAR Gen 2

28.68 srlg-group

srlg-group

Syntax

[**no**] **srlg-group** *group-name* [*group-name*]

no srlg-group

Context

[Tree] (config>service>ies>if>if-attribute srlg-group)

[Tree] (config>router>mpls>if srlg-group)

[Tree] (config>router>if>if-attribute srlg-group)

[Tree] (config>service>vprn>if>if-attribute srlg-group)

Full Context

configure service ies interface if-attribute srlg-group

configure router mpls interface srlg-group

configure router interface if-attribute srlg-group

configure service vprn interface if-attribute srlg-group

Description

This command configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface.

An interface can belong to up to 64 SRLG groups. However, each single operation of the **srlg-group** command allows a maximum of five (5) groups to be specified at a time. Once an SRLG group is bound to one or more interface, its value cannot be changed until all bindings are removed.

The configured SRLG membership is applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

Only the SRLGs bound to an MPLS interface are advertised area-wide in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The **no** form of this command deletes one or more of the SRLG memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.

Parameters

group-name

Specifies the name of the group, up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain. Each single operation of the **srlg-group** command allows a maximum of 5 groups to be specified at a time.

Platforms

7705 SAR Gen 2

srlg-group

Syntax

srlg-group *group-name* **value** *group-value* [**penalty-weight** *penalty-weight*]

no srlg-group *group-name*

Context

[Tree] (config>router>if-attribute srlg-group)

Full Context

configure router if-attribute srlg-group

Description

This command defines a Shared Risk Link Group (SRLG) which can be associated with an IP or MPLS interface.

SRLG is used to tag IP or MPLS interfaces which share a specific fate with the same identifier. For example, an SRLG group identifier could represent all links which use separate fibers but are carried in the

same fiber conduit. If the conduit is accidentally cut, all the fiber links are cut which means all interfaces using these fiber links will fail.

The user first configures locally on each router the name and identifier of each SRLG group. A maximum of 1024 SRLGs can be configured per system.

The user then configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface. A maximum of 64 SRLGs can be applied to a given interface.

When SRLGs are applied to MPLS interfaces, CSPF at an LER will exclude the SRLGs of interfaces used by the LSP primary path when computing the path of the secondary path. CSPF at an LER or LSR will also exclude the SRLGs of the outgoing interface of the primary LSP path in the computation of the path of the FRR backup LSP. This provides path disjointness between the primary path and the secondary path or FRR backup path of an LSP.

When SRLGs applied to IES, VPRN, or network IP interfaces, they are evaluated in the route next-hop selection by adding the **srlg-enable** option in a route next-hop policy template applied to an interface or a set of prefixes. For instance, the user can enable the SRLG constraint to select a LFA next-hop for a prefix which avoids all interfaces that share fate with the primary next-hop.

The following provisioning rules are applied to SRLG configuration. The system will reject the creation of a SRLG if it re-uses the same name but with a different group value than an existing group. The system will also reject the creation of an SRLG if it re-uses the same group value but with a different name than an existing group.

Only the SRLGs bound to an MPLS interface are advertised area-wide in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

A user may specify a penalty weight (**penalty-weight**) associated with an SRLG. This controls the likelihood of paths with links sharing SRLG values with a primary path being used by a bypass or detour LSP. The higher the penalty weight, the less desirable it is to use the link with a given SRLG.

Parameters

group-name

Specifies the name of the group, up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

group-value

Specifies the integer value associated with the group. The association of group name and value should be unique within an IP/MPLS domain.

Values 0 to 4294967295

penalty-weight

Specifies the integer value of the penalty weight that is assigned to the SRLG group

Values 0 to 65535

Default 0

Platforms

7705 SAR Gen 2

28.69 ssap

ssap

Syntax

ssap *ssap-value* [*ssap-mask*]

no ssap

Context

[Tree] (config>qos>sap-ingress>mac-criteria>entry>match ssap)

Full Context

configure qos sap-ingress mac-criteria entry match ssap

Description

This command configures an Ethernet 802.2 LLC SSAP value or range for an ingress SAP QoS policy match criterion.

This is a 1-byte field that is part of the 802.2 LLC header of the IEEE 802.3 Ethernet Frame.

The snap-pid field, etype field, ssap, and dsap fields are mutually exclusive and cannot be part of the same match criteria.

The **no** form of this command removes the ssap match criterion.

Default

no ssap

Parameters

ssap-value

The 8-bit ssap match criteria value in hex.

Values 0x00 to 0xFF (hex)

ssap-mask

This is optional and can be used when specifying a range of ssap values to use as the match criteria.

This 8-bit mask can be configured using the following formats.

Table 95: Format Styles to Configure Mask

Format Style	Format Syntax	Example
Decimal	DDD	240
Hexadecimal	0xHH	0xF0

Format Style	Format Syntax	Example
Binary	0bBBBBBBBB	0b11110000

Values 0x00 to 0xFF

Platforms

7705 SAR Gen 2

ssap

Syntax

ssap *ssap-value* [*ssap-mask*]

no ssap

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry>match ssap)

Full Context

configure system security management-access-filter mac-filter entry match ssap

Description

This command configures an Ethernet 802.2 LLC SSAP value or range for a MAC filter match criterion.

This is a one-byte field that is part of the 802.2 LLC header of the IEEE 802.3 Ethernet Frame.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and may not be part of the same match criteria. Refer to the *7705 SAR Gen 2 Router Configuration Guide* for information about MAC Match Criteria Exclusivity Rules fields that are exclusive based on the frame format.

The **no** form of this command removes the SSAP match criterion.

Default

no ssap

Parameters

ssap-value

Specifies the 8-bit SSAP match criteria value in hex.

Values 0x00 to 0xFF

ssap-mask

Specifies a range of SSAP values to use as the match criteria.

Platforms

7705 SAR Gen 2

28.70 ssh

ssh

Syntax

ssh *host* [**-l** *username*] [**-v** *ssh-version*] [{**router** *router-instance* | **service-name** *service-name*}] [**re-exchange-min** *minutes*] [**re-exchange-mbyte** *megabytes*] [**-i** *private-key-filename*] [**-p** *port*]

Context

[Tree] (ssh)

Full Context

ssh

Description

This command initiates a client SSH session with the remote host and is independent from the administrative or operational state of the SSH server. However, to be the target of an SSH session, the SSH server must be operational. This command also allows the user to initiate an SSH session, with a key reexchange, based on maximum megabytes or minutes, whichever occurs first. If the reexchange options are not set, the default behavior does not perform a key reexchange.

Quitting SSH while in the process of authentication is accomplished by either executing a ctrl-c or "~." (tilde and dot), assuming the "~" is the default escape character for the SSH session.

Parameters

host

Specifies the remote host for the SSH session.

Values

host: *user@hostname* - [up to 255 characters]	
user	up to 32 characters
hostname	[*dns-name* | *ipv4-address* | *ipv6-address*]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D
	interface: up to 32 characters, mandatory for link local addresses

dns-name up to128 characters

username

Specifies the user name to use when opening the SSH session, up to 32 characters.

router-instance

Specifies the router name or service ID.

Values *router-instance*: *router-name* or *vprn-svc-id*

router-name "Base", "management", "vpls-management"

vprn-svc-id 1 to 2147483647

Default Base

service-name

Specifies the service name, up to 64 characters.

minutes

Specifies the time interval after which the SSH client will initiate the key reexchange.

Values 1 to 1440 minutes

megabytes

Specifies the number of megabytes, on a SSH session, after which the SSH client will initiate the key reexchange.

Values 1 to 64000 MB

private-key-filename

Specifies the name of the file containing the private key for public-key authentication on the SR OS SSH client, up to 255 characters. (The public key must be provided to the SSH server.) When using the *private-key-filename* option, if the file containing the private key is encrypted, the system asks for the password to decrypt the file.

port

Specifies the listening port for the SR OS SSH client to establish the SSH session with the SSH server.

Values 1 to 65535

Platforms

7705 SAR Gen 2

ssh

Syntax

ssh

Context

[Tree] (config>system>login-control ssh)

[Tree] (config>system>security ssh)

Full Context

configure system login-control ssh

configure system security ssh

Description

Commands in this context configure the SSH parameters.

Platforms

7705 SAR Gen 2

28.71 ssh-authentication-method

ssh-authentication-method

Syntax

ssh-authentication-method

Context

[Tree] (config>system>security>user ssh-authentication-method)

Full Context

configure system security user ssh-authentication-method

Description

Commands in this context configure, at the user level, the authentication method accepted by the SSH server. The user-level configuration overrides the system-level configuration.

Platforms

7705 SAR Gen 2

28.72 ssh-max-sessions

ssh-max-sessions

Syntax

ssh-max-sessions *number-of-sessions*

no ssh-max-sessions

Context

[Tree] (config>system>security>cli-session-group ssh-max-sessions)

[Tree] (config>system>security>profile ssh-max-sessions)

Full Context

configure system security cli-session-group ssh-max-sessions

configure system security profile ssh-max-sessions

Description

This command is used to limit the number of SSH-based sessions available to all users that are part of a particular profile, or to all users of all profiles that are part of the same **cli-session-group**.

The **no** form of this command disables the command and the profile or group limit is not applied on the number of sessions.

Default

no ssh-max-sessions

Parameters

number-of-sessions

Specifies the maximum number of allowed SSH-based sessions.

Values 0 to 50

Platforms

7705 SAR Gen 2

28.73 ssh-reply

ssh-reply

Syntax

[**no**] **ssh-reply**

Context

[Tree] (config>service>ies>if>vrrp ssh-reply)

Full Context

configure service ies interface vrrp ssh-reply

Description

This command enables the non-owner master to reply to SSH Requests directed at the virtual router instances IP addresses. The SSH request can be received on any routed interface. SSH must not have been disabled at the management security level (either on the parental IP interface or based on the SSH source host address). Proper login and CLI command authentication is still enforced.

When ssh-reply is not enabled, SSH packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to SSH regardless of the ssh-reply configuration.

The ssh-reply command is only available in non-owner vrrp virtual-router-id nodal context. If the ssh-reply command is not executed, SSH packets to the virtual router instance IP addresses is silently discarded.

The **no** form of this command restores the default operation of discarding all SSH packets destined to the non-owner virtual router instance IP addresses.

Default

no ssh-reply

Platforms

7705 SAR Gen 2

ssh-reply

Syntax

[**no**] **ssh-reply**

Context

[Tree] (config>service>vprn>if>vrrp ssh-reply)

Full Context

configure service vprn interface vrrp ssh-reply

Description

This command enables the non-owner master to reply to SSH Requests directed at the virtual router instance's IP addresses. The SSH request can be received on any routed interface. SSH must not have been disabled at the management security level (either on the parental IP interface or based on the SSH source host address). Proper login and CLI command authentication is still enforced.

When ssh-reply is not enabled, SSH packets to non-owner master virtual IP addresses are silently discarded. Non-owner backup virtual routers never respond to SSH regardless of the ssh-reply configuration.

The ssh-reply command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the ssh-reply command is not executed, SSH packets to the virtual router instance IP addresses is silently discarded.

The **no** form of this command restores the default operation of discarding all SSH packets destined to the non-owner virtual router instance IP addresses.

Default

no ssh-reply

Platforms

7705 SAR Gen 2

ssh-reply

Syntax

[**no**] **ssh-reply**

Context

[Tree] (config>router>if>vrrp ssh-reply)

Full Context

configure router interface vrrp ssh-reply

Description

This command enables the non-owner master to reply to SSH requests directed at the virtual router instance IP addresses. This command is only applicable to IPv4.

Non-owner virtual router instances are limited by the VRRP specifications to responding to ARP requests destined to the virtual router IP addresses and routing IP packets not addressed to the virtual router IP addresses.

This limitation can be disregarded for certain applications. Ping, Telnet and SSH can be individually enabled or disabled on a per-virtual-router-instance basis.

The **ssh-reply** command enables the non-owner master to reply to SSH requests directed at the virtual router instances IP addresses. The SSH request can be received on any routed interface. SSH must not

have been disabled at the management security level (either on the parental IP interface or based on the SSH source host address). Correct login and CLI command authentication is still enforced.

When **ssh-reply** is not enabled, SSH requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to SSH requests regardless of the **ssh-reply** setting.

The **ssh-reply** command is only available in non-owner **vrrp** nodal context.

By default, SSH requests to the virtual router instance IP addresses are silently discarded.

The **no** form of the command discards all SSH request messages destined to the non-owner virtual router instance IP addresses.

Default

no ssh-reply — SSH requests to the virtual router instance IP addresses are discarded.

Platforms

7705 SAR Gen 2

28.74 ssm

ssm

Syntax

ssm

Context

[Tree] (config>port>ethernet ssm)

Full Context

configure port ethernet ssm

Description

This command enables the Ethernet Synchronization Messaging Channel (ESMC) for the Ethernet port. ESMC carries the Synchronization Status Message (SSM) code representing the quality level of the source of frequency of the central clock of the node.

Platforms

7705 SAR Gen 2

28.75 ssm-assert-compatible-mode

ssm-assert-compatible-mode

Syntax

ssm-assert-compatible-mode [**enable** | **disable**]

Context

[Tree] (config>service>vprn>pim ssm-assert-compatible-mode)

Full Context

configure service vprn pim ssm-assert-compatible-mode

Description

This command specifies whether SSM assert is enabled in compatibility mode for this PIM protocol instance. When enabled, for SSM groups, PIM will consider the SPT bit to be implicitly set to compute the value of CouldAssert (S,G,I) as defined in RFC 4601, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*. When disabled, for SSM groups, PIM will not assume the SPT bit to be set. The SPT bit is set by Update_SPTbit(S,G,iif) macro defined in RFC 4601.

Default

ssm-assert-compatible-mode disable

Parameters

enable

enables SSM assert in compatibility mode for this PIM protocol instance

disable

disabled SSM assert in compatibility mode for this PIM protocol instance

Platforms

7705 SAR Gen 2

28.76 ssm-default-range-disable

ssm-default-range-disable

Syntax

ssm-default-range-disable ipv4

Context

[Tree] (config>service>vprn>pim ssm-default-range-disable)

Full Context

configure service vprn pim ssm-default-range-disable

Description

This command specifies whether to disable the use of default range (232/8) for SSM so that it can be used by ASM to process (*,G). When enabled, the use of default range is disabled for SSM and it can be used by ASM. When disabled, the SSM default range is enabled.

Default

ssm-default-range-disable

Platforms

7705 SAR Gen 2

28.77 ssm-groups

ssm-groups

Syntax

[**no**] **ssm-groups**

Context

[Tree] (config>router>pim ssm-groups)

Full Context

configure router pim ssm-groups

Description

Commands in this context enable an ssm-group configuration instance.

Platforms

7705 SAR Gen 2

28.78 ssm-translate

ssm-translate

Syntax

ssm-translate

Context

[Tree] (config>service>vprn>igmp>if ssm-translate)

[Tree] (config>service>vprn>igmp ssm-translate)

Full Context

configure service vprn igmp interface ssm-translate

configure service vprn igmp ssm-translate

Description

Commands in this context configure group ranges which are translated to SSM (S,G) entries. If the static entry needs to be created, it has to be translated from a IGMPv1 IGMPv2 request to a Source Specific Multicast (SSM) join. An SSM translate source can only be added if the **starg** command is not enabled. An error message is generated if you try to configure the **source** command with **starg** command enabled.

Platforms

7705 SAR Gen 2

ssm-translate

Syntax

ssm-translate

Context

[Tree] (config>service>vprn>mld ssm-translate)

Full Context

configure service vprn mld ssm-translate

Description

Commands in this context configure group ranges which are translated to SSM (S,G) entries. If the static entry needs to be created, it has to be translated from a IGMPv1 IGMPv2 request to a Source Specific Multicast (SSM) join. An SSM translate source can only be added if the starg command is not enabled. An error message is generated if you try to configure the **source** command with **starg** command enabled.

Platforms

7705 SAR Gen 2

ssm-translate

Syntax

ssm-translate

Context

[Tree] (config>router>igmp ssm-translate)

[Tree] (config>router>igmp>if ssm-translate)

Full Context

configure router igmp ssm-translate

configure router igmp interface ssm-translate

Description

Commands in this context configure group ranges which are translated to SSM (S,G) entries. If the static entry needs to be created, it has to be translated from a IGMPv1 IGMPv2 request to a Source Specific Multicast (SSM) join. An SSM translate source can only be added if the starg command is not enabled. An error message is generated if you try to configure the **source** command with **starg** command enabled.

Platforms

7705 SAR Gen 2

ssm-translate

Syntax

ssm-translate

Context

[Tree] (config>router>mld>if ssm-translate)

[Tree] (config>router>mld ssm-translate)

Full Context

configure router mld interface ssm-translate

configure router mld ssm-translate

Description

Commands in this context configure group ranges which are translated to SSM (S,G) entries. If the static entry needs to be created, it has to be translated from a IGMPv1 IGMPv2 request to a Source Specific

Multicast (SSM) join. An SSM translate source can only be added if the starg command is not enabled. An error message is generated if you try to configure the **source** command with **starg** command enabled.

Platforms

7705 SAR Gen 2

28.79 stale-routes-time

stale-routes-time

Syntax

[**no**] **stale-routes-time** *time*

Context

[Tree] (config>service>vprn>bgp>graceful-restart stale-routes-time)

[Tree] (config>service>vprn>bgp>group>graceful-restart stale-routes-time)

[Tree] (config>service>vprn>bgp>group>neighbor>graceful-restart stale-routes-time)

Full Context

configure service vprn bgp graceful-restart stale-routes-time

configure service vprn bgp group graceful-restart stale-routes-time

configure service vprn bgp group neighbor graceful-restart stale-routes-time

Description

This command configures the time period to keep stale routes before the END-OF-RIB message is received from the restarting router.

Default

360 seconds

Parameters

time

1 to 3600 seconds

Platforms

7705 SAR Gen 2

stale-routes-time

Syntax

stale-routes-time *time*

no stale-routes-time

Context

[Tree] (config>router>bgp>group>graceful-restart stale-routes-time)

[Tree] (config>router>bgp>group>neighbor>graceful-restart stale-routes-time)

[Tree] (config>router>bgp>graceful-restart stale-routes-time)

Full Context

configure router bgp group graceful-restart stale-routes-time

configure router bgp group neighbor graceful-restart stale-routes-time

configure router bgp graceful-restart stale-routes-time

Description

This command configures the maximum amount of time in seconds that stale routes should be maintained after a graceful restart is initiated.

The **no** form of this command resets the stale routes time back to the default of 360 seconds.

Default

no stale-routes-time

Parameters

time

Specifies the amount of time that stale routes should be maintained after a graceful restart is initiated.

Values 1 to 3600 seconds

Platforms

7705 SAR Gen 2

28.80 stale-time

stale-time

Syntax

stale-time *seconds*

no stale-time

Context

[Tree] (config>service>vprn>ipv6 stale-time)

[Tree] (config>service>ies>if>ipv6 stale-time)

[Tree] (config>service>vprn>if>ipv6 stale-time)

Full Context

configure service vprn ipv6 stale-time

configure service ies interface ipv6 stale-time

configure service vprn interface ipv6 stale-time

Description

This command configures the time a neighbor discovery cache entry can remain stale before being removed.

The **no** form of this command removes the stale-time value.

Default

no stale-time

Parameters

seconds

The allowed stale time (in seconds) before a neighbor discovery cache entry is removed.

Values 60 to 65535

Platforms

7705 SAR Gen 2

stale-time

Syntax

stale-time *seconds*

no stale-time

Context

[Tree] (config>router>ipv6 stale-time)

Full Context

configure router ipv6 stale-time

Description

This command configures the time a neighbor discovery cache entry can remain stale before being removed.

The **no** form of this command removes the stale-time value.

Default

stale-time 14400

Parameters

seconds

Specifies the allowed stale time (in seconds) before a neighbor discovery cache entry is removed.

Values 60 to 65535

Platforms

7705 SAR Gen 2

stale-time

Syntax

stale-time *seconds*

no stale-time

Context

[Tree] (config>router>origin-validation>rpki-session stale-time)

Full Context

configure router origin-validation rpki-session stale-time

Description

This command configures the maximum length of time that prefix origin validation records learned from the cache server remain usable after the RPKI-Router session goes down. The default stale-time is 3600 seconds (1 hour). When the timer expires all remaining stale entries associated with the session are deleted.

Default

no stale-time

Parameters

seconds

Specifies a time, in seconds.

Values 60 to 3600

Platforms

7705 SAR Gen 2

stale-time

Syntax

stale-time *seconds*

no stale-time

Context

[Tree] (config>router>if>ipv6 stale-time)

Full Context

configure router interface ipv6 stale-time

Description

This command configures the time a neighbor discovery cache entry can remain stale before being removed.

The **no** form of this command removes the stale-time value.

Default

no stale-time

Parameters

seconds

The allowed stale time (in seconds) before a neighbor discovery cache entry is removed.

Values 60 to 65535

Platforms

7705 SAR Gen 2

28.81 standard-multi-instance

standard-multi-instance

Syntax

[**no**] **standard-multi-instance**

Context

[Tree] (config>service>vprn>isis standard-multi-instance)

Full Context

configure service vprn isis standard-multi-instance

Description

This command enables IS-IS multi-instance (MI) as described in draft-ginsberg-isis-mi-bis-01. Multiple instances allow instance-specific adjacencies to be formed that support multiple network topologies on the same physical interfaces. Each instance has an LSDB, and each PDU contains a TLV identifying the instance and the topology to which the PDU belongs. A single topology is supported in each instance, so the instance-specific topology identifier (ITID) is set to 0 and cannot be changed.

The **standard-multi-instance** (based on draft-ginsberg-isis-mi-bis-01) and **iid-tlv-enable** (based on draft-ietf-isis-mi-02) commands cannot be configured in the same instance, because the MAC addresses and PDUs from the two standards are incompatible.

The **no** form of this command removes the **standard-multi-instance** configuration.

Default

no standard-multi-instance

Platforms

7705 SAR Gen 2

standard-multi-instance

Syntax

[**no**] **standard-multi-instance**

Context

[Tree] (config>router>isis standard-multi-instance)

Full Context

configure router isis standard-multi-instance

Description

This command enables IS-IS multi-instance (MI) as described in *draft-ginsberg-isis-mi-bis-01*. Multiple instances allow instance-specific adjacencies to be formed that support multiple network topologies on the same physical interfaces. Each instance has an LSDB, and each PDU contains a TLV identifying the instance and the topology to which the PDU belongs. A single topology is supported in each instance, so the instance-specific topology identifier (ITID) is set to 0 and cannot be changed.

The **standard-multi-instance** (based on *draft-ginsberg-isis-mi-bis-01*) and **iid-tlv-enable** (based on *draft-ietf-isis-mi-02*) commands cannot be configured in the same instance, because the MAC addresses and PDUs from the two standards are incompatible.

The **no** form of this command removes the **standard-multi-instance** configuration.

Default

no standard-multi-instance

Platforms

7705 SAR Gen 2

28.82 standby

standby

Syntax

[**no**] **standby**

Context

[Tree] (config>router>mpls>lsp>secondary standby)

Full Context

configure router mpls lsp secondary standby

Description

The secondary path LSP is normally signaled once the primary path LSP fails. The **standby** keyword ensures that the secondary path LSP is signaled and maintained indefinitely in a hot standby state. Standby paths are selected in preference to non-standby secondary paths. When multiple standby secondary paths exist, then the path-preference is used to determine the order in which the paths are selected. If multiple standby secondary paths have the same, lowest, path-preference value then the system will select the path with the lowest up-time. When the primary path is re-established then the traffic is switched back to the primary path LSP.

The **no** form of this command specifies that the secondary LSP is signaled when the primary path LSP fails.

Platforms

7705 SAR Gen 2

28.83 standby-forwarding

standby-forwarding

Syntax

[**no**] **standby-forwarding**

Context

[Tree] (config>service>ies>if>ipv6>vrrp standby-forwarding)

Full Context

configure service ies interface ipv6 vrrp standby-forwarding

Description

This command allows the forwarding of packets by a standby router.

The **no** form of this command specifies that a standby router should not forward traffic sent to virtual router's MAC address. However, the standby router should forward traffic sent to the standby router's real MAC address.

Default

no standby-forwarding

Platforms

7705 SAR Gen 2

standby-forwarding

Syntax

[**no**] **standby-forwarding**

Context

[Tree] (config>service>ies>if>vrrp standby-forwarding)

Full Context

configure service ies interface vrrp standby-forwarding

Description

This command allows the forwarding of packets by a standby router.

The **no** form of this command specifies that a standby router should not forward traffic sent to virtual router's MAC address. However, the standby router should forward traffic sent to the standby router's real MAC address.

Default

no standby-forwarding

Platforms

7705 SAR Gen 2

standby-forwarding

Syntax

[**no**] **standby-forwarding**

Context

[Tree] (config>service>vprn>if>ipv6>vrrp standby-forwarding)

[Tree] (config>service>vprn>if>vrrp standby-forwarding)

Full Context

configure service vprn interface ipv6 vrrp standby-forwarding

configure service vprn interface vrrp standby-forwarding

Description

This command allows the forwarding of packets by a standby router.

The **no** form of this command specifies that a standby router should not forward traffic sent to virtual router's MAC address. However, the standby router should forward traffic sent to the standby router's real MAC address.

Default

no standby-forwarding

Platforms

7705 SAR Gen 2

standby-forwarding

Syntax

[**no**] **standby-forwarding**

Context

[Tree] (config>router>if>ipv6>vrrp standby-forwarding)

[Tree] (config>router>if>vrrp standby-forwarding)

Full Context

configure router interface ipv6 vrrp standby-forwarding

configure router interface vrrp standby-forwarding

Description

This command specifies whether this VRRP instance allows forwarding packets to a standby router. When disabled, a standby router should not forward traffic sent to virtual router's MAC address. However, the standby router should forward traffic sent to the standby router’s real MAC address. When enabled, a standby router should forward all traffic.

Default

no standby-forwarding

Platforms

7705 SAR Gen 2

28.84 standby-signaling

standby-signaling

Syntax

standby-signaling {**lacp** | **power-off**}

no standby-signaling

Context

[Tree] (config>lag standby-signaling)

Full Context

configure lag standby-signaling

Description

This command specifies how the state of a member port is signaled to the remote side when the status corresponding to this member port has the **standby** value.

Default

standby-signaling lacp

Platforms

7705 SAR Gen 2

28.85 standby-signaling-master

standby-signaling-master

Syntax

[**no**] **standby-signaling-master**

Context

[Tree] (config>service>epipe>endpoint standby-signaling-master)

Full Context

configure service epipe endpoint standby-signaling-master

Description

When this command is enabled, the pseudowire standby bit (value 0x00000020) is sent to T-LDP peer for each spoke SDP of the endpoint that is selected as a standby.

This command is mutually exclusive with a VLL mate SAP created on a mc-lag/mc-aps or ICB. It is also mutually exclusive with vc-switching.

Default

standby-signaling-master

Platforms

7705 SAR Gen 2

28.86 standby-signaling-slave

standby-signaling-slave

Syntax

[**no**] **standby-signaling-slave**

Context

[Tree] (config>service>epipe>spoke-sdp-fec standby-signaling-slave)

Full Context

configure service epipe spoke-sdp-fec standby-signaling-slave

Description

This command enables standby-signaling-slave for an Epipe.

Platforms

7705 SAR Gen 2

standby-signaling-slave

Syntax

[**no**] **standby-signaling-slave**

Context

[Tree] (config>service>epipe>spoke-sdp standby-signaling-slave)

[Tree] (config>service>epipe>endpoint standby-signaling-slave)

Full Context

configure service epipe spoke-sdp standby-signaling-slave

configure service epipe endpoint standby-signaling-slave

Description

When this command is enabled, the node will block the transmit forwarding direction of a spoke SDP based on the pseudowire standby bit received from a T-LDP peer.

This command is present at the endpoint level as well as the spoke SDP level. If the spoke SDP is part of an explicit-endpoint, it will not be possible to change this setting at the spoke SDP level. An existing spoke SDP can be made part of the explicit endpoint only if the settings do not conflict. A newly created spoke SDP, which is part of a specific explicit-endpoint, will inherit this setting from the endpoint configuration.

This command is mutually exclusive with an endpoint that is part of an mc-lag, mc-aps or an ICB.

If the command is disabled, the node assumes the existing independent mode of behavior for the forwarding on the spoke SDP.

Default

no standby-signaling-slave

Platforms

7705 SAR Gen 2

28.87 starg

starg

Syntax

[**no**] **starg**

Context

[Tree] (config>service>vpls>mesh-sdp>mld-snooping>static>group starg)

[Tree] (config>service>vpls>sap>mld-snooping>static>group starg)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping>static>group starg)

[Tree] (config>service>vpls>sap>igmp-snooping>static>group starg)

[Tree] (config>service>vpls>spoke-sdp>mld-snooping>static>group starg)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping>static>group starg)

Full Context

configure service vpls mesh-sdp mld-snooping static group starg

configure service vpls sap mld-snooping static group starg

configure service vpls spoke-sdp igmp-snooping static group starg

configure service vpls sap igmp-snooping static group starg

configure service vpls spoke-sdp mld-snooping static group starg

configure service vpls mesh-sdp igmp-snooping static group starg

Description

This command adds a static (*,g) entry to allow multicast traffic for the corresponding multicast group from any source. This command can only be enabled if no existing source addresses for this group are specified.

The **no** form of this command removes the starg entry from the configuration.

Default

no starg

Platforms

7705 SAR Gen 2

starg

Syntax

starg

Context

[Tree] (config>service>vprn>igmp>if>static>group starg)

Full Context

configure service vprn igmp interface static group starg

Description

This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.

Use the **no** form of this command to remove the starg entry from the configuration.

Platforms

7705 SAR Gen 2

starg

Syntax

[**no**] **starg**

Context

[Tree] (config>service>vprn>mld>if>static>group starg)

Full Context

configure service vprn mld interface static group starg

Description

This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.

Use the **no** form of this command to remove the **starg** entry from the configuration.

Platforms

7705 SAR Gen 2

starg

Syntax

[**no**] **starg**

Context

[Tree] (config>router>igmp>if>static>group starg)

Full Context

configure router igmp interface static group starg

Description

This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.

Use the **no** form of the command to remove the (*,G) entry from the configuration.

Platforms

7705 SAR Gen 2

starg

Syntax

[**no**] **starg**

Context

[Tree] (config>router>mld>if>static>group starg)

Full Context

configure router mld interface static group starg

Description

This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.

The **no** form of this command removes the starg entry from the configuration.

Platforms

7705 SAR Gen 2

28.88 start

start

Syntax

start *start-week start-day start-month hours-minutes*

Context

[Tree] (config>system>time>dst-zone start)

Full Context

configure system time dst-zone start

Description

This command configures start of summer time settings.

Default

start first sunday january 00:00

Parameters

start-week

Specifies the starting week of the month when the summer time takes effect.

Values first, second, third, fourth, last

Default first

start-day

Specifies the starting day of the week when the summer time takes effect.

Values sunday, monday, tuesday, wednesday, thursday, friday, saturday

Default sunday

start-month

Specifies the starting month of the year when the summer time takes effect.

Values january, february, march, april, may, june, july, august, september, october, november, december

Default january

hours-minutes

Specifies the time at which the summer time takes effect, in hh:mm format.

Values hours: 00 to 23
minutes: 00 to 59

Default 00:00

Platforms

7705 SAR Gen 2

28.89 start-label

start-label

Syntax

start-label *start-value* **end-label** *end-value*

no start-label

Context

[Tree] (config>router>mpls-labels>reserved-label-block start-label)

Full Context

configure router mpls-labels reserved-label-block start-label

Description

This command configures start and end labels for a reserved label block. This command must be configured for a reserved label block to be created.

Default

start-label 0, end-label 0

Parameters

start-value

Specifies a starting value.

Values 18432 to 524287 within dynamic label range | 1048575 (FP4 or FP5 only)

end-value

Specifies an ending value.

Values 18432 to 524287 within dynamic label range | 1048575 (FP4 or FP5 only)

Platforms

7705 SAR Gen 2

28.90 startup-wait-time

startup-wait-time

Syntax

startup-wait-time [**min** *minutes*] [**sec** *seconds*] [**hrs** *hours*]

no startup-wait-time [**min** *minutes*] [**sec** *seconds*]

Context

[Tree] (config>router>dhcp6>server>pool>failover startup-wait-time)

[Tree] (config>router>dhcp>server>pool>failover startup-wait-time)

[Tree] (config>router>dhcp>server>failover startup-wait-time)

[Tree] (config>router>dhcp6>server>failover startup-wait-time)

Full Context

configure router dhcp6 local-dhcp-server pool failover startup-wait-time

configure router dhcp local-dhcp-server pool failover startup-wait-time

configure router dhcp local-dhcp-server failover startup-wait-time

configure router dhcp6 local-dhcp-server failover startup-wait-time

Description

This command enables the startup wait time during which each peer waits after the initialization process before assuming the active role for the prefix designated as local or access-driven. This is to avoid transient issues during the initialization process.

The **startup-wait-time** should be configured to an interval in which, after boot, both nodes can set up an MCS TCP link and start MCS. The timer is restarted each time the server downloads a lease from the MCS database and stops when the last state record from the peer is synchronized. The next state is (PRE-)NORMAL, unless the timer times out or is forced to stop via the tools command (**tools>perform>router>dhcp** or **dhcp6>local-dhcp-server** *server-name*>**pool**/**failover**>**abort-startup-wait**), in which case the local DHCP server transitions immediately to the COMMUNICATIONS-INTERRUPTED state.

Default

startup-wait-time min 2

Parameters

minutes

Specifies the startup wait time, in minutes.

Values 1 to 59

seconds

Specifies the startup wait time, in seconds.

Values 1 to 59

hours

Specifies the startup wait time, in hours.

Values 1

Platforms

7705 SAR Gen 2

28.91 stat-mode

stat-mode

Syntax

stat-mode *stat-mode*

no stat mode

Context

[Tree] (config>card>fp>ingress>access>qgrp>policer-over>plcr stat-mode)

[Tree] (config>card>fp>ingress>network>qgrp>policer-over>plcr stat-mode)

Full Context

configure card fp ingress access queue-group policer-override policer stat-mode

configure card fp ingress network queue-group policer-override policer stat-mode

Description

This command configures the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. An ingress policer has multiple types of offered packets (explicit in-profile, explicit out-of-profile, high priority or low priority) and each of these offered types is interacting with the policer's metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the large number of policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly profiled offered packets. The **stat-mode** command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a **no-stats** mode is supported which prevents any packet accounting, the use of the policer's **parent** command requires at the policer's **stat-mode** to be set at least to the **minimal** setting so that offered stats

are available for the policer's Fair Information Rate (FIR) to be calculated. Once a policer has been made a child to a parent policer, the **stat-mode** cannot be changed to **no-stats** unless the policer parenting is first removed.

Each time the policer's **stat-mode** is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. You can view the total/allocated/free stats by using the **tools dump resource-usage** command. If insufficient counters exist to implement a mode on any policer instance, the **stat-mode** change will fail and the previous mode will continue unaffected for all instances of the policer.

The default **stat-mode** when a policer is created within the policy is **minimal**.

The **stat-mode** setting defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the **stat-mode** override command will fail. The previous **stat-mode** setting active for the policer will continue to be used by the policer.

The **no** form of this command attempts to return the policer's stat-mode setting to minimal. The command will fail if insufficient policer counter resources exist to implement minimal where the QoS policer is currently applied and has a forwarding class mapping.

Parameters

See the *7705 SAR Gen 2 Router Configuration Guide* for details on the policer stat-mode parameters.

Platforms

7705 SAR Gen 2

stat-mode

Syntax

stat-mode st*at-mode*

no stat-mode

Context

[Tree] (config>service>epipe>sap>ingress>policer-over>plcr stat-mode)

[Tree] (config>service>epipe>sap>egress>policer-over>plcr stat-mode)

Full Context

configure service epipe sap ingress policer-override policer stat-mode

configure service epipe sap egress policer-override policer stat-mode

Description

The SAP QoS policy's **policer stat-mode** command is used to configure the forwarding plane counters that allow offered, output, and discard accounting to occur for the policer. A policer has multiple types of offered packets (for example, soft in-profile and out-of-profile from ingress and hard in-profile and out-

of-profile due to egress profile overrides) and each of these offered types is interacting with the policers metering and profiling functions resulting in colored output packets (green, yellow, and red). Due to the potentially large number of egress policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly re-profiled offered packets. The **stat-mode** command allows provisioning of the number of counters each policer requires and indicates how the offered packet types and output conditions should be mapped to the counters.

While a **no-stats** mode is supported that prevents any packet accounting, the use of the policer's parent command requires that the policer's **stat-mode** to be set at least to the minimal setting so that offered statistics are available for the policer's Fair Information Rate (FIR) to be calculated.

Each time the policer's stat mode is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. The total/allocated/free statistics can be viewed by using the **tools dump resource-usage card** *slot-num* **fp** *fp-number* command. If insufficient counters exist to implement a mode on any policer instance, the stat-mode change will fail and the previous mode will continue unaffected for all instances of the policer.

The **stat-mode** setting defined for the policer in the QoS policy may be overridden on a SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the stat-mode override command will fail. The current active stat mode setting will continue to be used by the policer.

The command will fail if insufficient policer counter resources exist to implement **minimal** where the QoS policer is currently applied and has a forwarding class mapping.

The **no** form of this command attempts to return the policer's stat-mode setting to **minimal**.

Refer to the *7705 SAR Gen 2 Quality of Service Guide* for detailed information about the supported parameters for the **policer stat-mode** command.

Platforms

7705 SAR Gen 2

stat-mode

Syntax

stat-mode st*at-mode*

no stat-mode

Context

[Tree] (config>service>vpls>sap>egress>policer-override>plcr stat-mode)

[Tree] (config>service>vpls>sap>ingress>policer-override>plcr stat-mode)

Full Context

configure service vpls sap egress policer-override policer stat-mode

configure service vpls sap ingress policer-override policer stat-mode

Description

The SAP-egress QoS policy's policer stat-mode command is used to configure the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. A policer has multiple types of offered packets (for example, soft in-profile and out-of-profile from ingress and hard in-profile and out-of-profile due to egress profile overrides) and each of these offered types is interacting with the policers metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the potential large number of egress policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly re-profiled offered packets. The stat-mode command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a no-stats mode is supported which prevents any packet accounting, the use of the policer's parent command requires that the policer's stat-mode to be set at least to the minimal setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated.

Each time the policer's stat-mode is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. You can view the total/allocated/free stats by using the **tools dump resource-usage card** *slot-num* **fp** *fp-number* command. If insufficient counters exist to implement a mode on any policer instance, the stat-mode change will fail and the previous mode will continue unaffected for all instances of the policer.

The default stat-mode when a policer is created within the policy is minimal.

The stat-mode setting defined for the policer in the QoS policy may be overridden on a SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the stat-mode override command will fail. The previous stat-mode setting active for the policer will continue to be used by the policer.

The **no** form of the command returns the policer's stat-mode setting to minimal.

Refer to the *7705 SAR Gen 2 Quality of Service Guide* for detailed information about the **policer stat-mode** command parameters.

Platforms

7705 SAR Gen 2

stat-mode

Syntax

stat-mode st*at-mode*

no stat-mode

Context

[Tree] (config>service>ies>if>sap>egress>policer-override>plcr stat-mode)

[Tree] (config>service>ies>if>sap>ingress>policer-override>plcr stat-mode)

Full Context

configure service ies interface sap egress policer-override policer stat-mode

configure service ies interface sap ingress policer-override policer stat-mode

Description

The SAP-egress QoS policy's policer stat-mode command is used to configure the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. A policer has multiple types of offered packets (for example, soft in-profile and out-of-profile from ingress and hard in-profile and out-of-profile due to egress profile overrides) and each of these offered types is interacting with the policers metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the potential large number of egress policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly re-profiled offered packets. The stat-mode command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a no-stats mode is supported which prevents any packet accounting, the use of the policer's parent command requires that the policer's stat-mode to be set at least to the minimal setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated.

Each time the policer's stat-mode is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. You can view the total/allocated/free stats by using the **tools dump resource-usage card** *slot-num* **fp** *fp-number* command. If insufficient counters exist to implement a mode on any policer instance, the stat-mode change will fail and the previous mode will continue unaffected for all instances of the policer.

The default stat-mode when a policer is created within the policy is minimal.

The stat-mode setting defined for the policer in the QoS policy may be overridden on a SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the stat-mode override command will fail. The previous stat-mode setting active for the policer will continue to be used by the policer.

The **no** form of this command returns the policer's stat-mode setting to minimal.

Refer to the *7705 SAR Gen 2 Quality of Service Guide* for detailed information about the **policer stat-mode** command parameters.

Platforms

7705 SAR Gen 2

stat-mode

Syntax

stat-mode st*at-mode*

no stat-mode

Context

[Tree] (config>service>vprn>if>sap>egress>policer-override>plcr stat-mode)

[Tree] (config>service>vprn>if>sap>ingress>policer-override>plcr stat-mode)

Full Context

configure service vprn interface sap egress policer-override policer stat-mode

configure service vprn interface sap ingress policer-override policer stat-mode

Description

The SAP-egress QoS policy's policer stat-mode command is used to configure the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. A policer has multiple types of offered packets (for example, soft in-profile and out-of-profile from ingress and hard in-profile and out-of-profile due to egress profile overrides) and each of these offered types is interacting with the policers metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the potential large number of egress policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly re-profiled offered packets. The stat-mode command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a no-stats mode is supported which prevents any packet accounting, the use of the policer's parent command requires that the policer's stat-mode to be set at least to the minimal setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated.

Each time the policer's stat-mode is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. You can view the total/allocated/free stats by using the **tools dump resource-usage card** *slot-num* **fp** *fp-number* command. If insufficient counters exist to implement a mode on any policer instance, the stat-mode change will fail and the previous mode will continue unaffected for all instances of the policer.

The default stat-mode when a policer is created within the policy is minimal.

The stat-mode setting defined for the policer in the QoS policy may be overridden on a SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the stat-mode override command will fail. The previous stat-mode setting active for the policer will continue to be used by the policer.

The **no** form of this command returns the policer's stat-mode setting to minimal.

Refer to the *7705 SAR Gen 2 Quality of Service Guide* for detailed information about the **policer stat-mode** command parameters.

Platforms

7705 SAR Gen 2

stat-mode

Syntax

stat-mode {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-priority-no-cir** | **offered-profile-cir** | **offered-priority-cir** | **offered-total-cir** | **offered-limited-profile-cir** | **offered-profile-capped-cir** | **offered-limited-capped-cir**}

no stat mode

Context

[Tree] (config>qos>sap-ingress>policer stat-mode)

Full Context

configure qos sap-ingress policer stat-mode

Description

This command is used to configure the forwarding plane counters that allow offered, forwarded, and dropped accounting to occur for the policer. An ingress policer has multiple types of offered packets (explicit in-profile, explicit out-of-profile, uncolored, high-priority, or low priority) and each of these offered types is interacting with the policer's metering and profiling functions resulting in colored output packets (green, yellow, and red). Due to the large number of policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers, for example, will not be configured with a CIR profiling rate and not all policers will receive explicitly profiled offered packets. The **stat-mode** command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a **no-stats** mode is supported that prevents any packet accounting, the use of the policer's **parent** command requires that the policer's **stat-mode** be set at least to the **minimal** setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated. When a policer has been made a child to a parent policer, the **stat-mode** cannot be changed to **no-stats** unless the policer parenting is first removed.

Each time the policer's **stat-mode** is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. The total/allocated/free stats can be viewed by using the **tools dump resource-usage card fp** command. If insufficient counters exist to implement a mode on any policer instance, the **stat-mode** change will fail and the previous mode will continue unaffected for all instances of the policer.

The ingress policer stat-modes are described in Table 96: Ingress Policer Stat Mode Summary .

Table 96: Ingress Policer Stat Mode Summary

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
no-stats	0	—	—	—
Minimal	1	Single counter entering policer	Single counter for dropped/forwarded exiting policer	—
offered-profile-no-cir	2	In/out entering policer	In/out entering policer	Intended for when the policer does not change the profile of packets. Includes only in-profile and out-of-profile.
offered-priority-no-cir	2	High/low entering policer	High/low entering policer	Intended for when only packet priority stats are required.
offered-profile-cir	4	In/out/uncolored entering policer	In/out exiting policer	Intended for when the policer can change the profile of packets to in-profile and out-of-profile.
offered-priority-cir	4	High/low entering policer	In/out exiting policer	Intended for when packet priority entering the policer and profile exiting the policer is required.
offered-total-cir	2	Single counter entering policer	In/out exiting policer	—
offered-limited-profile-cir	3	Out/uncolored entering policer	In/out exiting policer	Intended for when the policer can change the profile of packet to in-profile and out-of-profile. The information is limited compared to offered-profile-capped-cir with the benefit of using one less stat resource.

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
offered-profile-capped-cir	5	In/out/uncolored entering policer	In/out exiting policer	Intended for when the policer has profile-capped configured.
offered-limited-capped-cir	4	In/uncolored entering policer	In/out exiting policer	Intended for when the policer has profile-capped configured. The information is limited compared to offered-profile-capped-cir with the benefit of using one less stat resource.

The default **stat-mode** when a policer is created within the policy is **minimal**.

The **stat-mode** setting defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the **stat-mode** override command will fail. The previous **stat-mode** setting active for the policer will continue to be used by the policer.

The **no** form of this command attempts to return the policer's stat-mode setting to minimal. The command will fail if insufficient policer counter resources exist to implement minimal where the QoS policer is currently applied and has a forwarding class mapping.

Parameters

no-stats

Counter resource allocation: 0

The policer does not have any forwarding plane counters allocated and cannot provide offered, dropped, and forwarded statistics. A policer using no-stats cannot be a child to a parent policer and the policer's parent command will fail.

When **collect-stats** is enabled, no statistics are generated.

minimal

Counter resource allocation: 1

This stat-mode provides the minimal accounting resource usage and counter information, and includes the total offered, dropped and forwarded packet and octet counters for traffic entering (offered) and exiting (dropped/forwarded) the policer.

The default **stat-mode** for a policer is **minimal**. The **minimal** mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types (profile or priority) and do not count in-profile or out-of-

profile output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate.

This counter mode is useful when only the most basic accounting information is required.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 97: Ingress Accounting Statistics Collected in minimal stat-mode .

Table 97: Ingress Accounting Statistics Collected in minimal stat-mode

Show Output	Accounting Statistics Collected	
	Field	**Field Description**
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. All	apd	AllPacketsDropped
	aod	AllOctetsDropped
For. All	apf	AllPacketsForwarded
	aof	AllOctetsForwarded

offered-profile-no-cir

Counter resource allocation: 2

This **stat-mode** provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering the policer.

The **offered-profile-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-profile-no-cir** mode is most useful when the policer is receiving only in-profile and out-of-profile premarked (and trusted) packets. It is expected that, in this instance, a CIR rate will not be defined since all packets are already premarked. This mode does not prevent the policer from receiving untrusted (color undefined) traffic nor does it prevent the policer from being configured with a CIR rate.

This mode is intended to be used without profile-capped configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 98: Ingress Accounting Statistics Collected in offered-profile-no-cir stat-mode .

Table 98: Ingress Accounting Statistics Collected in offered-profile-no-cir stat-mode

Show Output	Accounting Statistics Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-priority-no-cir

Counter resource allocation: 2

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the packet priority of traffic entering the policer.

The **offered-priority-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-priority-no-cir** mode is most useful when the policer is receiving only untrusted packets and the ingress priority high and priority low classification options are being used without a CIR profiling rate defined. This mode does not prevent the policer from receiving trusted packets that are premarked in-profile or out-of-profile nor does it prevent the policer from being configured with a CIR rate.

This mode is intended to be used without profile-capped configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 99: Ingress Accounting Statistics Collected in offered-priority-no-cir stat-mode .

Table 99: Ingress Accounting Statistics Collected in offered-priority-no-cir stat-mode

Show Output	Accounting Statistics Collected	
	Field	Field Description
Off. HiPrio	hpo	HighPriorityPacketsOffered
	hoo	HighPriorityOctetsOffered
Off. LowPrio	lpo	LowPriorityPacketsOffered
	loo	LowPriorityOctetsOffered
Dro. HiPrio	hpd	HighPriorityPacketsDropped
	hod	HighPriorityOctetsDropped
Dro. LowPrio	lpd	LowPriorityPacketsDropped
	lod	LowPriorityOctetsDropped
For. HiPrio	hpf	HighPriorityPacketsForwarded
	hof	HighPriorityOctetsForwarded
For. LowPrio	lpf	LowPriorityPacketsForwarded
	lof	LowPriorityOctetsForwarded

offered-profile-cir

Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed so that the traffic entering the policer comprises hard in/out and uncolored traffic. The offered counters cover traffic explicitly profiled to in-profile, traffic explicitly profiled to out-of-profile, and traffic that has not been explicitly profiled at ingress (uncolored).

The **offered-profile-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile and in-profile traffic and is also receiving untrusted packets that are being applied to a defined CIR profiling rate. This mode differs from **offered-limited-profile-cir** mode in that it expects both trusted in-profile and out-of-profile packets while still performing CIR profiling on packets with untrusted markings. If trusted in-profile packets are not being received, the **offered-limited-profile-cir** stat-mode could be used instead, which has the benefit of using a reduced number of stat resources.

This mode is intended to be used without **profile-capped** configured within the policer as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 100: Ingress Accounting Statistics Collected in offered-profile-cir stat-mode .

Table 100: Ingress Accounting Statistics Collected in offered-profile-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-priority-cir

Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the priority of traffic entering the policer and the profile exiting the policer.

The **offered-priority-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-priority-cir** mode is most useful when the policer is receiving only untrusted packets that are being classified as high priority or low priority and are being applied to a defined CIR profiling rate. This mode differs from **offered-profile-cir** mode in that it does not expect trusted in-profile and out-of-profile packets but does not exclude the ability of the policer to receive them.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 101: Ingress Accounting Statistics Collected in offered-priority-cir stat-mode .

Table 101: Ingress Accounting Statistics Collected in offered-priority-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. HiPrio	hpo	HighPriorityPacketsOffered
	hoo	HighPriorityOctetsOffered
Off. LowPrio	lpo	LowPriorityPacketsOffered
	loo	LowPriorityOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-total-cir

Counter resource allocation: 2

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter.

The **offered-total-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-total-cir** mode is most useful when the policer is not receiving trusted in-profile or out-of-profile traffic and both high- and low-priority classifications are not being used on the untrusted packets and the offered packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile or out-of-profile packets and does not prevent the use of priority high or low classifications on the untrusted packets.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 102: Ingress Accounting Statistics Collected in offered-total-cir stat-mode .

Table 102: Ingress Accounting Statistics Collected in offered-total-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-limited-profile-cir

Counter resource allocation: 3

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed so that the traffic entering the policer comprises of hard out and uncolored. The offered counters cover traffic explicitly profiled to out-of-profile and traffic that has not been explicitly profiled at ingress (Uncolor). The traffic explicitly profiled to in-profile is counted with the uncolored traffic.

The **offered-limited-profile-cir** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-limited-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile (profile out but no profile in) traffic and untrusted packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile packets. If trusted in-profile packets are not being received, the **offered-limited-profile-cir** is preferred over **offered-profile-cir** because it uses a reduced number of stat resources.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 103: Ingress Accounting Statistics Collected in offered-limited-profile-cir stat-mode .

Table 103: Ingress Accounting Statistics Collected in offered-limited-profile-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-profile-cir

Counter resource allocation: 4

offered-profile-capped-cir

Counter resource allocation: 5

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed so that the traffic entering the policer comprises of hard in/out and uncolored. The offered counters cover traffic explicitly profiled to in-profile, traffic explicitly profiled to out-of-profile, and traffic that has not been explicitly profiled at ingress (Uncolor).

When **offered-profile-capped-cir** is defined, the system creates five offered-output counters in the forwarding plane and five discard counters in the traffic manager.

The **offered-profile-capped-cir** mode is similar to the offered-profile-cir mode except that it includes support for profile in and **soft-in-profile** that may be output as out-of-profile due to enabling **profile-capped** mode on the ingress policer.

The impact of using **offered-profile-capped-cir** stat-mode while **profile-capped** mode is disabled is that one of the counting resources in the forwarding plane and traffic manager will not be used.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 104: Ingress Accounting Statistics Collected in offered-profile-capped-cir stat-mode .

Table 104: Ingress Accounting Statistics Collected in offered-profile-capped-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-limited-capped-cir

Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed resulting in the traffic entering the policer comprising of hard in/out and uncolored. The offered counters cover in-profile traffic and traffic that has not been explicitly profiled at ingress (Uncolor). The traffic explicitly profiled to out-of-profile is counted with the uncolored traffic.

When **offered-limited-capped-cir** is defined, the system creates four forwarding plane offered-output counters in the network processor and four discard counters in the traffic manager.

The **offered-limited-capped-cir** mode is similar to the **offered-profile-capped-cir** mode except that it combines **soft in-profile** with **profile in** (InProf) and **profile out** (OutProf) with **soft-out-of-profile** (Uncolor) and eliminates the "offered undefined" statistic. If trusted out-of-profile packets are not being received, the **offered-limited-capped-cir** is preferred over **offered-profile-capped-cir** because it uses a reduced number of stat resources.

This mode is intended to be used with **profile-capped** configured within the policer.

The impact of using **offered-limited-capped-cir** stat-mode while **profile-capped** mode is disabled is that one of the counting resources in the forwarding plane and traffic manager will not be used.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 105: Ingress Accounting Statistics Collected in offered-limited-capped-cir stat-mode .

Table 105: Ingress Accounting Statistics Collected in offered-limited-capped-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	**Field Description**
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

Platforms

7705 SAR Gen 2

stat-mode

Syntax

stat-mode {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-profile-cir** | **offered-total-cir** | **offered-limited-capped-cir** | **offered-profile-capped-cir** | **offered-total-cir-exceed** | **offered-four-profile-no-cir** | **offered-total-cir-four-profile**}

no stat mode

Context

[Tree] (config>qos>sap-egress>policer stat-mode)

Full Context

configure qos sap-egress policer stat-mode

Description

The **sap-egress** QoS policy's policer **stat-mode** command is used to configure the forwarding plane counters that allow offered, forwarded, and dropped accounting to occur for the policer. An egress policer has multiple types of offered packets (soft in-profile and out-of-profile from ingress and hard in-profile, out-of-profile, and exceed-profile due to egress profile overrides) and each of these offered types is interacting with the policer's metering and profiling functions resulting in colored output packets (green, yellow, and red). Due to the potential large number of egress policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers, for example, will not be configured with a CIR profiling rate and not all policers will receive explicitly reprofiled offered packets. The **stat-mode** command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a **no-stats** mode is supported that prevents any packet accounting, the use of the policer's **parent** command requires that the policer's **stat-mode** be set at least to the **minimal** setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated. When a policer has been made a child to a parent policer, the **stat-mode** cannot be changed to **no-stats** unless the policer parenting is first removed.

Each time the policer's **stat-mode** is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. The total, allocated, and free statistics can be viewed by using the **tools dump resource-usage card fp** command. If insufficient counters exist to implement a mode on any policer instance, the **stat-mode** change will fail and the previous mode will continue unaffected for all instances of the policer.

The egress policer stat-modes are described in Table 106: Egress Policer Stat-mode Summary.

Table 106: Egress Policer Stat-mode Summary

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
no-stats	0	—	—	—
minimal	1	Single counter entering policer	Single counter for dropped/forwarded exiting policer	—
offered-profile-no-cir	2	In or out entering policer	In/out entering policer	Intended for when the policer does not change the profile of packets. Includes only in-profile and out-of-profile.
offered-profile-cir	4	In, out, or uncolored (which corresponds to hard in-profile, hard out-of-profile, or soft in- or out-of-profile) entering policer	In/out exiting policer	Intended for when the policer can change the profile of packets to in-profile and out-of-profile.
offered-total-cir	2	Single counter entering policer	In/out exiting policer	—
offered-limited-capped-cir	4	In or out entering policer	In/out exiting policer	Intended for when the policer has **profile-capped** configured. The information is limited compared to **offered-profile-capped-cir** with the benefit of using one less stat resource.
offered-profile-capped-cir	5	In, out, or uncolored (which corresponds to hard in-profile, hard out-of-profile, or soft in- or out-of-profile) entering policer	In/out exiting policer	Intended for when the policer has **profile-capped** configured
offered-total-cir-exceed	3	Single counter entering policer	In/out/exceed exiting policer	Intended for when the policer is

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
				configured with **enable-exceed-pir** to forward packets that exceed its configured PIR or when traffic is reclassified at egress to exceed-profile
offered-four-profile-no-cir	4	Inplus, in, out, or exceed entering policer	Inplus/in/out/exceed entering policer	Intended to be used when the policer does not change the profile of the packets and traffic is reclassified at egress to inplus and/ or exceed-profile
offered-total-cir-four-profile	4	Single counter entering policer	Inplus, in, out, or exceed exiting policer	Intended to be used when the policer can change the profile of the packet and traffic is reclassified at egress to profile inplus

When a policer is created within the policy, the default setting for **stat-mode** is **minimal**.

The **stat-mode** setting defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the **stat-mode** override command will fail. The previous **stat-mode** setting active for the policer will continue to be used by the policer.

The **no** form of this command attempts to return the policer's **stat-mode** setting to **minimal**. The command will fail if insufficient policer counter resources exist to implement **minimal** where the QoS policer is currently applied and has a forwarding class mapping.

Parameters

no-stats

Counter resource allocation: 0

The policer does not have any forwarding plane counters allocated and cannot provide offered, discard, and forward statistics. A policer using **no-stats** cannot be a child to a parent policer and the policer's **parent** command will fail.

When **collect-stats** is enabled, no statistics are generated.

minimal

Counter resource allocation: 1

This stat-mode provides the minimal accounting resource usage and counter information, and includes only the total offered, dropped and forwarded packet and octet counters for traffic entering (offered) and exiting (dropped/forwarded) the policer.

The default **stat-mode** for a policer is **minimal**. The **minimal** mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types and do not count different profile output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate or using exceed PIR.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 107: Egress Accounting Statistics Collected in minimal stat-mode .

Table 107: Egress Accounting Statistics Collected in minimal stat-mode

Show Output	Accounting Stats Collected	
	Field	**Field Description**
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. All	apd	AllPacketsDropped
	aod	AllOctetsDropped
For. All	apf	AllPacketsForwarded
	aof	AllOctetsForwarded

offered-profile-no-cir

Counter resource allocation: 2

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering the policer. inplus-profile traffic is counted with the in-profile counters and exceed-profile traffic is counted with the out-of-profile counters.

The **offered-profile-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-profile-no-cir** mode is most useful when profile-based offered, dropped, and forwarded stats are required from the egress policer, but a CIR or **enable-exceed-pir** is not being used to recolor the soft in-profile and out-of-profile packets. This mode does not prevent the policer from being configured with a CIR rate or using **enable-exceed-pir**.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 108: Egress Accounting Statistics Collected in offered-profile-no-cir stat-mode .

Table 108: Egress Accounting Statistics Collected in offered-profile-no-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-profile-cir

Counter resource allocation: 4

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer is made up of traffic that is inplus-profile, in-profile, out-of-profile, exceed-profile, soft in-profile, and soft out-of-profile. The offered counters cover traffic reclassified to in-profile (which includes traffic reclassified to inplus-profile), traffic reclassified to out-of-profile (which includes traffic reclassified to exceed-profile) and traffic which has not been reclassified at egress (Uncolor). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

The **offered-profile-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-profile-cir** mode is most useful when profile-based offered, dropped and forwarded stats are required from the egress policer and a CIR rate is being used to recolor the soft in-profile and out-of-profile packets.

This mode is intended to be used without **profile-capped** or **enable-exceed-pir** configured within the policer as these could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 109: Egress Accounting Statistics Collected in offered-profile-cir stat-mode .

Table 109: Egress Accounting Statistics Collected in offered-profile-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-total-cir

Counter resource allocation: 2

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

The **offered-total-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-total-cir** mode is most useful when the policer is not receiving trusted in-profile or out-of-profile traffic, and both high- and low- priority classifications are not being used on the untrusted packets, and the offered packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile or out-of-profile packets and does not prevent the use of priority high or low classifications on the untrusted packets.

This mode is intended to be used without **profile-capped** or **enable-exceed-pir** configured within the policer as these could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 110: Egress Accounting Statistics Collected in offered-total-cir stat-mode .

Table 110: Egress Accounting Statistics Collected in offered-total-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	**Field Description**
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-limited-capped-cir

Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer is made up of traffic that is inplus-profile, in-profile, out-of-profile, exceed-profile, soft in-profile, and soft out-of-profile. The offered counters cover in-profile traffic (which includes traffic reclassified to inplus-profile) and out-of-profile traffic (which includes traffic reclassified to exceed-profile). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

offered-limited-capped-cir is defined, the system creates four forwarding plane offered-output counters in the network processor and three discard counters in the traffic manager.

The **offered-limited-capped-cir** mode is similar to the **offered-profile-capped-cir** mode except that it combines **soft in-profile** with **profile in** and **soft-out-of-profile** with **profile out** and eliminates the offered-undefined statistic.

The impact of using **offered-limited-capped-cir** stat-mode while profile-capped mode is disabled is that one of the counting resources in the forwarding plane and traffic manager will not be used and **soft-in-profile** will be treated as offered-in instead of offered-undefined.

This mode is intended to be used with **profile-capped** configured within the policer but without **enable-exceed-pir** configured as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 111: Egress Accounting Statistics Collected in offered-limited-capped-cir stat-mode .

Table 111: Egress Accounting Statistics Collected in offered-limited-capped-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-profile-capped-cir

Counter resource allocation: 5

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded)

the policer when egress reclassification is performed so that the traffic entering the policer is made up of traffic that is inplus-profile, in-profile, out-of-profile, exceed-profile, soft in-profile, and soft out-of-profile. The offered counters cover traffic reclassified to in-profile (which includes traffic reclassified to inplus-profile), traffic reclassified to out-of-profile (which includes traffic reclassified to exceed-profile) and traffic that has not been reclassified at egress (uncolored). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

When **offered-profile-capped-cir** is defined, the system creates five offered-output counters in the forwarding plane and five discard counters in the traffic manager.

The **offered-profile-capped-cir** mode is similar to the **offered-profile-cir** mode except that it includes support for **profile inplus**, **profile in** and **soft-in-profile** that may be output as out-of-profile due to enabling **profile-capped** mode on the ingress policer.

The impact of using o**ffered-profile-capped-cir** stat-mode while **profile-capped** mode is disabled is that one of the counting resources in the forwarding plane and traffic manager will not be used and soft-in-profile will be treated as offered-in (hard in-profile) instead of offered-undefined (uncolored).

This mode is intended to be used with **profile-capped** configured within the policer but without **enable-exceed-pir** configured as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 112: Egress Accounting Statistics Collected in offered-profile-capped-cir stat-mode .

Table 112: Egress Accounting Statistics Collected in offered-profile-capped-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	**Field Description**
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped

Show Output	Accounting Stats Collected	
	Field	Field Description
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-total-cir-exceed

Counter resource allocation: 3

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter. The **offered-total-cir-exceed** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-total-cir-exceed** mode is similar to the **offered-total-cir** mode except that it includes support for forwarded and dropped counters for **profile exceed**.

This mode is intended to be used when the policer is configured with **enable-exceed-pir** to forward packets that exceed its configured PIR or when traffic is egress reclassified to profile exceed. The mode gives the forwarded and dropped counters per profile (in, out, exceed). It is also intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer. This stat-mode is not supported for dynamic policers.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 113: Egress Accounting Statistics Collected in offered-total-cir-exceed stat-mode .

Table 113: Egress Accounting Statistics Collected in offered-total-cir-exceed stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped

Show Output	Accounting Stats Collected	
	Field	Field Description
Dro. ExcProf	xpd	ExceedProfilePktsDropped
	xod	ExceedProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded
For. ExcProf	xpf	ExceedProfilePktsForwarded
	xof	ExceedProfileOctetsForwarded

offered-four-profile-no-cir

Counter resource allocation: 4

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering the policer. Offered, dropped, and forwarded counters are provided for inplus, in, out and exceed-profile traffic.

The **offered-four-profile-no-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-four-profile-no-cir** mode is similar to the **offered-profile-no-cir** mode except that it includes support for offered, dropped, and forwarded counters for both inplus-profile and exceed-profile.

This mode is intended to be used when traffic is egress reclassified to inplus and/or exceed-profile. It is also intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer. This stat-mode is not supported for dynamic policers.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 114: Egress Accounting Statistics Collected in offered-four-profile-no-cir stat-mode .

Table 114: Egress Accounting Statistics Collected in offered-four-profile-no-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered

Show Output	Accounting Stats Collected	
	Field	Field Description
	ooo	OutOfProfileOctetsOffered
Off. ExcProf	xpo	ExceedProfilePacketsOffered
	xoo	ExceedProfileOctetsOffered
Off. InplusProf	ppo	InplusProfilePacketsOffered
	poo	InplusProfileOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
Dro. ExcProf	xpd	ExceedProfilePktsDropped
	xod	ExceedProfileOctetsDropped
Dro. InprofProf	ppd	InplusProfilePktsDropped
	pod	InplusProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded
For. ExcProf	xpf	ExceedProfilePktsForwarded
	xof	ExceedProfileOctetsForwarded
For. InplusProf	ppf	InplusProfilePktsForwarded
	pof	InplusProfileOctetsForwarded

offered-total-cir-four-profile

Counter resource allocation: 4

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. There is a separate dropped and forwarded counter for inplus, in, out and exceed-profile traffic.

The **offered-total-cir-four-profile** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-total-cir-four-profile** mode is similar to the **offered-total-cir** except that it includes support for forwarded and dropped counters for both **profile inplus** and **profile exceed**.

This mode is intended to be used when traffic is reclassified at egress to inplus-profile. It is also intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer. This stat-mode is not supported for dynamic policers.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 115: Egress Accounting Statistics Collected in offered-total-cir-four-profile stat-mode .

Table 115: Egress Accounting Statistics Collected in offered-total-cir-four-profile stat-mode

Show Output	Accounting Stats Collected	
	Field	**Field Description**
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
Dro. ExcProf	xpd	ExceedProfilePktsDropped
	xod	ExceedProfileOctetsDropped
Dro. InprofProf	ppd	InplusProfilePktsDropped
	pod	InplusProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded
For. ExcProf	xpf	ExceedProfilePktsForwarded
	xof	ExceedProfileOctetsForwarded
For. InplusProf	ppf	InplusProfilePktsForwarded
	pof	InplusProfileOctetsForwarded

Platforms

7705 SAR Gen 2

stat-mode

Syntax

stat-mode {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-priority-no-cir** | **offered-profile-cir** | **offered-priority-cir** | **offered-total-cir** | **offered-limited-profile-cir** | **offered-profile-capped-cir** | **offered-limited-capped-cir**}

no stat mode

Context

[Tree] (config>qos>qgrps>ing>qgrp>policer stat-mode)

Full Context

configure qos queue-group-templates ingress queue-group policer stat-mode

Description

This command is used to configure the forwarding plane counters that allow offered, forwarded, and dropped accounting to occur for the policer. An ingress policer has multiple types of offered packets (explicit in-profile, explicit out-of-profile, uncolored, high-priority or low-priority) and each of these offered types is interacting with the policer's metering and profiling functions resulting in colored output packets (green, yellow, and red). Due to the large number of policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers, for example, will not be configured with a CIR profiling rate and not all policers will receive explicitly profiled offered packets. The **stat-mode** command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a **no-stats** mode is supported that prevents any packet accounting, the use of the policer's **parent** command requires that the policer's **stat-mode** be set at least to the **minimal** setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated. When a policer has been made a child to a parent policer, the **stat-mode** cannot be changed to **no-stats** unless the policer parenting is first removed.

Each time the policer's **stat-mode** is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. The total/allocated/free stats can be viewed by using the **tools dump resource-usage card fp** command. If insufficient counters exist to implement a mode on any policer instance, the **stat-mode** change will fail and the previous mode will continue unaffected for all instances of the policer.

The stat-modes are described in Table 116: Stat Mode Descriptions.

Table 116: Stat Mode Descriptions

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
no-stats	0	None	None	—
Minimal	1	Single counter entering policer	Single counter for dropped/forwarded exiting policer	—
offered-profile-no-cir	2	In/out entering policer	In/out entering policer	Intended for when the policer does not change the profile of packets. Includes only in- and out-of-profile.
offered-priority-no-cir	2	High/low entering policer	High/low entering policer	Intended for when only packet priority stats are required.
offered-profile-cir	4	In/out/uncolored entering policer	In/out exiting policer	Intended for when the policer can change the profile of packets to in- and out-of-profile.
offered-priority-cir	4	High/low entering policer	In/out exiting policer	Intended for when packet priority entering the policer and profile exiting the policer is required.
offered-total-cir	2	Single counter entering policer	In/out exiting policer	—
offered-limited-profile-cir	3	Out/uncolored entering policer	In/out exiting policer	Intended for when the policer can change the profile of packet to in- and out-of-profile. The information is limited compared to offered-profile-capped-cir with the benefit of using one less stat resource.

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
offered-profile-capped-cir	5	In/out/uncolored entering policer	In/out exiting policer	Intended for when the policer has profile-capped configured.
offered-limited-capped-cir	4	In/uncolored entering policer	In/out exiting policer	Intended for when the policer has profile-capped configured. The information is limited compared to offered-profile-capped-cir with the benefit of using one less stat resource.

The default **stat-mode** when a policer is created within the policy is **minimal**.

The **stat-mode** setting defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the **stat-mode** override command will fail. The previous **stat-mode** setting active for the policer will continue to be used by the policer.

The **no** form of this command attempts to return the policer's stat-mode setting to minimal. The command will fail if insufficient policer counter resources exist to implement minimal where the QoS policer is currently applied and has a forwarding class mapping.

Parameters

no-stats

Counter resource allocation: 0

The policer does not have any forwarding plane counters allocated and cannot provide offered, dropped and forwarded statistics. A policer using no-stats cannot be a child to a parent policer and the policer's parent command will fail.

When **collect-stats** is enabled, no statistics are generated.

minimal

Counter resource allocation: 1

This stat-mode provides the minimal accounting resource usage and counter information, and includes the total offered, dropped and forwarded packet and octet counters for traffic entering (offered) and exiting (dropped/forwarded) the policer.

The default **stat-mode** for a policer is **minimal**. The **minimal** mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types (profile or priority) and do not count in-profile or out-of-

profile output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate.

This counter mode is useful when only the most basic accounting information is required.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 117: Ingress Accounting Statistics Collected in minimal stat-mode .

Table 117: Ingress Accounting Statistics Collected in minimal stat-mode

Show Output	Accounting Stats Collected	
	Field	**Field Description**
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. All	apd	AllPacketsDropped
	aod	AllOctetsDropped
For. All	apf	AllPacketsForwarded
	aof	AllOctetsForwarded

offered-profile-no-cir

Counter resource allocation: 2

This **stat-mode** provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering the policer.

The **offered-profile-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-profile-no-cir** mode is most useful when the policer is receiving only in-profile and out-of-profile premarked (and trusted) packets. It is expected that, in this instance, a CIR rate will not be defined since all packets are already premarked. This mode does not prevent the policer from receiving untrusted (color undefined) nor does it prevent the policer from being configured with a CIR rate.

This mode is intended to be used without profile-capped configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 118: Ingress Accounting Statistics Collected in offered-profile-no-cir stat-mode .

Table 118: Ingress Accounting Statistics Collected in offered-profile-no-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-priority-no-cir

Counter resource allocation: 2

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the packet priority of traffic entering the policer.

The **offered-priority-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-priority-no-cir** mode is most useful when the policer is receiving only untrusted packets and the ingress priority high and priority low classification options are being used without a CIR profiling rate defined. This mode does not prevent the policer from receiving trusted packets that are premarked in-profile or out-of-profile nor does it prevent the policer from being configured with a CIR rate.

This mode is intended to be used without profile-capped configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 119: Ingress Accounting Statistics Collected in offered-priority-no-cir stat-mode .

Table 119: Ingress Accounting Statistics Collected in offered-priority-no-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. HiPrio	hpo	HighPriorityPacketsOffered
	hoo	HighPriorityOctetsOffered
Off. LowPrio	lpo	LowPriorityPacketsOffered
	loo	LowPriorityOctetsOffered
Dro. HiPrio	hpd	HighPriorityPacketsDropped
	hod	HighPriorityOctetsDropped
Dro. LowPrio	lpd	LowPriorityPacketsDropped
	lod	LowPriorityOctetsDropped
For. HiPrio	hpf	HighPriorityPacketsForwarded
	hof	HighPriorityOctetsForwarded
For. LowPrio	lpf	LowPriorityPacketsForwarded
	lof	LowPriorityOctetsForwarded

offered-profile-cir

Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed so that the traffic entering the policer comprises of hard in/out and uncolored. The offered counters cover traffic explicitly profiled to in-profile, traffic explicitly profiled to out-of-profile, and traffic that has not been explicitly profiled at ingress (uncolored).

The **offered-profile-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile and in-profile traffic and is also receiving untrusted packets that are being applied to a defined CIR profiling rate. This mode differs from **offered-limited-profile-cir** mode in that it expects both trusted in-profile and out-of-profile packets while still performing CIR profiling on packets with untrusted markings. If trusted in-profile packets are not being received, the **offered-limited-profile-cir** stat-mode could be used instead, which has the benefit of using a reduced number of stat resources.

This mode is intended to be used without **profile-capped** configured within the policer as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 120: Ingress Accounting Statistics Collected in offered-profile-cir stat-mode .

Table 120: Ingress Accounting Statistics Collected in offered-profile-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-priority-cir

Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the priority of traffic entering the policer and the profile exiting the policer.

The **offered-priority-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-priority-cir** mode is most useful when the policer is receiving only untrusted packets that are being classified as high priority or low priority and are being applied to a defined CIR profiling rate. This mode differs from **offered-profile-cir** mode in that it does not expect trusted in-profile and out-of-profile packets but does not exclude the ability of the policer to receive them.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 121: Ingress Accounting Statistics Collected in offered-priority-cir stat-mode .

Table 121: Ingress Accounting Statistics Collected in offered-priority-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. HiPrio	hpo	HighPriorityPacketsOffered
	hoo	HighPriorityOctetsOffered
Off. LowPrio	lpo	LowPriorityPacketsOffered
	loo	LowPriorityOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-total-cir

Counter resource allocation: 2

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter.

The **offered-total-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-total-cir** mode is most useful when the policer is not receiving trusted in-profile or out-of-profile traffic and both high- and low-priority classifications are not being used on the untrusted packets and the offered packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile or out-of-profile packets and does not prevent the use of priority high or low classifications on the untrusted packets.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 122: Ingress Accounting Statistics collected in offered-total-cir stat-mode .

Table 122: Ingress Accounting Statistics collected in offered-total-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-limited-profile-cir

Counter resource allocation: 3

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed so that the traffic entering the policer comprises of hard out and uncolored. The offered counters cover traffic explicitly profiled to out-of-profile and traffic that has not been explicitly profiled at ingress (uncolored). The traffic explicitly profiled to in-profile is counted with the uncolored traffic.

The **offered-limited-profile-cir** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-limited-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile (profile out but no profile in) traffic and untrusted packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile packets. If trusted in-profile packets are not being received, the **offered-limited-profile-cir** is preferred over **offered-profile-cir** because it uses a reduced number of stat resources.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 123: Ingress Accounting Statistics Collected in offered-limited-profile-cir stat-mode .

Table 123: Ingress Accounting Statistics Collected in offered-limited-profile-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-profile-capped-cir

Counter resource allocation: 5

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed so that the traffic entering the policer comprises of hard in/out and uncolored. The offered counters cover traffic explicitly profiled to in-profile, traffic explicitly profiled to out-of-profile, and traffic that has not been explicitly profiled at ingress (uncolored).

When **offered-profile-capped-cir** is defined, the system creates five offered-output counters in the forwarding plane and five discard counters in the traffic manager.

The **offered-profile-capped-cir** mode is similar to the offered-profile-cir mode except that it includes support for profile in and **soft-in-profile** that may be output as 'out-of-profile' due to enabling **profile-capped** mode on the ingress policer.

The impact of using **offered-profile-capped-cir** stat-mode while **profile-capped** mode is disabled is that one of the counting resources in the forwarding plane and traffic manager will not be used.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 124: Ingress Accounting Statistics Collected in offered-profile-capped-cir stat-mode .

Table 124: Ingress Accounting Statistics Collected in offered-profile-capped-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-limited-capped-cir

Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when ingress reclassification is performed resulting in the traffic entering the policer comprising of hard in/out and uncolored. The offered counters cover in-profile traffic and traffic that has not been explicitly profiled at ingress (uncolored). The traffic explicitly profiled to out-of-profile is counted with the uncolored traffic.

offered-limited-capped-cir is defined, the system creates four forwarding plane offered-output counters in the network processor and four discard counters in the traffic manager.

The **offered-limited-capped-cir** mode is similar to the **offered-profile-capped-cir** mode except that it combines **soft in-profile** with **profile in** (InProf) and **profile out** (OutProf) with **soft-out-of-profile** (Uncolor) and eliminates the 'offered undefined' statistic. If trusted out-of-profile packets are not being received, the **offered-limited-capped-cir** is preferred over **offered-profile-capped-cir** because it uses a reduced number of stat resources.

This mode is intended to be used with **profile-capped** configured within the policer.

The impact of using **offered-limited-capped-cir** stat-mode while **profile-capped** mode is disabled is that one of the counting resources in the forwarding plane and traffic manager will not be used.

The counters displayed in the show output and those collected when **collect-stats** is enabled (the actual fields collected depends on the **record** configured in the applied accounting policy) are described in Table 125: Ingress Accounting Statistics Collected in offered-limited-capped-cir stat-mode .

Table 125: Ingress Accounting Statistics Collected in offered-limited-capped-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

Platforms

7705 SAR Gen 2

stat-mode

Syntax

stat-mode {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-profile-cir** | **offered-total-cir** | **offered-limited-capped-cir** | **offered-profile-capped-cir** | **offered-total-cir-exceed** | **offered-four-profile-no-cir** | **offered-total-cir-four-profile**}

no stat mode

Context

[Tree] (cfg>qos>qgrps>egr>qgrp>policer stat-mode)

Full Context

configure qos queue-group-templates egress queue-group policer stat-mode

Description

The **sap-egress** QoS policy's policer **stat-mode** command is used to configure the forwarding plane counters that allow offered, forwarded, and dropped accounting to occur for the policer. An egress policer has multiple types of offered packets (soft in-profile and out-of-profile from ingress and hard in-profile, out-of-profile, and exceed-profile due to egress profile overrides) and each of these offered types is interacting with the policer's metering and profiling functions resulting in colored output packets (green, yellow, and red). Due to the potential large number of egress policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers, for example, will not be configured with a CIR profiling rate and not all policers will receive explicitly reprofiled offered packets. The **stat-mode** command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a **no-stats** mode is supported that prevents any packet accounting, the use of the policer's **parent** command requires that the policer's **stat-mode** to be set at least to the **minimal** setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated. When a policer has been made a child to a parent policer, the **stat-mode** cannot be changed to **no-stats** unless the policer parenting is first removed.

Each time the policer's **stat-mode** is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. The total/allocated/free stats can be viewed by using the **tools dump resource-usage card fp** command. If insufficient counters exist to implement a mode on any policer instance, the **stat-mode** change will fail and the previous mode will continue unaffected for all instances of the policer.

The ingress policer stat-modes are described in Table 126: Egress Policer Stat Mode Summary.

Table 126: Egress Policer Stat Mode Summary

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
no-stats	0	None	None	—
Minimal	1	Single counter entering policer	Single counter for dropped/forwarded exiting policer	—
offered-profile-no-cir	2	In/out entering policer	In/out entering policer	Intended for when the policer does not change the profile of packets. Includes only in- and out-of-profile.
offered-profile-cir	4	In/out/uncolored (that corresponds to in- or out-of-profile from the ingress processing) entering policer	In/out exiting policer	Intended for when the policer can change the profile of packets to in- and out-of-profile.
offered-total-cir	2	Single counter entering policer	In/out exiting policer	—
offered-limited-capped-cir	4	In/out entering policer	In/out exiting policer	Intended for when the policer has **profile-capped** configured. The information is limited compared to **offered-profile-capped-cir** with the benefit of using one less stat resource.
offered-profile-capped-cir	5	In/out/uncolored (that corresponds to in- or out-of-profile from the ingress processing) entering policer	In/out exiting policer	Intended for when the policer has **profile-capped** configured.
offered-total-cir-exceed	3	Single counter entering policer	In/out/exceed exiting policer	Intended for when the policer is configured with **enable-exceed-pir**

Stat Mode	Stat Resources	Traffic Counters (Packet/Octets)		Comments
		Offered	Dropped/ Forwarded	
				to forward packets that exceed its configured PIR or when traffic is egress reclassified to profile exceed.
offered-four-profile-no-cir	4	Inplus/in/out/exceed entering policer	Inplus/in/out/exceed entering policer	Intended to be used when the policer does not change the profile of the packets and traffic is egress reclassified to profile inplus and/ or exceed.
offered-total-cir-four-profile	4	Single counter entering policer	Inplus/in/out/exceed exiting policer	Intended to be used when the policer can change the profile of the packet and traffic is egress reclassified to profile inplus.

The default **stat-mode** when a policer is created within the policy is **minimal**.

The **stat-mode** setting defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the **stat-mode** override command will fail. The previous **stat-mode** setting active for the policer will continue to be used by the policer.

The **no** form of this command attempts to return the policer's **stat-mode** setting to **minimal**. The command will fail if insufficient policer counter resources exist to implement **minimal** where the QoS policer is currently applied and has a forwarding class mapping.

Parameters

no-stats

Counter resource allocation: 0

The policer does not have any forwarding plane counters allocated and cannot provide offered, discard, and forward statistics. A policer using **no-stats** cannot be a child to a parent policer and the policer's **parent** command will fail.

When **collect-stats** is enabled, no statistics are generated.

minimal

Counter resource allocation: 1

This stat-mode provides the minimal accounting resource usage and counter information, and includes only the total offered, dropped and forwarded packet and octet counters for traffic entering (offered) and exiting (dropped/forwarded) the policer.

The default **stat-mode** for a policer is **minimal**. The **minimal** mode allocates one forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types and do not count different profile output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate or using exceed PIR.

The counters displayed in the **show** output and those collected when **collect-stats** is enabled are described in Table 127: Egress Accounting Statistics Collected in minimal stat-mode (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 127: Egress Accounting Statistics Collected in minimal stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. All	apd	AllPacketsDropped
	aod	AllOctetsDropped
For. All	apf	AllPacketsForwarded
	aof	AllOctetsForwarded

offered-profile-no-cir

Counter resource allocation: 2

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering the policer. inplus-profile traffic is counted with the in-profile counters and exceed-profile traffic is counted with the out-of-of profile counters.

The **offered-profile-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-profile-no-cir** mode is most useful when profile-based offered, dropped and forwarded statistics are required from the egress policer, but a CIR or **enable-exceed-pir** is not being used to recolor the soft in-profile and out-of-profile packets. This mode does not prevent the policer from being configured with a CIR rate or using **enable-exceed-pir**.

This mode is intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in Table 128: Egress Accounting Statistics Collected in offered-

profile-no-cir stat-mode (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 128: Egress Accounting Statistics Collected in offered-profile-no-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	**Field Description**
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-profile-cir

Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer comprises of hard inplus/in/out/exceed and soft in/out. The offered counters cover traffic reclassified to in-profile (that includes traffic reclassified to inplus-profile), traffic reclassified to out-of-profile (that includes traffic reclassified to exceed-profile), and traffic that has not been reclassified at egress (Uncolor). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

The **offered-profile-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-profile-cir** mode is most useful when profile-based offered, dropped and forwarded stats are required from the egress policer and a CIR rate is being used to recolor the soft in-profile and out-of-profile packets.

This mode is intended to be used without **profile-capped** or **enable-exceed-pir** configured within the policer as these could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in Table 129: Egress Accounting Statistics Collected in offered-profile-cir stat-mode (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 129: Egress Accounting Statistics Collected in offered-profile-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	**Field Description**
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-total-cir

Counter resource allocation: 2

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

The **offered-total-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-total-cir** mode is most useful when the policer is not receiving trusted in-profile or out-of-profile traffic and both high- and low-priority classifications are not being used on the untrusted packets and the offered packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile or

out-of-profile packets and does not prevent the use of priority high or low classifications on the untrusted packets.

This mode is intended to be used without **profile-capped** or **enable-exceed-pir** configured within the policer as these could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in Table 130: Egress Accounting Statistics Collected in offered-total-cir stat-mode (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 130: Egress Accounting Statistics Collected in offered-total-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	**Field Description**
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-limited-capped-cir

Counter resource allocation: 4

This stat-mode provides offered, dropped and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer comprises of hard inplus/in/out/exceed and soft in/out. The offered counters cover in-profile traffic (that includes traffic reclassified to inplus-profile) and out-of-profile traffic (that includes traffic reclassified to exceed-profile). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

When **offered-limited-capped-cir** is defined, the system creates four forwarding plane offered-output counters in the network processor and three discard counters in the traffic manager.

The **offered-limited-capped-cir** mode is similar to the **offered-profile-capped-cir** mode except that it combines **soft-in-profile** with **profile in** and **soft-out-of-profile** with **profile out** and eliminates the offered-undefined statistic.

The impact of using **offered-limited-capped-cir** stat-mode while profile-capped mode is disabled is that one of the counting resources in the forwarding plane and traffic manager will not be used and **soft-in-profile** will be treated as offered-in instead of offered-undefined.

This mode is intended to be used with **profile-capped** configured within the policer but without **enable-exceed-pir** configured as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in Table 131: Egress Accounting Statistics Collected in offered-limited-capped-cir stat-mode (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 131: Egress Accounting Statistics Collected in offered-limited-capped-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	**Field Description**
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded

offered-profile-capped-cir

Counter resource allocation: 5

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer when egress reclassification is performed so that the traffic entering the policer is comprised of hard inplus, hard in, hard out, and hard exceed, as well as soft in and soft out. The offered counters cover traffic reclassified to in-profile (that includes traffic

reclassified to inplus-profile), traffic reclassified to out-of-profile (that includes traffic reclassified to exceed-profile), and traffic that has not been reclassified at egress (uncolor). In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter and exceed-profile traffic is counted with the out-of-profile counter.

When **offered-profile-capped-cir** is defined, the system creates five offered-output counters in the forwarding plane and five discard counters in the traffic manager.

The **offered-profile-capped-cir** mode is similar to the **offered-profile-cir** mode except that it includes support for **profile inplus**, **profile in**, and **soft-in-profile** that may be output as out-of-profile due to enabling **profile-capped** mode on the ingress policer.

The impact of using **offered-profile-capped-cir** stat-mode while **profile-capped** mode is disabled is that one of the counting resources in the forwarding plane and traffic manager will not be used and soft-in-profile will be treated as offered-in (hard in-profile) instead of offered-undefined (uncolored).

This mode is intended to be used with **profile-capped** configured within the policer but without **enable-exceed-pir** configured as this could cause the traffic profile to be modified by the policer in a way that is not accounted for in the statistics.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in Table 132: Egress Accounting Statistics Collected in offered-profile-capped-cir stat-mode (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 132: Egress Accounting Statistics Collected in offered-profile-capped-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	**Field Description**
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. Uncolor	ucp	UncoloredPacketsOffered
	uco	UncoloredOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded

Show Output	Accounting Stats Collected	
	Field	Field Description
	oof	OutOfProfileOctetsForwarded

offered-total-cir-exceed

Counter resource allocation: 3

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. In the dropped and forwarded counters, inplus-profile traffic is counted with the in-profile counter. The **offered-total-cir-exceed** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-total-cir-exceed** mode is similar to the **offered-total-cir** mode except that it includes support for forwarded and dropped counters for **profile exceed**.

This mode is intended to be used when the policer is configured with **enable-exceed-pir** to forward packets that exceed its configured PIR or when traffic is egress reclassified to profile exceed. The mode gives the forwarded and dropped counters per profile (in, out, exceed). It is also intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer. This stat-mode is not supported for dynamic policers.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in Table 133: Egress Accounting Statistics Collected in offered-total-cir-exceed stat-mode (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 133: Egress Accounting Statistics Collected in offered-total-cir-exceed stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
Dro. ExcProf	xpd	ExceedProfilePktsDropped
	xod	ExceedProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded

Show Output	Accounting Stats Collected	
	Field	Field Description
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded
For. ExcProf	xpf	ExceedProfilePktsForwarded
	xof	ExceedProfileOctetsForwarded

offered-four-profile-no-cir

Counter resource allocation: 4

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering the policer. Offered, dropped, and forwarded counters are provided for inplus-profile, in-profile, out-of-profile, and exceed-profile traffic.

The **offered-four-profile-no-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-four-profile-no-cir** mode is similar to the **offered-profile-no-cir** mode except that it includes support for offered, dropped and forwarded counters for both profile inplus and profile exceed.

This mode is intended to be used when traffic is egress reclassified to profile inplus and/or exceed. It is also intended to be used without **profile-capped** configured within the policer as it could cause the traffic profile to be modified by the policer. This stat-mode is not supported for dynamic policers.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in Table 134: Egress Accounting Statistics Collected in offered-four-profile-no-cir stat-mode (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 134: Egress Accounting Statistics Collected in offered-four-profile-no-cir stat-mode

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InProf	ipo	InProfilePacketsOffered
	ioo	InProfileOctetsOffered
Off. OutProf	opo	OutOfProfilePacketsOffered
	ooo	OutOfProfileOctetsOffered
Off. ExcProf	xpo	ExceedProfilePacketsOffered
	xoo	ExceedProfileOctetsOffered

Show Output	Accounting Stats Collected	
	Field	Field Description
Off. InplusProf	ppo	InplusProfilePacketsOffered
	poo	InplusProfileOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
Dro. ExcProf	xpd	ExceedProfilePktsDropped
	xod	ExceedProfileOctetsDropped
Dro. InplusProf	ppd	InplusProfilePktsDropped
	pod	InplusProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded
For. ExcProf	xpf	ExceedProfilePktsForwarded
	xof	ExceedProfileOctetsForwarded
For. InplusProf	ppf	InplusProfilePktsForwarded
	pof	InplusProfileOctetsForwarded

offered-total-cir-four-profile

Counter resource allocation: 4

This stat-mode provides offered, dropped, and forwarded packet and octet counters corresponding to the profile of traffic entering (offered) and exiting (dropped/forwarded) the policer. All offered traffic is provided in a single counter. There is a separate dropped and forwarded counter for inplus, in, out, and exceed-profile traffic.

The **offered-total-cir-four-profile** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-total-cir-four-profile** mode is similar to the **offered-total-cir** except that it includes support for forwarded and dropped counters for both inplus-profile and exceed-profile.

This mode is intended to be used when traffic is egress reclassified to inplus-profile. It is also intended to be used without **profile-capped** configured within the policer as it could

cause the traffic profile to be modified by the policer. This stat-mode is not supported for dynamic policers.

The counters displayed in the show output and those collected when **collect-stats** is enabled are described in Table 135: Egress Accounting Statistics Collected in offered-total-cir-four-profile stat-mode (the actual fields collected depends on the **record** configured in the applied accounting policy).

Table 135: Egress Accounting Statistics Collected in offered-total-cir-four-profile stat-mode

Show Output	Accounting Stats Collected	
	Field	**Field Description**
Off. All	apo	AllPacketsOffered
	aoo	AllOctetsOffered
Dro. InProf	ipd	InProfilePacketsDropped
	iod	InProfileOctetsDropped
Dro. OutProf	opd	OutOfProfilePacketsDropped
	ood	OutOfProfileOctetsDropped
Dro. ExcProf	xpd	ExceedProfilePktsDropped
	xod	ExceedProfileOctetsDropped
Dro. InprofProf	ppd	InplusProfilePktsDropped
	pod	InplusProfileOctetsDropped
For. InProf	ipf	InProfilePacketsForwarded
	iof	InProfileOctetsForwarded
For. OutProf	opf	OutOfProfilePacketsForwarded
	oof	OutOfProfileOctetsForwarded
For. ExcProf	xpf	ExceedProfilePktsForwarded
	xof	ExceedProfileOctetsForwarded
For. InplusProf	ppf	InplusProfilePktsForwarded
	pof	InplusProfileOctetsForwarded

Platforms

7705 SAR Gen 2

29 s Commands – Part III

29.1 state

state

Syntax

state *state*

no state

Context

[Tree] (config>router>policy-options>policy-statement>entry>from state)

Full Context

configure router policy-options policy-statement entry from state

Description

This command identifies in resilient gateways which routes are associated with an active context and which routes are associated with a standby context.

Default

no state

Parameters

state

Specifies the state.

Values srrp-master — This is used in non-CUPS BNG resiliency to identify routes associated with an active SRRP instance.

srrp-non-master — This is used in non-CUPS BNG resiliency to identify routes associated with a standby SRRP instance.

ipsec-master-with-peer — This is used in stateful multi-chassis IPsec (MC-IPsec) redundancy to identify routes associated with an active MC-IPsec node with a reachable peer.

ipsec-non-master — This is used in stateful MC-IPsec redundancy to identify routes associated with a standby MC-IPsec node.

ipsec-master-without-peer — This is used in stateful MC-IPsec redundancy to identify routes associated with an active MC-IPsec node without a reachable peer.

fsg-active — This is used in BNG CUPS inter-BNG-UP resiliency to identify routes associated with an FSG on the active BNG-UP. This covers all session-related routes, including framed routes, IPv6 gateway addresses, and aggregated routes. It does not include loopback addresses.

fsg-active-path-restoration — This is used in BNG CUPS inter-BNG-UP resiliency to identify routes associated with an FSG on an active BNG-UP that is in a headless state. This covers all session-related routes, including framed routes, IPv6 gateway addresses, and aggregated routes. It does not include loopback addresses.

fsg-standby — This is used in BNG CUPS inter-BNG-UP resiliency to identify routes associated with an FSG on the standby BNG-UP. This covers all session-related routes, including framed routes, IPv6 gateway addresses, and aggregated routes. It does not include loopback addresses.

Platforms

7705 SAR Gen 2

29.2 state-timer

state-timer

Syntax

state-timer *seconds* [**action** *action*]

no state-timer

Context

[Tree] (config>router>pcep>pcc state-timer)

Full Context

configure router pcep pcc state-timer

Description

This command configures the state timer for PCE-initiated LSPs. The state timer must be set to a value greater than the redelegation timer.

The **no** form of the command sets this value to the default.

Default

state-timer 180 action remove

Parameters

seconds

Specifies the number of seconds before the state timer expires.

Values 1 to 3600

action

Specifies the actions that are taken on undelegated LSPs upon the state timer expiration.

Values remove, none

Default remove

Platforms

7705 SAR Gen 2

29.3 static

static

Syntax

static

Context

[Tree] (config>service>vpls>sap>mld-snooping static)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping static)

[Tree] (config>service>vpls>mesh-sdp>mld-snooping static)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping static)

[Tree] (config>service>vpls>sap>igmp-snooping static)

[Tree] (config>service>vpls>spoke-sdp>mld-snooping static)

Full Context

configure service vpls sap mld-snooping static

configure service vpls spoke-sdp igmp-snooping static

configure service vpls mesh-sdp mld-snooping static

configure service vpls mesh-sdp igmp-snooping static

configure service vpls sap igmp-snooping static

configure service vpls spoke-sdp mld-snooping static

Description

Commands in this context configure static group addresses. Static group addresses can be configured on a SAP or SDP. When present, either as a (*, g) or a (s,g) entry, multicast packets matching the configuration are forwarded even if no join message was registered for the specific group.

Platforms

7705 SAR Gen 2

static

Syntax

static *ip-address ieee-address*

no static *ip-address*

Context

[Tree] (config>service>vpls>proxy-arp static)

Full Context

configure service vpls proxy-arp static

Description

This command configures static entries to be added to the table. A static MAC-IP entry requires the addition of the MAC address to the FDB as either learned or CStatic (conditional static MAC) in order to become active.

Parameters

ip-address

Specifies the IPv4 address for the static entry.

ieee-address

Specifies a 48-bit MAC address in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx represents a hexadecimal number.

Platforms

7705 SAR Gen 2

static

Syntax

static *ipv6-address ieee-address* {**host** | **router**}

no static *ipv6-address*

Context

[Tree] (config>service>vpls>proxy-nd static)

Full Context

configure service vpls proxy-nd static

Description

This command configures static entries to be added to the table. A static MAC-IP entry requires the addition of the MAC address to the FDB as either dynamic or CStatic (Conditional Static MAC) in order to become active. Along with the IPv6 and MAC, the entry must also be configured as either host or router. This will determine if the received NS for the entry will be replied with the R flag set to 1 (router) or 0 (host).

Parameters

ipv6-address

Specifies the IPv6 address for the static entry.

ieee-address

Specifies a 48-bit MAC address in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx represents a hexadecimal number.

host

Specifies that the entry is type "host”.

router

Specifies that the entry is type "router”.

Platforms

7705 SAR Gen 2

static

Syntax

static

Context

[Tree] (config>service>vprn>igmp>if static)

Full Context

configure service vprn igmp interface static

Description

This command tests forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without receiving membership reports from host members.

Platforms

7705 SAR Gen 2

static

Syntax

static

Context

[Tree] (config>service>vprn>mld>if static)

Full Context

configure service vprn mld interface static

Description

This command tests multicast forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without receiving membership reports from host members.

Platforms

7705 SAR Gen 2

static

Syntax

static

Context

[Tree] (config>service>vprn>pim>rp static)

Full Context

configure service vprn pim rp static

Description

This command enables access to the context to configure a static rendezvous point (RP) of a PIM-SM protocol instance.

Platforms

7705 SAR Gen 2

static

Syntax

static

Context

[Tree] (config>router>igmp>if static)

Full Context

configure router igmp interface static

Description

This command tests multicast forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without receiving membership reports from host members.

Platforms

7705 SAR Gen 2

static

Syntax

static

Context

[Tree] (config>router>mld>if static)

Full Context

configure router mld interface static

Description

This command tests multicast forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without receiving membership reports from host members.

Platforms

7705 SAR Gen 2

static

Syntax

static

Context

[Tree] (config>router>pim>rp static)

[Tree] (config>router>pim>rp>ipv6 static)

Full Context

configure router pim rp static

configure router pim rp ipv6 static

Description

Commands in this context configure static Rendezvous Point (RP) addresses for a multicast group range.

Entries can be created or destroyed. If no IP addresses are configured in the **config>router>pim>rp>static>address** context, then the multicast group to RP mapping is derived from the RP-set messages received from the Bootstrap Router.

Platforms

7705 SAR Gen 2

static

Syntax

static *microseconds*

no static

Context

[Tree] (config>router>if>if-attribute>delay static)

Full Context

configure router interface if-attribute delay static

Description

This command configures the unidirectional link delay. By default there is no configured delay, the link delay metric TLV is pruned in the IGP.

The **no** form of this command removes the configured unidirectional link delay.

Default

no static

Parameters

microseconds

Specifies the unidirectional link delay in microseconds.

Values 1 to 16777214

Platforms

7705 SAR Gen 2

static

Syntax

[**no**] **static**

Context

[Tree] (config>router>bgp>next-hop-res>use-leaked-routes static)

Full Context

configure router bgp next-hop-resolution use-leaked-routes static

Description

This command configures the router to resolve any non-leaked, unlabeled unicast IPv4 or IPv6 route in the base router BGP RIB by using a static route with direct next hops leaked from any VPRN instance. A BGP route resolved this way cannot resolve other routes (including BGP routes) and cannot be redistributed into non-BGP protocols, such as IGP.

The **no** form of this command prevents the use of leaked static routes to resolve BGP routes of the base router.

Default

no static

Platforms

7705 SAR Gen 2

static

Syntax

[**no**] **static**

Context

[Tree] (config>service>vprn>bgp>next-hop-res>use-leaked-routes static)

Full Context

configure service vprn bgp next-hop-resolution use-leaked-routes static

Description

This command configures the router to resolve any non-leaked, unlabeled unicast IPv4 or IPv6 route in the VPRN BGP RIB by using a static route with direct next hops leaked from the GRT. A BGP route resolved this way cannot resolve other routes (including BGP routes) and cannot be redistributed into non-BGP protocols, such as IGP.

The **no** form of this command prevents the use of leaked static routes to resolve BGP routes of the VPRN.

Default

no static

Platforms

7705 SAR Gen 2

29.4 static-arp

static-arp

Syntax

static-arp *ieee-mac-address* **unnumbered**

static-arp *ip-address ieee-mac-address*

no static-arp [*ieee-mac-address*] **unnumbered**

no static-arp *ip-address* [*ieee-mac-address*]

Context

[Tree] (config>service>ies>if static-arp)

[Tree] (config>service>vprn>if static-arp)

Full Context

configure service ies interface static-arp

configure service vprn interface static-arp

Description

This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP appears in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of this command removes a static ARP entry.

Parameters

ip-address

Specifies the IP address for the static ARP in IP address dotted decimal notation.

ieee-mac-address

Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff, where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

unnumbered

Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. Once this command is configured, it overrides any dynamic ARP.

Platforms

7705 SAR Gen 2

static-arp

Syntax

static-arp *ieee-mac-addr* **unnumbered**

static-arp *ip-address ieee-mac-address*

no static-arp [*ieee-mac-addr*] **unnumbered**

no static-arp *ip-address* [*ieee-mac-address*]

Context

[Tree] (config>service>vpls>interface static-arp)

Full Context

configure service vpls interface static-arp

Description

This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of this command removes a static ARP entry.

Parameters

ip-address

Specifies the IP address for the static ARP in dotted decimal notation

ieee-mac-address

Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

unnumbered

Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. Once this command is configured, it overrides any dynamic ARP.

Platforms

7705 SAR Gen 2

static-arp

Syntax

static-arp *ip-address ieee-mac-address*

no static-arp *ip-address*

Context

[Tree] (config>service>vprn>nw-if static-arp)

Full Context

configure service vprn network-interface static-arp

Description

This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP will appear in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface. If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of this command removes a static ARP entry.

Parameters

ip-address

Specifies the IP address for the static ARP in IP address dotted decimal notation.

ieee-mac-address

Specifies the 48-bit MAC address for the static ARP in the form *aa*:*bb*:*cc*:*dd* :*ee*:*ff* or *aa-bb-cc-dd -ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

Platforms

7705 SAR Gen 2

static-arp

Syntax

static-arp *ip-address ieee-address*

no static-arp *ip-address*

static-arp *ieee-address* **unnumbered**

no static-arp unnumbered

Context

[Tree] (config>router>if static-arp)

Full Context

configure router interface static-arp

Description

This command configures a static Address Resolution Protocol (ARP) entry associating an IP address with a MAC address for the core router instance. This static ARP appears in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a specific IP address already exists and a new MAC address is configured for the IP address, the existing MAC address is replaced by the new MAC address.

The number of static-arp entries that can be configured on a single node is limited to 1000.

Static ARP is used when a router needs to know about a device on an interface that cannot or does not respond to ARP requests. Therefore, the router configuration can state that if it has a packet that has a certain IP address to send it to the corresponding ARP address. Use proxy ARP so the router responds to ARP requests on behalf of another device.

The **no** form of this command removes a static ARP entry.

Parameters

ieee-address

Specifies the 48-bit MAC address for the static ARP in the form *aa*:*bb*:*cc*:*dd* :*ee*:*ff* or *aa*-*bb*-*cc*-*dd* -*ee*-*ff*, where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

unnumbered

Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. Once this command is configured, it overrides any dynamic ARP.

Platforms

7705 SAR Gen 2

29.5 static-blackhole-first

static-blackhole-first

Syntax

[**no**] **static-blackhole-first**

Context

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel static-blackhole-first)

Full Context

configure service vprn bgp-ipvpn mpls auto-bind-tunnel static-blackhole-first

Description

This command configures the router to use a modified next-hop resolution sequence for each imported VPN-IP route. The router first checks for a static route in the Base routing table that matches the BGP next-hop address. If at least one such static route exists, and the route that is the longest match of the BGP next-hop address is a blackhole static route, the router resolves the VPN-IP route and programs it into the VPRN IP FIB table with a next-hop action that discards all matching packets. If there is no matching static route, or the longest matching static route is not a blackhole, the router resolves the VPN-IP route in the Base routing table as normal, that is, according to the configured VPRN auto-bind filter options.

The **no** form of this command configures the router to resolve VPN-IP routes in the Base routing table according to the configured VPRN auto-bind filter options.

Default

no static-blackhole-first

Platforms

7705 SAR Gen 2

29.6 static-cak

static-cak

Syntax

[**no**] **static-cak**

Context

[Tree] (config>macsec>connectivity-association static-cak)

Full Context

configure macsec connectivity-association static-cak

Description

This command allows the configuration of a Connectivity Association Key (CAK). The CAK is responsible for managing the MKA.

Platforms

7705 SAR Gen 2

29.7 static-entry

static-entry

Syntax

static-entry *ip-prefix/ip-prefix-length* **upto** *prefix-length2* **origin-as** *as-number* [{**valid** | **invalid**}]

no static-entry *ip-prefix/ip-prefix-length* **upto** *prefix-length2* **origin-as** *as-number*

Context

[Tree] (config>router>origin-validation static-entry)

Full Context

configure router origin-validation static-entry

Description

This command configures a static VRP entry indicating that a specific origin AS is either valid or invalid for a specific IP prefix range. Static VRP entries are stored along with dynamic VRP entries (learned from local cache servers using the RPKI-Router protocol) in the origin validation database of the router. This database is used for determining the **origin-validation** state of IPv4 and/or IPv6 BGP routes received over sessions with the **enable-origin-validatio**n command configured.

Static entries can only be configured under the **config**>**router**>**origin-validation** context of the base router.

Parameters

ip-prefix/ip-prefix-length

Specifies an IPv4 or IPv6 address with a minimum prefix length value.

Values 60 to 3600

prefix-length2

Specifies the maximum prefix length.

Values 1 to 128

as-number

Specifies as-number.

Values 0 to 4294967295

valid

Specifies a keyword meaning the static entry expresses a valid combination of origin AS and prefix range.

invalid

Specifies a keyword meaning the static entry expresses an invalid combination of origin AS and prefix range.

Platforms

7705 SAR Gen 2

29.8 static-label-range

static-label-range

Syntax

static-label-range *static-range*

no static-label-range

Context

[Tree] (config>router>mpls-labels static-label-range)

Full Context

configure router mpls-labels static-label-range

Description

This command configures the range of MPLS static label values shared among static LSP, MPLS-TP LSP, and static service VC label. Once this range is configured, it is reserved and cannot be used by other protocols such as RSVP, LDP, BGP, or Segment Routing to assign a label dynamically.

Default

static-label-range 18400

Parameters

static-range

Specifies the size of the static label range in number of labels. The minimum label value in the range is 32. The maximum label value is therefore computed as {32+ static-range-1}.

Values 0 to 262112

Default 18400

Platforms

7705 SAR Gen 2

29.9 static-lsp

static-lsp

Syntax

[**no**] **static-lsp** *lsp-name*

Context

[Tree] (config>router>mpls static-lsp)

Full Context

configure router mpls static-lsp

Description

This command is used to configure a static LSP on the ingress router. The static LSP is a manually set up LSP where the nexthop IP address and the outgoing label (push) must be specified.

The **no** form of this command deletes this static LSP and associated information.

The LSP must be shutdown first in order to delete it. If the LSP is not shut down, the **no static-lsp** *lsp-name* command does nothing except generate a warning message on the console indicating that the LSP is administratively up.

Parameters

lsp-name

Specifies the name that identifies the LSP.

Values Up to 32 alphanumeric characters.

Platforms

7705 SAR Gen 2

29.10 static-lsp-fast-retry

static-lsp-fast-retry

Syntax

static-lsp-fast-retry *seconds*

no static-lsp-fast-retry

Context

[Tree] (config>router>mpls static-lsp-fast-retry)

Full Context

configure router mpls static-lsp-fast-retry

Description

This command specifies the value used as the fast retry timer for a static LSP.

When a static LSP is trying to come up, the MPLS request for the ARP entry of the LSP next-hop may fail when it is made while the next-hop is still down or unavailable. In that case, MPLS starts a retry timer before making the next request. This enhancement allows the user to configure the retry timer, so that the LSP comes up as soon as the next-hop is up.

The **no** form of this command reverts to the default.

Default

no static-lsp-fast-retry

Parameters

seconds

Specifies the value (in s), used as the fast retry timer for a static LSP.

Values 1 to 30

Platforms

7705 SAR Gen 2

29.11 static-mac

static-mac

Syntax

static-mac *ieee-mac-address* [**create**]

no static-mac *ieee-mac-address*

Context

[Tree] (config>service>vpls>spoke-sdp static-mac)

[Tree] (config>service>vpls>mesh-sdp static-mac)

[Tree] (config>service>vpls>sap static-mac)

Full Context

configure service vpls spoke-sdp static-mac

configure service vpls mesh-sdp static-mac

configure service vpls sap static-mac

Description

This command creates a remote static MAC entry in the Virtual Private LAN Service (VPLS) forwarding database (FDB) associated with the service destination point (SDP).

In a VPLS service, MAC addresses are associated with a SAP or with an SDP. MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Local and remote static MAC entries create a permanent MAC address to SDP association in the forwarding database for the VPLS instance so that MAC address is not learned on the edge device.

Note:
Static MAC definitions on one edge device are not propagated to other edge devices participating in the VPLS instance, that is, each edge device has an independent forwarding database for the VPLS.

Only one static MAC entry (local or remote) can be defined per MAC address per VPLS instance.

By default, no static MAC address entries are defined for the SDP.

The **no** form of this command deletes the static MAC entry with the specified MAC address associated with the SDP from the VPLS forwarding database.

Parameters

ieee-mac-address

Specifies the 48-bit MAC address for the static ARP in the form *aa*:*bb*:*cc*:*dd* :*ee*:*ff* or *aa-bb-cc-dd -ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

create

Keyword used to create the static MAC instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

static-mac

Syntax

static-mac

Context

[Tree] (config>service>vpls static-mac)

Full Context

configure service vpls static-mac

Description

A set of conditional static MAC addresses can be created within a VPLS supporting BGP-EVPN. Conditional Static Macs are also supported in B-VPLS with SPBs. Unless they are configured as **black-hole**, conditional Static Macs are dependent on the SAP/SDP state.

This command allows the assignment of a set of conditional Static MAC addresses to a SAP/ spoke-SDP or **black-hole**. In the FDB, the static MAC is then associated with the active SAP or spoke-SDP.

When configured in conjunction with SPBM services, Static MACs are used for PBB Epipe and I-VPLS services that may terminate external to SPBM. If this is configured under a Control B-VPLS the interface referenced will not use IS-IS for this neighbor. This may also be configured under a User B-VPLS where the corresponding interface is not supported under the Control B-VPLS.

Static MACs configured in a BGP-EVPN service are advertised as protected (EVPN will signal the MAC as protected).

Platforms

7705 SAR Gen 2

static-mac

Syntax

static-mac *ieee-address* [**create**]

no static-mac *ieee-address*

Context

[Tree] (config>service>vpls>endpoint static-mac)

Full Context

configure service vpls endpoint static-mac

Description

This command assigns a static MAC address to the endpoint. In the FDB, the static MAC is then associated with the active spoke-SDP.

Parameters

ieee-address

Specifies the static MAC address to the endpoint

Values 6-byte mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx) Cannot be all zeros

create

This keyword is mandatory while creating a static MAC

Platforms

7705 SAR Gen 2

29.12 static-policer

static-policer

Syntax

[**no**] **static-policer** *policer-name* [**create**]

Context

[Tree] (config>sys>security>dist-cpu-protection>policy static-policer)

Full Context

configure system security dist-cpu-protection policy static-policer

Description

Configures a static enforcement policer that can be referenced by one or more protocols in the policy. Once this **policer-name** is referenced by a protocol, then this policer will be instantiated for each object (for example, a SAP or network interface) that is created and references this policy. If there is no policer resource available on the associated card or fp then the object is be blocked from being created. Multiple protocols can use the same **static-policer**.

Parameters

policy-name

Specifies the name of the policy, up to 32 characters.

Platforms

7705 SAR Gen 2

29.13 static-policy

static-policy

Syntax

static-policy *name* [**create**]

no static-policy *name*

Context

[Tree] (conf>router>segment-routing>sr-policies static-policy)

Full Context

configure router segment-routing sr-policies static-policy

Description

This command creates a context to configure a segment routing policy. The resulting segment routing policy is targeted for local installation or propagation by BGP to another router.

The **no** form of this command deletes the statically defined segment routing policy.

Default

no static-policy

Parameters

name

Specifies the name assigned to the statically defined segment routing policy, up to 64 characters.

create

Keyword used to create the policy.

Platforms

7705 SAR Gen 2

29.14 static-route

static-route

Syntax

[**no**] **static-route** *route-name*

Context

[Tree] (config>service>pw-routing static-route)

Full Context

configure service pw-routing static-route

Description

This command configures a static route to a next hop S-PE or T-PE. Static routes may be configured on either S-PEs or T-PEs.

A default static route is entered as follows:

static-route 0:0:next_hop_ip_addresss

or

static-route 0:0.0.0.0:next_hop_ip_address

The **no** form of this command removes a previously configured static route.

Parameters

route-name

Specifies the static pseudowire route.

Values

route-name	<global-id>:<prefix>:<next-hop-ip_addr>
global-id	0 to 4294967295
prefix	a.b.c.d | 0 to 4294967295
next-hop-ip_addr	a.b.c.d

Platforms

7705 SAR Gen 2

static-route

Syntax

[**no**] **static-route** *ip-prefix/ip-prefix-length* **next-hop** *ip-address*

Context

[Tree] (bof static-route)

Full Context

bof static-route

Description

This command creates a **static route** entry for the CPM management Ethernet port in the running configuration and the Boot Option File (BOF).

This command allows manual configuration of static routing table entries. These static routes are only used by traffic generated by the CPM Ethernet port. To reduce configuration, manual address aggregation should be applied where possible.

A maximum of 10 static routes can be configured on the CPM port.

The **no** form of this command deletes the static route.

Default

no static-route

Parameters

ip-prefix/ip-prefix-length

Specifies the destination address of the static route in dotted decimal notation.

Values

ip-prefix/ip-prefix-length	ipv4-prefix	a.b.c.d (host bits must be 0)
	ipv4-prefix-le	0 to 32
	ipv6-prefix	*x:x:x:x:x:x:x:x* (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x: [0 to FFFF]H
		d: [0to 255]D
	ipv6-prefix-le	0 to128
ip-address	ipv4-address	a.b.c.d
	ipv6-address	*x:x:x:x:x:x:x:x* (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x: [0 to FFFF]H
		d: [0 to 255]D

mask

Specifies the subnet mask, expressed as an integer or in dotted decimal notation.

Values 1 to 32 (mask length), 128.0.0.0 to255.255.255.255 (dotted decimal)

ip-address

Specifies the next hop IP address used to reach the destination.

Platforms

7705 SAR Gen 2

29.15 static-route-entry

static-route-entry

Syntax

static-route-entry *ip-prefixIprefix-length* [**mcast**]

no static-route-entry *ip-prefixIprefix-length* [**mcast**]

Context

[Tree] (config>service>vprn static-route-entry)

Full Context

configure service vprn static-route-entry

Description

This command creates a static route entry for both the network and access routes. A prefix and netmask must be specified.

Once the static route context for the specified prefix and netmask has been created, additional parameters associated with the static route(s) may be specified through the inclusion of additional static-route parameter commands.

The **no** form of this command deletes the static route entry. If a static route needs to be removed when multiple static routes exist to the same destination, then as many parameters to uniquely identify the static route must be entered.

Default

No static routes are defined.

Parameters

ip-prefix/prefix-length

The destination address of the static route.

Values	ipv4-prefix	a.b.c.d (host bits must be 0)
	ipv4-prefix-length	0 to 32
	ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x: [0 to FFFF]H
		d: [0 to 255]D
	ipv6-prefix-length	0 to 128

mcast

Specifies that the associated static route should be populated in the associated VPRN multicast route table.

Platforms

7705 SAR Gen 2

static-route-entry

Syntax

[**no**] **static-route-entry** *ip-prefix*/*prefix-length* [**mcast**]

Context

[Tree] (config>router static-route-entry)

Full Context

configure router static-route-entry

Description

This command creates a static route entry for both the network and access routes. A prefix and netmask must be specified.

After the static route context for the specified prefix and netmask has been created, additional parameters associated with the static routes may be specified through the inclusion of additional static route parameter commands.

The **no** form of this command deletes the static route entry. If a static route needs to be removed when multiple static routes exist to the same destination, then as many parameters to uniquely identify the static route must be entered.

Default

No static routes are defined.

Parameters

ip-prefix/prefix-length

Specifies the destination address of the static route.

Values	ipv4-prefix	a.b.c.d (host bits must be 0)	
	ipv4-prefix-length	0 to 32	
	ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)	
		x:x:x:x:x:x:d.d.d.d	
		x	[0 to FFFF]H
		d	[0 to 255]D
	ipv6-prefix-length	0 to 128	

ip-address

Specifies the IP address of the IP interface. The *ip-addr* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values

ipv4-address	a.b.c.d (host bits must be 0)
ipv6-address	x:x:x:x:x:x:x:x[-interface]
	x:x:x:x:x:x:d.d.d.d[-interface]
	x: [0..FFFF]H
	d: [0..255]D
	interface: 32 characters maximum, mandatory for link local addresses

mcast

Indicates that static route being configured is used for multicast table only.

Platforms

7705 SAR Gen 2

29.16 static-route-hold-down

static-route-hold-down

Syntax

static-route-hold-down initial *initial* **multiplier** *multiplier* **max-value** *max-value*

no static-route-hold-down

Context

[Tree] (config>router static-route-hold-down)

Full Context

configure router static-route-hold-down

Description

This command enables the hold down time feature globally for static routes in the system.

The static route hold-down time is a mechanism to protect from rapid, fluctuating state changes of static routes resulting from issues with reachability because of link flap.

This command applies to all static routes in the VPRN and the base router instance in which this hold-down time is configured.

The **no** form of this command disables the hold down time feature globally for static routes in the system.

Default

no static-route-hold-down

Parameters

initial

Specifies the initial value of the hold down time, in seconds, globally for static routes in the system.

Values 1 to 65535

multiplier

Specifies the multiplier value of the hold down time feature globally for static routes in the system.

Values 1 to 10

max-value

Specifies the maximum value of the hold down time, in seconds, globally for static routes in the system.

Values 1 to 65535

Platforms

7705 SAR Gen 2

29.17 static-sa

static-sa

Syntax

static-sa *sa-name* [**create**]

no static-sa *sa-name*

Context

[Tree] (config>ipsec static-sa)

Full Context

configure ipsec static-sa

Description

This command configures an IPsec static SA.

Platforms

7705 SAR Gen 2

29.18 static-tunnel-redundant-next-hop

static-tunnel-redundant-next-hop

Syntax

static-tunnel-redundant-next-hop *ip-address*

no static-tunnel-redundant-next-hop

Context

[Tree] (config>service>vprn>if static-tunnel-redundant-next-hop)

[Tree] (config>service>ies>if static-tunnel-redundant-next-hop)

Full Context

configure service vprn interface static-tunnel-redundant-next-hop

configure service ies interface static-tunnel-redundant-next-hop

Description

This command specifies redundant next-hop address on public or private IPsec interface (with public or private tunnel-sap) for static IPsec tunnel. The specified next-hop address will be used by standby node to shunt traffic to master in case of it receives them. Refer to the *7705 SAR Gen 2 Multiservice ISA and ESA Guide* for information about IPsec commands and descriptions.

The next-hop address will be resolved in routing table of corresponding service.

The **no** form of this command removes the address from the interface configuration.

Parameters

ip-address

Specifies the static ISA tunnel redundant next-hop address.

Platforms

7705 SAR Gen 2

29.19 stats-collection

stats-collection

Syntax

stats-collection

Context

[Tree] (config>isa>tunnel-grp stats-collection)

Full Context

configure isa tunnel-group stats-collection

Description

Commands in this context configure ISA statistics collection parameters.

Platforms

7705 SAR Gen 2

29.20 status-verify

status-verify

Syntax

status-verify

Context

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>dyn>cert status-verify)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>dyn>cert status-verify)

[Tree] (config>service>vprn>if>sap>ipsec-tun>dyn>cert status-verify)

[Tree] (config>ipsec>trans-mode-prof>dyn>cert status-verify)

[Tree] (config>service>vprn>if>sap>ipsec-gw>cert status-verify)

[Tree] (config>service>ies>if>sap>ipsec-gw>cert status-verify)

[Tree] (config>router>if>ipsec>ipsec-tun>dyn>cert status-verify)

Full Context

configure service ies interface ipsec ipsec-tunnel dynamic-keying cert status-verify

configure service vprn interface ipsec ipsec-tunnel dynamic-keying cert status-verify

configure service vprn interface sap ipsec-tunnel dynamic-keying cert status-verify

configure ipsec ipsec-transport-mode-profile dynamic-keying cert status-verify

configure service vprn interface sap ipsec-gw cert status-verify

configure service ies interface sap ipsec-gw cert status-verify

configure router interface ipsec ipsec-tunnel dynamic-keying cert status-verify

Description

Commands in this context configure Certificate Status Verification (CSV) parameters.

Platforms

7705 SAR Gen 2

status-verify

Syntax

status-verify default-result {revoked | good}

no status-verify

Context

[Tree] (config>system>security>tls>client-tls-profile status-verify)

[Tree] (config>system>security>tls>server-tls-profile status-verify)

Full Context

configure system security tls client-tls-profile status-verify

configure system security tls server-tls-profile status-verify

Description

This command configures the certificate revocation status verification parameters for end-entity (EE) certificates in the TLS client or server. This configuration overrides the existing revocation check policy.

By default the router checks the certification revocation status, but if this command is set to **good**, the end-entity certificate revocation status is overwritten and a good revocation status is returned for the EE certificate.

If this command is set to **revoked**, the router returns the actual revocation status of the end-entity certificate.

The **no** form of this command returns the actual revocation status to that of the end entity certificate.

Default

status-verify default-result revoked

Parameters

good

Specifies that the certificate is considered acceptable.

revoked

Specifies that the certificate is considered revoked.

Platforms

7705 SAR Gen 2

29.21 sticky-dest

sticky-dest

Syntax

sticky-dest *hold-time-up*

sticky-dest no-hold-time-up

no sticky-dest

Context

[Tree] (config>filter>redirect-policy sticky-dest)

[Tree] (config>filter>ip-filter>entry sticky-dest)

[Tree] (config>filter>ipv6-filter>entry sticky-dest)

Full Context

configure filter redirect-policy sticky-dest

configure filter ip-filter entry sticky-dest

configure filter ipv6-filter entry sticky-dest

Description

This command configures sticky destination behavior for redundant PBR/PBF actions. Configuring sticky destination has an effect on PBR/PBF actions whether a secondary action is configured.

The *hold-time-up* parameter allows the operator to delay programming of a PBR/PBF action for a specified amount of time. The timer is only started when transitioning from all configured targets being down (that is, the primary target if no secondary target is configured, or both the primary and secondary targets when both are configured) to at least one target being up.

When the timer expires, the primary PBR/PBF action is programmed if its target is up. If the primary PBR/PBF target is down and a secondary PBR/PBF action has been configured and its target is up, then this secondary PBR/PBF action is programmed. In all other cases, no specific programming occurs when the timer expires.

When sticky destination is configured and the secondary PBR/PBF target is up and its associated action is programmed, it is not automatically replaced by the primary PBR/PBF action when its target transitions from down to up. In this situation, programming the primary PBR/PBF action can be forced using the **activate-primary-action** tools command.

Changing the value of the timer while the timer is running takes effect immediately (that is, the timer is restarted immediately using the new value).

The **no** form of the command disables sticky destination behavior.

Default

no sticky-dest

Parameters

hold-time-up

Specifies the initial delay in seconds. Zero is equivalent to **no-hold-time-up** (no delay).

Values 0 to 65535 seconds

Platforms

7705 SAR Gen 2

29.22 sticky-dr

sticky-dr

Syntax

sticky-dr [**priority** *dr-priority*]

no sticky-dr

Context

[Tree] (config>service>vprn>pim>if sticky-dr)

Full Context

configure service vprn pim interface sticky-dr

Description

This command enables sticky-dr operation on this interface. When enabled, the priority in PIM hellos sent on this interface when elected as the designated router (DR) is modified to the value configured in *dr-priority*. This is done to avoid the delays in forwarding caused by DR recovery, when switching back to the old DR on a LAN when it comes back up.

By enabling **sticky-dr** on this interface, it will continue to act as the DR for the LAN even after the old DR comes back up.

The **no** form of this command disables sticky-dr operation on this interface.

Default

no sticky-dr

Parameters

priority *dr-priority*

Sets the DR priority to be sent in PIM Hello messages following the election of that interface as the DR, when sticky-dr operation is enabled.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

sticky-dr

Syntax

sticky-dr [**priority** *dr-priority*]

no sticky-dr

Context

[Tree] (config>router>pim>interface sticky-dr)

Full Context

configure router pim interface sticky-dr

Description

This command enables sticky-dr operation on this interface. When enabled, the priority in PIM hellos sent on this interface when elected as the designated router (DR) will be modified to the value configured in *dr-priority*. This is done to avoid the delays in forwarding caused by DR recovery, when switching back to the old DR on a LAN when it comes back up.

By enabling **sticky-dr** on this interface, it will continue to act as the DR for the LAN even after the old DR comes back up.

The **no** form of this command disables sticky-dr operation on this interface.

Default

no sticky-dr

Parameters

priority *dr-priority*

Sets the DR priority to be sent in PIM Hello messages following the election of that interface as the DR, when sticky-dr operation is enabled.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

29.23 sticky-ecmp

sticky-ecmp

Syntax

sticky-ecmp

no sticky-ecmp

Context

[Tree] (config>router>policy-options>policy-statement>default-action sticky-ecmp)

[Tree] (config>router>policy-options>policy-statement>entry>action sticky-ecmp)

Full Context

configure router policy-options policy-statement default-action sticky-ecmp

configure router policy-options policy-statement entry action sticky-ecmp

Description

This command specifies that BGP routes matching an entry or default-action of a route policy should be tagged internally as requiring sticky ECMP behavior. When a BGP route with multiple equal-cost BGP next-hops is programmed for sticky ECMP the failure of one or more of its BGP next-hops causes only the affected traffic flows to be re-distributed to the remaining next-hops; by default (without sticky-ECMP) all flows are potentially affected, even those using a next-hop that did not fail.

Default

no sticky-ecmp

Platforms

7705 SAR Gen 2

29.24 stp

stp

Syntax

stp

Context

[Tree] (config>service>vpls stp)

[Tree] (config>service>template>vpls-sap-template stp)

[Tree] (config>service>vpls>spoke-sdp stp)

[Tree] (config>service>template>vpls-template stp)

[Tree] (config>service>vpls>sap stp)

Full Context

configure service vpls stp

configure service template vpls-sap-template stp

configure service vpls spoke-sdp stp

configure service template vpls-template stp

configure service vpls sap stp

Description

Commands in this context configure the Spanning Tree Protocol (STP) parameters. Nokia's STP is simply the Spanning Tree Protocol (STP) with a few modifications to better suit the operational characteristics of VPLS services. The most evident change is to the root bridge election. Since the core network operating between Nokia's service routers should not be blocked, the root path is calculated from the core perspective.

Platforms

7705 SAR Gen 2

stp

Syntax

[**no**] **stp**

Context

[Tree] (debug>service>id stp)

Full Context

debug service id stp

Description

Commands in this context debug STP.

The **no** form of the command disables debugging.

Platforms

7705 SAR Gen 2

stp

Syntax

stp

Context

[Tree] (config>service>pw-template stp)

Full Context

configure service pw-template stp

Description

Commands in this context configure the Spanning Tree Protocol (STP) parameters. The STP is simply the Spanning Tree Protocol (STP) with a few modifications to better suit the operational characteristics of VPLS services. The most evident change is to the root bridge election. Since the core network operating between service routers should not be blocked, the root path is calculated from the core perspective.

Platforms

7705 SAR Gen 2

29.25 streaming

streaming

Syntax

streaming

Context

[Tree] (config>system>snmp streaming)

Full Context

configure system snmp streaming

Description

This command enables the proprietary SNMP request/response bundling and TCP-based transport mechanism for optimizing network management of the router nodes. In higher latency networks, synchronizing router MIBs from network management via streaming takes less time than synchronizing via classic SNMP UDP requests. Streaming operates on TCP port 1491 and runs over IPv4 or IPv6.

Platforms

7705 SAR Gen 2

29.26 strict-adjacency-check

strict-adjacency-check

Syntax

[**no**] **strict-adjacency-check**

Context

[Tree] (config>service>vprn>isis strict-adjacency-check)

Full Context

configure service vprn isis strict-adjacency-check

Description

This command enables strict checking of address families (IPv4 and IPv6) for IS-IS adjacencies. When enabled, adjacencies do not come up unless both routers have exactly the same address families configured. If there is an existing adjacency with unmatched address families, it is torn down.

This command is used to prevent black-holing traffic when IPv4 and IPv6 topologies are different. When disabled (**no strict-adjacency-check**) a BFD session failure for either IPv4 or IPv6 will cause the routes for the other address family to be removed as well.

When disabled (**no strict-adjacency-check**), both routers only need to have one common address family to establish the adjacency.

Default

no strict-adjacency-check

Platforms

7705 SAR Gen 2

strict-adjacency-check

Syntax

[**no**] **strict-adjacency-check**

Context

[Tree] (config>router>isis strict-adjacency-check)

Full Context

configure router isis strict-adjacency-check

Description

This command enables strict checking of address families (IPv4 and IPv6) for IS-IS adjacencies. When enabled, adjacencies will not come up unless both routers have exactly the same address families configured. If there is an existing adjacency with unmatched address families, it will be torn down. This command is used to prevent black-holing traffic when IPv4 and IPv6 topologies are different. When disabled (no strict-adjacency-check) a BFD session failure for either IPv4 or Ipv6 will cause the routes for the other address family to be removed as well.

When disabled (**no strict-adjacency-check**), both routers only need to have one common address family to establish the adjacency.

Platforms

7705 SAR Gen 2

29.27 strict-ero-nhop-direct-resolution

strict-ero-nhop-direct-resolution

Syntax

[**no**] **strict-ero-nhop-direct-resolution**

Context

[Tree] (config>router>mpls strict-ero-nhop-direct-resolution)

Full Context

configure router mpls strict-ero-nhop-direct-resolution

Description

This command enables the strict Explicit Route Object (ERO) next-hop direct resolution. The feature restricts the routes used to resolve the next hop of an ERO address to local and host routes. This command avoids using a next hop over a parallel link when a half link is up in the routing table.

When enabled, this command applies to an ERO when all of the following conditions are met:

- the ERO next hop is an IPv4 address
- the ERO object is a strict hop
- the IPv4 address matches the primary subnet of a local numbered interface

An ERO that meets the preceding conditions restricts resolution of the next hop to a LOCAL or a HOST route. If no such route exists, RSVP rejects the PATH message with ErrCode = Routing Error (24) and SubErrCode = Bad Strict Node (2).

The **no** form of this command disables the strict ERO next-hop direct resolution.

Default

no strict-ero-nhop-direct-resolution

Platforms

7705 SAR Gen 2

29.28 strict-lsa-checking

strict-lsa-checking

Syntax

[**no**] **strict-lsa-checking**

Context

[Tree] (config>service>vprn>ospf>graceful-restart strict-lsa-checking)

[Tree] (config>service>vprn>ospf3>graceful-restart strict-lsa-checking)

Full Context

configure service vprn ospf graceful-restart strict-lsa-checking

configure service vprn ospf3 graceful-restart strict-lsa-checking

Description

This command indicates whether an OSPF restart helper should terminate graceful restart when there is a change to an LSA that would be flooded to the restarting router during the restart process.

The default OSPF behavior is to terminate a graceful restart if an LSA changes, which causes the OSPF neighbor to go down.

The **no strict-lsa-checking** command disables strict LSA checking.

Default

strict-lsa-checking

Platforms

7705 SAR Gen 2

strict-lsa-checking

Syntax

[**no**] **strict-lsa-checking**

Context

[Tree] (config>router>ospf3>graceful-restart strict-lsa-checking)

[Tree] (config>router>ospf>graceful-restart strict-lsa-checking)

Full Context

configure router ospf3 graceful-restart strict-lsa-checking

configure router ospf graceful-restart strict-lsa-checking

Description

This command indicates whether an OSPF restart helper should terminate graceful restart when there is a change to an LSA that would be flooded to the restarting router during the restart process.

The default OSPF behavior is to terminate a graceful restart if an LSA changes, which causes the OSPF neighbor to go down.

The **no** form of this command disables strict LSA checking.

Default

strict-lsa-checking

Platforms

7705 SAR Gen 2

29.29 string

string

Syntax

string *string*

no string

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>host-ident string)

Full Context

configure subscriber-mgmt local-user-db ipoe host host-identification string

Description

This command specifies the string from the Nokia vendor-specific sub-option (VSO) in Option 82 to match when the LUDB is accessed using a DHCPv4 server.

Note:
This command is only used when **string** is configured as one of the **match-list** parameters.

The **no** form of this command removes the host identification string from the configuration.

Parameters

string

Specifies the VSO string of this host, up to 255 characters.

Platforms

7705 SAR Gen 2

string

Syntax

[**no**] **string** *text*

Context

[Tree] (config>service>vpls>sap>dhcp>option>vendor string)

[Tree] (config>service>vprn>if>dhcp>option>vendor string)

Full Context

configure service vpls sap dhcp option vendor-specific-option string

configure service vprn interface dhcp option vendor-specific-option string

Description

This command specifies the string in the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command reverts to the default.

Parameters

text

Specifies a string that can be any combination of ASCII characters, up to 32 characters. If spaces are used in the string, enclose the entire string in quotation marks (" ").

Platforms

7705 SAR Gen 2

string

Syntax

[**no**] **string** *text*

Context

[Tree] (config>router>if>dhcp>option>vendor-specific-option string)

Full Context

configure router interface dhcp option vendor-specific-option string

Description

This command specifies the vendor-specific sub-option string of the DHCP relay packet.

The **no** form of this command returns the default value.

Default

no string

Parameters

text

Specifies a string that can be any combination of ASCII characters, up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (" ").

Platforms

7705 SAR Gen 2

29.30 stub

stub

Syntax

[**no**] **stub**

Context

[Tree] (config>service>vprn>ospf>area stub)

[Tree] (config>service>vprn>ospf3>area stub)

Full Context

configure service vprn ospf area stub

configure service vprn ospf3 area stub

Description

This command enables access to the context to configure an OSPF stub area and adds/removes the stub designation from the area. External routing information is not flooded into stub areas. All routers in the stub area must be configured with the **stub** command. An OSPF area cannot be both an NSSA and a stub area. Existing virtual links of a non STUB or NSSA area will be removed when its designation is changed to NSSA or STUB.

By default, an area is not a stub area.

The **no** form of this command removes the stub designation and configuration context from the area.

Default

no stub — The area is not configured as a stub area.

Platforms

7705 SAR Gen 2

stub

Syntax

[**no**] **stub**

Context

[Tree] (config>router>ospf3>area stub)

[Tree] (config>router>ospf>area stub)

Full Context

configure router ospf3 area stub

configure router ospf area stub

Description

This command enables access to the context to configure an OSPF or OSPF3 stub area and adds/removes the stub designation from the area.

External routing information is not flooded into stub areas. All routers in the stub area must be configured with the **stub** command. An OSPF or OSPF3 area cannot be both an NSSA and a stub area.

Existing virtual links of a non STUB or NSSA area will be removed when its designation is changed to NSSA or STUB.

By default, an area is not a stub area.

The **no** form of this command removes the stub designation and configuration context from the area.

Default

no stub

Platforms

7705 SAR Gen 2

29.31 sub-port

sub-port

Syntax

sub-port *port-id* [**create**]

no sub-port *port-id*

Context

[Tree] (config>port>ethernet>dot1x>macsec sub-port)

Full Context

configure port ethernet dot1x macsec sub-port

Description

This command creates a MACsec instance on a physical port, targeting the specific subset of traffic defined by the **encap-match** command.

The **no** form of this command removes the MACsec instance.

Parameters

port-id

Specifies the sub-port id index.

Values 1 to 1023

create

Creates a new sub-port.

Platforms

7705 SAR Gen 2

29.32 subject

subject

Syntax

subject {**eq** | **neq**} *subject* [*regexp*]

no subject

Context

[Tree] (config>service>vprn>log>filter>entry>match subject)

Full Context

configure service vprn log filter entry match subject

Description

This command adds an event subject as a match criterion.

The subject is the entity for which the event is reported, such as a port. In this case the port-id string would be the subject. Only one **subject** command can be entered per event filter entry. The latest **subject** command overwrites the previous command.

The **no** form of this command removes the subject match criterion.

Default

no subject

Parameters

eq | neq

This operator specifies the type of match. Valid operators are listed below.

Values

Table 136: Valid Operators

Operator	Notes
eq	equal to
neq	not equal to

subject

A string used as the subject match criterion.

regexp

Specifies the type of string comparison to use to determine if the log event matches the value of **subject** command parameters. When the **regexp** keyword is specified, the string in the **subject** command is a regular expression string that will be matched against the subject string in the log event being filtered.

When **regexp** keyword is not specified, the **subject** command string is matched exactly by the event filter.

Platforms

7705 SAR Gen 2

subject

Syntax

subject {**eq** | **neq**} *subject* [*regexp*]

no subject

Context

[Tree] (config>log>filter>entry>match subject)

Full Context

configure log filter entry match subject

Description

This command adds an event subject as a match criterion.

The subject is the entity for which the event is reported, such as a port. In this case the port-id string would be the subject. Only one **subject** command can be entered per event filter entry. The latest **subject** command overwrites the previous command.

The **no** form of this command removes the subject match criterion.

Parameters

eq | neq

Specifies the match type. Valid operators are listed in Table 137: Valid Operators.

Table 137: Valid Operators

Operator	Notes
eq	equal to
neg	not equal to

subject

Specifies a string up to 32 characters, used as the subject match criterion.

regexp

Specifies the type of string comparison to use to determine if the log event matches the value of **subject** command parameters. When the **regexp** keyword is specified, the string in the **subject** command is a regular expression string that will be matched against the subject string in the log event being filtered. When the **regexp** keyword is not specified, the **subject** command string is matched exactly by the event filter.

Platforms

7705 SAR Gen 2

29.33 subnet

subnet

Syntax

subnet {*ip-address*/*mask* | *ip-address netmask*} [**create**]

no subnet {*ip-address*/*mask* | *ip-address netmask*}

Context

[Tree] (config>service>vprn>dhcp>server>pool subnet)

[Tree] (config>router>dhcp>server>pool subnet)

Full Context

configure service vprn dhcp local-dhcp-server pool subnet

configure router dhcp local-dhcp-server pool subnet

Description

This command creates a subnet of IP addresses to be served from the pool. The subnet cannot include any addresses that were assigned to subscribers without those addresses specifically excluded. When the subnet is created, no IP addresses are made available until a range is defined.

The **no** form of the removes the subnet parameters from the configuration.

Parameters

ip-prefix/mask

Specifies the address prefix and mask. A mask of 255.255.255.255 is reserved for system IP addresses.

Values ip-prefix: a.b.c.d

mask: 8 to 32

netmask

Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.

Values a.b.c.d, any mask expressed as dotted quad

create

Keyword used to create the subnet. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

29.34 subnet-check

subnet-check

Syntax

[**no**] **subnet-check**

Context

[Tree] (config>service>vprn>igmp>if subnet-check)

Full Context

configure service vprn igmp interface subnet-check

Description

This command enables subnet checking for IGMP messages received on this interface. All IGMP packets with a source address that is not in the local subnet are dropped.

The **no** form of this command disables local subnet checking for IGMP.

Platforms

7705 SAR Gen 2

subnet-check

Syntax

[**no**] **subnet-check**

Context

[Tree] (config>router>igmp>if subnet-check)

Full Context

configure router igmp interface subnet-check

Description

This command enables subnet checking for IGMP messages received on this interface. All IGMP packets with a source address that is not in the local subnet are dropped.

Default

subnet-check

Platforms

7705 SAR Gen 2

29.35 subnet-mask

subnet-mask

Syntax

subnet-mask *ip-address*

no subnet-mask

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>options subnet-mask)

[Tree] (config>router>dhcp>server>pool>subnet>options subnet-mask)

Full Context

configure subscriber-mgmt local-user-db ipoe host options subnet-mask

configure router dhcp local-dhcp-server pool subnet options subnet-mask

Description

This command specifies the subnet-mask option to the client. The mask can either be defined (for supernetting) or taken from the pool address.

The **no** form of this command removes the address from the configuration.

Parameters

ip-address

Specifies the IP address of the subnet mask. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

Values a.b.c.d

Platforms

7705 SAR Gen 2

29.36 subscriber-limit

subscriber-limit

Syntax

subscriber-limit *limit*

no subscriber-limit

Context

[Tree] (config>service>vprn>nat>outside>pool subscriber-limit)

Full Context

configure service vprn nat outside pool subscriber-limit

Description

This command configures the maximum number of subscribers per outside IP address.

If multiple port blocks per subscriber are used, the block size is typically small; all blocks assigned to a given subscriber belong to the same IP address; the subscriber limit guarantees that any subscriber can get a minimum number of ports.

The subscribers are counted per protocol (UDP, TCP and ICMP). For example, in LSN44 a source IPv4 address that uses ports on each of the three protocols (UDP, TCP and ICMP) on an outside IP address count as 3 subscribers on that outside IP address. The ‘no subscriber-limit’ removes the limit for the number of subscribers per outside IP address.

This command is not applicable to pools with:

- arbitrary address pooling enabled
- flexible port allocations (application configured under a pool)

Parameters

limit

Specifies the maximum number of subscribers per outside IP address.

Values 1 to 65535

Platforms

7705 SAR Gen 2

29.37 subscriber-mgmt

subscriber-mgmt

Syntax

subscriber-mgmt

Context

[Tree] (config subscriber-mgmt)

Full Context

configure subscriber-mgmt

Description

Commands in this context configure subscriber management entities. A subscriber is uniquely identified by a subscriber identification string. Each subscriber can have several DHCP sessions active at any time. Each session is referred to as a subscriber host and is identified by its IP address and MAC address.

All subscriber hosts belonging to the same subscriber are subject to the same hierarchical QoS (HQoS) processing. The HQoS processing is defined in the sub-profile (the subscriber profile). A sub-profile refers to an existing scheduler policy (configured in **the config>qos>scheduler-policy** context) and offers the possibility to overrule the rate of individual schedulers within this policy.

Because all subscriber hosts use the same scheduler policy instance, they must all reside on the same complex.

Platforms

7705 SAR Gen 2

29.38 subscription

subscription

Syntax

subscription *percentage*

no subscription

Context

[Tree] (config>router>rsvp>interface subscription)

Full Context

configure router rsvp interface subscription

Description

This command configures the percentage of the link bandwidth that RSVP can use for reservation and sets a limit for the amount of over-subscription or under-subscription allowed on the interface.

When the **subscription** is set to zero, no new sessions are permitted on this interface. If the *percentage* is exceeded, the reservation is rejected and a log message is generated.

The **no** form of this command reverts the *percentage* to the default value.

Default

subscription 100

Parameters

percentage

Specifies the percentage of the interface's bandwidth that RSVP allows to be used for reservations.

Values 0 to 1000

Platforms

7705 SAR Gen 2

subscription

Syntax

subscription *subscription-id* **cancel**

subscription cancel-all

Context

[Tree] (admin>system>telemetry>grpc subscription)

Full Context

admin system telemetry grpc subscription

Description

This command cancels an active telemetry subscription.

Parameters

subscription-id

Specifies the ID of the telemetry subscription to cancel.

Values 0 to 4294967295

Platforms

7705 SAR Gen 2

subscription

Syntax

subscription *name* [**create**]

no subscription *name*

Context

[Tree] (config>system>telemetry>persistent-subscriptions subscription)

Full Context

configure system telemetry persistent-subscriptions subscription

Description

Commands in this context configure persistent subscription commands.

The **no** form of this command removes the configuration.

Parameters

name

Specifies the subscription name, up to 32 characters.

create

Keyword used to create the subscription.

Platforms

7705 SAR Gen 2

29.39 suggest-internal-objects

suggest-internal-objects

Syntax

[**no**] **suggest-internal-objects**

Context

[Tree] (environment suggest-internal-objects)

Full Context

environment suggest-internal-objects

Description

This command enables suggesting of internally created objects while auto completing.

The **no** form of the command disables the command.

Platforms

7705 SAR Gen 2

29.40 summaries

summaries

Syntax

[**no**] **summaries**

Context

[Tree] (config>service>vprn>ospf>area>stub summaries)

[Tree] (config>service>vprn>ospf3>area>nssa summaries)

[Tree] (config>service>vprn>ospf>area>nssa summaries)

Full Context

configure service vprn ospf area stub summaries

configure service vprn ospf3 area nssa summaries

configure service vprn ospf area nssa summaries

Description

This command enables sending summary (type 3) advertisements into a stub area or Not So Stubby Area (NSSA) on an Area Border Router (ABR). This parameter is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub or nssa area. By default, summary route advertisements are sent into the stub area or NSSA.

The **no** form of this command disables sending summary route advertisements and, for stub areas, only the default route is advertised by the ABR.

Default

summaries — Summary routes are advertised by the ABR into the stub area or NSSA.

Platforms

7705 SAR Gen 2

summaries

Syntax

[**no**] **summaries**

Context

[Tree] (config>router>ospf3>area>nssa summaries)

[Tree] (config>router>ospf>area>nssa summaries)

[Tree] (config>router>ospf>area>stub summaries)

[Tree] (config>router>ospf3>area>stub summaries)

Full Context

configure router ospf3 area nssa summaries

configure router ospf area nssa summaries

configure router ospf area stub summaries

configure router ospf3 area stub summaries

Description

This command enables sending summary (type 3) advertisements into a stub area or Not So Stubby Area (NSSA) on an Area Border Router (ABR).

This parameter is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub or NSSA area (default: summary).

By default, summary route advertisements are sent into the stub area or NSSA.

The **no** form of this command disables sending summary route advertisements and, for stub areas; only the default route is advertised by the ABR.

Default

summaries

Platforms

7705 SAR Gen 2

29.41 summary

summary

Syntax

summary

Context

[Tree] (config>filter>log summary)

Full Context

configure filter log summary

Description

Commands in this context configure log summarization. These settings will only be taken into account when syslog is the log destination.

Platforms

7705 SAR Gen 2

summary

Syntax

summary [*ip-address*]

no summary

Context

[Tree] (debug>router>isis summary)

Full Context

debug router isis summary

Description

This command enables debugging for ISIS summary addresses.

The **no** form of the command disables the debugging.

Parameters

ip-address

When specified, only packets with the specified address are debugged.

Platforms

7705 SAR Gen 2

29.42 summary-address

summary-address

Syntax

summary-address {*ip-prefix/mask* | *ip-prefix* [*netmask*]} [*level*] [**tag** *tag*]

no summary-address {*ip-prefix/mask* | *ip-prefix* [*netmask*]}

Context

[Tree] (config>service>vprn>isis summary-address)

Full Context

configure service vprn isis summary-address

Description

This command creates summary-addresses for the specified router or VPRN instance.

Parameters

ip-prefix/mask

Specifies information for the specified IP prefix and mask length.

Values		
	ip-prefix	a.b.c.d (host bits must be 0)
	ipv4-prefix-length	0 to 32
	ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x: [0 to FFFF]H
		d: [0 to 255]D
	ipv6-prefix-length	0 to 128

netmask

The subnet mask in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

level

Specifies IS-IS level area attributes. If no level parameter is specified, the default is level-1/2.

Values level-1, level-2, level-1/2

tag *tag*

Assigns a route tag to the summary address.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

summary-address

Syntax

summary-address {*ip-prefix/ip-prefix-length* | *ip-prefix netmask*} [*level*] [**tag** *tag*] [**algorithm** *algo-id*]

summary-address {*ip-prefix/ip-prefix-length* | *ip-prefix netmask*} [*level*] [**tag** *tag*] [**algorithm** *algo-id*] **advertise-unreachable** [**match-route-tag** *tag*] [**advertise-route-tag** *tag*]

no summary-address {*ip-prefix/ip-prefix-length* | *ip-prefix netmask*}

Context

[Tree] (config>router>isis summary-address)

Full Context

configure router isis summary-address

Description

This command creates a summary IPv4, IPv6, or SRv6 locator address.

When an IS-IS domain exists out of multiple areas, the user must redistribute IP addresses and SRv6 locators between areas for inter-area SRv6-based transport services.

Scaling may be impacted if all existing IPv4, IPv6, and SRv6 locators are redistributed between all existing areas. SRv6 locators and IP addresses can be summarized when they are redistributed from one area into another area. Summarization reduces the number of entries redistributed, which reduces the size of the Link State Database (LSDB) and increases network stability.

The **no** form of this command reverts to the default.

Default

no summary-address

Parameters

ip-prefix/ip-prefix-length

Specifies the IP prefix and prefix length of the summary address.

Values

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 32
ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)

	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6-prefix-length	0 to 128

netmask

Specifies the subnet mask in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

level

Specifies IS-IS level area attributes.

Values level-1, level-2, level-1/2

Default level-1/2

tag

Specifies the route tag to assign for the summary address.

Values 1 to 4294967295

algo-id

Specifies the algorithm topology applied for the summary address.

Values 0, 128 to 255

Default 0

match-route-tag *tag*

Specifies the route tag to match the Unreachable Prefix Announcements (UPAs). This selects a subset of summary member prefixes to monitor for reachability.

Values 1 to 4294967295

Default no match-route-tag

advertise-route-tag *tag*

Specifies the route tag to advertise in the UPA. The UPA tag can be used when there are multiple ASBR redistributing prefixes between two IGP areas.

Values 1 to 4294967295

Default no advertise-route-tag

Platforms

7705 SAR Gen 2

29.43 summary-crit

summary-crit

Syntax

summary-crit dst-addr

summary-crit src-addr

no summary-crit

Context

[Tree] (config>filter>log>summary summary-crit)

Full Context

configure filter log summary summary-crit

Description

This command defines the key of the index of the mini-table. If key information is changed while summary is administratively enabled (no shutdown), the filter summary mini-table is flushed and recreated with different key information. Log packets received during the reconfiguration time will be handled as if summary was not active.

The **no** form of the command reverts to the default parameter.

Default

summary-crit src-addr

Parameters

dst-addr

Specifies that received log packets are summarized based on the destination IPv4, IPv6, or MAC address.

src-addr

Specifies that received log packets are summarized based on the source IPv4, IPv6 or MAC address.

Platforms

7705 SAR Gen 2

29.44 super-backbone

super-backbone

Syntax

[**no**] **super-backbone**

Context

[Tree] (config>service>vprn>ospf super-backbone)

Full Context

configure service vprn ospf super-backbone

Description

This command specifies whether CE-PE functionality is required or not. The OSPF super backbone indicates the type of the LSA generated as a result of routes redistributed into OSPF. When enabled, the redistributed routes are injected as summary, external or NSSA LSAs. When disabled, the redistributed routes are injected as either external or NSSA LSAs only.

Default

no super-backbone

Platforms

7705 SAR Gen 2

29.45 supplicant-timeout

supplicant-timeout

Syntax

supplicant-timeout *seconds*

no supplicant-timeout

Context

[Tree] (config>port>ethernet>dot1x supplicant-timeout)

Full Context

configure port ethernet dot1x supplicant-timeout

Description

This command configures the period during which the router waits for a client to respond to its EAPOL messages. When the supplicant-timeout expires, the 802.1x authentication session is considered to have failed.

The **no** form of this command returns the value to the default.

Default

supplicant-timeout 30

Parameters

seconds

Specifies the server timeout period in seconds.

Values 1 to 300

Platforms

7705 SAR Gen 2

29.46 suppress

suppress

Syntax

suppress *integer*

no suppress

Context

[Tree] (config>router>policy-options>damping suppress)

Full Context

configure router policy-options damping suppress

Description

This command configures the suppression parameter for the route policy damping profile.

A route is suppressed when it has flapped frequently enough to increase the Figure of Merit (FoM) value to exceed the **suppress** threshold limit. When the **FoM** value exceeds the **suppress** threshold limit, the route is removed from the route table or inclusion in advertisements.

The **no** form of this command removes the suppress parameter from the damping profile.

Default

no suppress

Parameters

integer

Specifies the suppress value expressed as a decimal integer.

Values 1 to 20000

Platforms

7705 SAR Gen 2

29.47 suppress-attached-bit

suppress-attached-bit

Syntax

[**no**] **suppress-attached-bit**

Context

[Tree] (config>service>vprn>isis suppress-attached-bit)

Full Context

configure service vprn isis suppress-attached-bit

Description

This command configures IS-IS to suppress setting the attached bit on originated Level 1 LSPs to prevent all L1 routers in the area from installing a default route to it.

Platforms

7705 SAR Gen 2

suppress-attached-bit

Syntax

[**no**] **suppress-attached-bit**

Context

[Tree] (config>router>isis suppress-attached-bit)

Full Context

configure router isis suppress-attached-bit

Description

This command configures IS-IS to suppress setting the attached bit on originated Level 1 LSPs to prevent all L1 routers in the area from installing a default route to it.

Default

no suppress-attached-bit

Platforms

7705 SAR Gen 2

29.48 suppress-dn-bit

suppress-dn-bit

Syntax

[**no**] **suppress-dn-bit**

Context

[Tree] (config>service>vprn>ospf suppress-dn-bit)

[Tree] (config>service>vprn>ospf3 suppress-dn-bit)

Full Context

configure service vprn ospf suppress-dn-bit

configure service vprn ospf3 suppress-dn-bit

Description

This command specifies whether to suppress the setting of the DN bit for OSPF LSA packets generated by this instance of OSPF on the router. When enabled, the DN bit for OSPF LSA packets generated by this instance of the OSPF router will not be set. When disabled, this instance of the OSPF router will follow the normal procedure to determine whether to set the DN bit.

Default

no suppress-dn-bit

Platforms

7705 SAR Gen 2

29.49 suppress-lsn-events

suppress-lsn-events

Syntax

[**no**] **suppress-lsn-events**

Context

[Tree] (config>isa>nat-group suppress-lsn-events)

Full Context

configure isa nat-group suppress-lsn-events

Description

This command suppresses the generation of Large Scale NAT (LSN) events when RADIUS accounting is enabled.

By default, only one logging facility for tracking subscribers in LSN44, DS-Lite, and NAT64 can be enabled at the time: either the SR OS event logging facility or the RADIUS logging facility. SR OS event logs can be sent to multiple destinations, such as the console session, a telnet or SSH session, memory logs, file destinations, SNMP trap groups, and syslog destinations.

If RADIUS logging is enabled, the NAT logs are sent to the RADIUS destination and the NAT logs are suppressed in the SR OS event logging facility, for example, NAT logs are not sent to the syslog server.

If RADIUS logging is disabled, the NAT logs are sent to the SR OS event logging facility; for example, syslog, assuming that the events are enabled via the event-control command (**configure log event-control nat** *event* **generate**).

By explicitly disabling this command (**no suppress-lsn-events**), the NAT logs can be sent to both logging facilities simultaneously, the SR OS event logging facility, and the RADIUS logging facility.

Default

suppress-lsn-events

Platforms

7705 SAR Gen 2

29.50 suppress-lsn-sub-blks-free

suppress-lsn-sub-blks-free

Syntax

[**no**] **suppress-lsn-sub-blks-free**

Context

[Tree] (config>isa>nat-group suppress-lsn-sub-blks-free)

Full Context

configure isa nat-group suppress-lsn-sub-blks-free

Description

This command suppresses the tmnxNatLsnSubBlksFree summary notification and use the tmnxNatPlBlockAllocationLsn notifications. When the SR OS node is in a state of excessive logging, the queue associated with the transmission of logs on the MS-ISA can become congested. This event further delays the generation of logs, and with this, further allocations and deallocations of NAT resources (port-blocks) will be stalled until the queue is relieved of congestion. For example, an excessive logging state in the system can be caused by issuing a command to clear a large number of NAT subscribers where a large number of resources (port-blocks) are released at once.

The **suppress-lsn-sub-blks-free** command enables the generation of individual logs carried in event-id 2012 for every released port block regardless of the state of the transmission queue (whether congested or not). If NAT subscribers have a large number of allocated port blocks (this could be hundreds of port blocks per subscriber), generating individual logs per port-block release contributes to the congestion.

To alleviate transmission queue congestion, this behavior can be changed by disabling this command (**no suppress-lsn-sub-blks-free**). This causes the suppression of logs related to the release of individual port blocks of a NAT subscriber when the transmission queue is congested. As a result, only a summarized release log via event-id 2021 for the subscriber is generated. The purpose of this new log is to inform the operator in a single message that all ports blocks for the subscriber are released. For example, the log message for LSN will be "LSN subscriber all blocks freed”. The benefit of such summarization (or log aggregation) is to alleviate the congestion of the transmission queue and consequently accelerate resource releases. An effect is the decreased granularity of information.

If summarization is enabled (**no suppress-lsn-sub-blks-free**) while there is no logging congestion in the system, the port block releases continue to be logged individually via the event-id 2012 (assuming that this is enabled in the event control), except for the last port block of the subscriber. When the last port block is released, the log with event-id 2021 is generated indicating that all port blocks for the subscriber are now released without carrying the specific information about this last port block that is released.

Default

no suppress-lsn-sub-blks-free

Platforms

7705 SAR Gen 2

29.51 suppress-standby-signaling

suppress-standby-signaling

Syntax

[**no**] **suppress-standby-signaling**

Context

[Tree] (config>service>vpls>endpoint suppress-standby-signaling)

Full Context

configure service vpls endpoint suppress-standby-signaling

Description

When this command is enabled, the pseudowire standby bit (value 0x00000020) will not be sent to T-LDP peer when the specified spoke is selected as a standby. This allows faster switchover as the traffic will be sent over this SDP and discarded at the blocking side of the connection. This is particularly applicable to multicast traffic.

Default

suppress-standby-signaling

Platforms

7705 SAR Gen 2

29.52 suppress-threshold

suppress-threshold

Syntax

suppress-threshold *suppress-penalties* **reuse-threshold** *reuse-penalties*

Context

[Tree] (config>port>ethernet>dampening suppress-threshold)

Full Context

configure port ethernet dampening suppress-threshold

Description

This command configures the penalties thresholds at which the port state events to the upper layer are dampened (suppress threshold) and then permitted (reuse threshold).

Parameters

suppress-penalties

Specifies the threshold at which the port up state is suppressed until the accumulated penalties drop below the reuse threshold again.

Values 1 to 20000

Default 2000

reuse-penalties

Specifies the threshold at which the port up state is no longer suppressed, after the port has been in a suppressed state and the accumulated penalties decay drops below this threshold. The reuse threshold value must be less than the suppress threshold value.

Values 1 to 20000

Default 1000

Platforms

7705 SAR Gen 2

29.53 svc-id

svc-id

Syntax

svc-id *service-id*

no svc-id

Context

[Tree] (config>system>security>mgmt-access-filter>mac-filter>entry>match svc-id)

Full Context

configure system security management-access-filter mac-filter entry match svc-id

Description

This command specifies an existing svc-id to use as a match condition.

Parameters

service-id

Specifies a service-id to match.

Values *service-id*: 1 to 2147483647*svc-name*: 64 characters maximum

Platforms

7705 SAR Gen 2

29.54 swap

swap

Syntax

swap {*out-label* | **implicit-null-label**} **nexthop** *ip-address*

no swap

Context

[Tree] (config>router>mpls>if>label-map swap)

Full Context

configure router mpls interface label-map swap

Description

This command swaps the incoming label and specifies the outgoing label and next hop IP address on an LSR for a static LSP.

The **no** form of this command removes the swap action associated with the *in-label*.

Parameters

implicit-null-label

Specifies the use of the implicit label value for the outgoing label of the swap operation.

out-label

Specifies the label value to be swapped with the in-label. Label values 16 through 1,048,575 are defined as follows:

- label values 16 through 31 are reserved
- label values 32 through 1,023 are available for static assignment
- label values 1,024 through 2,047 are reserved for future use
- label values 2,048 through 18,431 are statically assigned for services
- label values 28,672 through 131,071 are dynamically assigned for both MPLS and services

- label values 131,072 through 1,048,575 are reserved for future use

Values 16 to 1048575

nexthop ***ip-address***

Specifies the IP address to forward to. If an ARP entry for the next hop exists, then the static LSP will be marked operational. If ARP entry does not exist, software will set the operational status of the static LSP to down and continue to ARP for the configured nexthop. Software will continuously try to ARP for the configured nexthop at a fixed interval.

Platforms

7705 SAR Gen 2

29.55 sweep

sweep

Syntax

sweep start *dispersion-start* **end** *dispersion-end*

Context

[Tree] (config>port>dwdm>coherent sweep)

Full Context

configure port dwdm coherent sweep

Description

This command allows users to configure the dispersion sweep 'start' and 'end' values for the automatic mode of coherent control. If the user knows the approximate or theoretical residual dispersion of the link, this command can be used to limit the range of sweeping for the automatic control mode and thus achieve faster link up.

Parameters

dispersion-start

Specifies the lower range limit for the dispersion compensation.

Values -50000 to 50000

Default -25500

dispersion-end

Specifies the upper range limit for the dispersion compensation.

Values -50000 to 50000

Default 2000

Platforms

7705 SAR Gen 2

29.56 switchover-exec

switchover-exec

Syntax

switchover-exec *file-url*

no switchover-exec

Context

[Tree] (config>system switchover-exec)

Full Context

configure system switchover-exec

Description

This command specifies the location and name of the CLI script file executed following a redundancy switchover from the previously active CPM card. A switchover can happen because of a fatal failure or by manual action.

The CLI script file can contain commands for environment settings, classic CLI debug configuration (excluding mirroring settings), and other commands not maintained by the configuration redundancy.

The following commands are not supported in the switchover-exec file: clear, configure, candidate, oam, tools, oam, ping, traceroute, mstat, mtrace and mrinfo.

Default

no switch-over-exec

Parameters

file-url

Specifies the location and name of the CLI script file.

Values

local-url | *remote-url*	
local-url	[*cflash-id*/][*file-path*] 200 chars max, including cflash-id
	directory length 99 chars max each

remote-url	[{ftp:// | tftp://}*login:pswd@remote-locn/*][*file-path*]
	243 chars max
	directory length 99 chars max each
remote-locn	[*hostname* | *ipv4-address* | *ipv6-address*]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D
	interface - 32 chars max, for link local addresses
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

29.57 sync

sync

Syntax

[**no**] **sync**

Context

[Tree] (config>redundancy>multi-chassis>peer sync)

Full Context

configure redundancy multi-chassis peer sync

Description

Commands in this context configure synchronization parameters.

Default

no sync

Platforms
7705 SAR Gen 2

29.58 sync-e

sync-e

Syntax
[**no**] **sync-e**

Context
[Tree] (config>card>mda sync-e)

Full Context
configure card mda sync-e

Description
This command enables synchronous Ethernet on the MDA. Then any port on the MDA can be used as a source port in the sync-if-timing configuration.

The **no** form of this command disables synchronous Ethernet on the MDA.

Platforms
7705 SAR Gen 2

29.59 synchronize

synchronize

Syntax
synchronize {**boot-env** | **config**}

Context
[Tree] (config>redundancy synchronize)

Full Context
configure redundancy synchronize

Description

This command enables the automatic synchronization of the standby CPM's images and/or config files from the active CPM. Either the **boot-env** or **config** parameter must be specified. When the standby CPM takes over operation following a failure or reset of the active CPM, it is important to ensure that the active and standby CPMs have identical software images and configuration files. This includes the saved configuration, saved incremental configuration files in model-driven configuration mode, CPM, XCM, and IOM images.

The active CPM ensures that the active configuration is maintained on the standby CPM. However, to ensure smooth operation under all circumstances, runtime images and system initialization configurations must also be automatically synchronized between the active and standby CPM.

If synchronization fails, alarms and log messages that indicate the type of error that caused the failure of the synchronization operation are generated. When the error condition ceases to exist, the alarm is cleared.

Only files stored on the router are synchronized. If a configuration file or image is stored in a location other than on a local compact flash, the file is not synchronized (for example, storing a configuration file on an FTP server).

Default

synchronize config

Parameters

boot-env

Synchronizes all files required for the boot process (boot loader, BOF configuration, SR OS images, and all configuration files).

config

Synchronizes the primary, secondary, and tertiary configuration files, SSH keys, the password history and the model-driven commit history.

Default config

Platforms

7705 SAR Gen 2

synchronize

Syntax

synchronize cert

synchronize {**boot-env** | **config**}

Context

[Tree] (admin>redundancy synchronize)

Full Context

admin redundancy synchronize

Description

This command performs a synchronization of the standby CPM's images and/or configuration files to the active CPM. Either the **boot-env** or **config** parameter must be specified.

In the **admin>redundancy** context, this command performs a manually triggered standby CPM synchronization. When the standby CPM takes over operation following a failure or reset of the active CPM, it is important to ensure that the active and standby CPM have identical operational parameters. This includes the saved configuration, CPM, XCM, and IOM images.

The active CPM ensures that the active configuration is maintained on the standby CPM. However, to ensure smooth operation under all circumstances, runtime images and system initialization configurations must also be automatically synchronized between the active and standby CPM. If synchronization fails, alarms and log messages that indicate the type of error that caused the failure of the synchronization operation are generated. When the error condition ceases to exist, the alarm is cleared.

Only files stored on the router are synchronized. If a configuration file or image is stored in a location other than on a local compact flash, the file is not synchronized (for example, storing a configuration file on an FTP server).

The **no** form of the command removes the parameter from the configuration.

Default

no synchronize

Parameters

cert

Synchronizes the imported certificate, key, and CRL files.

boot-env

Synchronizes all files required for the boot process (boot loader, BOF, images, and configuration).

config

Synchronizes the primary, secondary, and tertiary configuration files.

Platforms

7705 SAR Gen 2

29.60 synchronous-execution

synchronous-execution

Syntax

synchronous-execution *seconds*

synchronous-execution never

Context

[Tree] (config>system>management-interface>ops>global-timeouts synchronous-execution)

Full Context

configure system management-interface operations global-timeouts synchronous-execution

Description

This command configures the period of time that operations launched as "'synchronous" (the default method for all operations) are allowed to execute before they are automatically stopped, and their associated data is deleted.

If a specific execution timeout is not included in the request for a particular synchronous operation, this system-level timeout applies.

Note:
This execution timeout is part of the general global operations infrastructure and is separate and independent from any operation-specific timeouts (for example, the **ping** operation also has its own **timeout** parameter).

Caution:
This timeout also applies to operations requested in the MD-CLI interface (for example, ping, file dir, and so on). If **synchronous-execution** is enabled with a specific time value, MD-CLI operations are subject to this timeout and are interrupted if they execute longer than the configured **synchronous-execution** time.

Default

synchronous-execution never

Parameters

seconds

Specifies the period of time, in seconds, that synchronous operations are allowed to execute.

Values 1 to 604800

never

Keyword to specify that an execution timeout is not applied to synchronous operations.

Platforms

7705 SAR Gen 2

29.61 syslog

syslog

Syntax

syslog *syslog-id* [**name** *syslog-name*]

no syslog *syslog-id*

Context

[Tree] (config>service>vprn>log syslog)

Full Context

configure service vprn log syslog

Description

This command creates the context to configure a Syslog target host that is capable of receiving selected Syslog messages from this network element.

A valid *syslog-id* must have the target Syslog host address configured.

A maximum of 30 Syslog IDs can be configured.

No log events are sent to a Syslog target address until the syslog-id has been configured as the log destination (**to**) in the log-id node.

The Syslog ID configured in the **configure**>**service**>**vprn** context has a local VPRN scope and only needs to be unique within the specific VPRN instance. The same ID can be reused under a different VPRN service or in the global log context under **config>log**.

Default

No syslog IDs are defined.

Parameters

syslog-id

Specifies the Syslog ID for the Syslog destination.

Values 1 to 30

name syslog-name

Specifies an optional Syslog name, up to 64 characters, that can be used to refer to the Syslog destination after it is created.

Platforms

7705 SAR Gen 2

syslog

Syntax

syslog *syslog-id* [**name** *syslog-name*]

no syslog *syslog-id*

Context

[Tree] (config>log syslog)

Full Context

configure log syslog

Description

Commands in this context configure a Syslog target host capable of receiving selected syslog messages from this network element.

A valid *syslog-id* must have the target Syslog host address configured.

A maximum of 10 Syslog IDs can be configured.

Log events are not sent to a Syslog target address until the *syslog-id* is configured as the log destination (**to**) in the node specified by the Log ID.

The Syslog ID configured in the **config**>**service**>**vprn** context has a local VPRN scope and only needs to be unique within the specific VPRN instance. The same ID can be reused under a different VPRN service or in the global log context under **config**>**log**.

The **no** form of this command removes the Syslog configuration.

Parameters

syslog-id

Specifies the Syslog ID for the Syslog destination.

Values 1 to 10

name syslog-name

Configures an optional Syslog name, up to 64 characters, that can be used to refer to the Syslog destination after it is created.

Platforms

7705 SAR Gen 2

29.62 system

system

Syntax

[**no**] **system**

Context

[Tree] (debug system)

Full Context

debug system

Description

This command displays system debug information.

Platforms

7705 SAR Gen 2

29.63 system-base-mac

system-base-mac

Syntax

system-base-mac *mac-address*

no system-base-mac

Context

[Tree] (bof system-base-mac)

Full Context

bof system-base-mac

Description

This command is used to specify the base MAC address. The specified MAC address is used as the first MAC address by the system to assign MAC addresses to individual interfaces.

It is strongly recommended that a unique base MAC address is assigned to each instance with a minimum gap of 1024 between base addresses to avoid a MAC address overlap.

The **no** form of this command removes the configured system base MAC address.

Default

no system-base-mac

Parameters

mac-address

Specifies the MAC address.

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

Platforms

7705 SAR Gen 2

29.64 system-filter

system-filter

Syntax

system-filter

Context

[Tree] (config>filter system-filter)

Full Context

configure filter system-filter

Description

Commands in this context activate system filter policies.

Platforms

7705 SAR Gen 2

29.65 system-id

system-id

Syntax

system-id *system-id*

no system-id

Context

[Tree] (config>subscr-mgmt>loc-user-db>ipoe>host>host-ident system-id)

Full Context

configure subscriber-mgmt local-user-db ipoe host host-identification system-id

Description

This command specifies the system ID to match for a host lookup. When the LUDB is accessed through a DHCPv4 server, the system ID is matched against the Nokia vendor specific sub-option in DHCP Option 82.

Note:
This command is only used when **system-id** is configured as one of the **match-list** parameters.

The **no** form of this command removes the system ID from the configuration.

Parameters

system-id

Specifies the system ID, up to 255 characters

Platforms

7705 SAR Gen 2

system-id

Syntax

[**no**] **system-id**

Context

[Tree] (config>service>vprn>if>dhcp>option>vendor system-id)

[Tree] (config>service>vpls>sap>dhcp>option>vendor system-id)

Full Context

configure service vprn interface dhcp option vendor-specific-option system-id

configure service vpls sap dhcp option vendor-specific-option system-id

Description

This command specifies whether the system-id is encoded in the Nokia vendor-specific sub-option of Option 82.

The **no** form of this command reverts to the default.

Platforms

7705 SAR Gen 2

system-id

Syntax

system-id *isis-system-id*

no system-id

Context

[Tree] (config>service>vprn>isis system-id)

Full Context

configure service vprn isis system-id

Description

This command configures the IS-IS system ID. The system ID has a fixed length of 6 octets; it is determined using the following preference order:

1. **config>service>vprn>isis>system-id**
2. **config>service>vprn>isis>router-id**
3. **config>service>vprn>router-id**
4. **config>service>vprn>if>address**
5. The default system ID 2550.0000.0000, based on the default router ID 255.0.0.0

The system ID is integral to IS-IS; therefore, for the **system-id** command to take effect, a **shutdown** and then **no shutdown** must be performed on the IS-IS instance. This will ensure that the configured and operational system ID are always the same.

The **no** form of this command removes the system ID from the configuration. The router ID is used when no system ID is specified.

Default

no system-id

Parameters

isis-system-id

12 hexadecimal characters in dotted-quad notation.

Values aaaa.bbbb.cccc, where aaaa, bbbb, and cccc are hexadecimal numbers

Platforms

7705 SAR Gen 2

system-id

Syntax

[**no**] **system-id**

Context

[Tree] (config>router>if>dhcp>option>vendor-specific-option system-id)

Full Context

configure router interface dhcp option vendor-specific-option system-id

Description

This command specifies whether the system-id is encoded in the Nokia vendor-specific sub-option of Option 82.

Default

no system-id

Platforms

7705 SAR Gen 2

system-id

Syntax

system-id *isis-system-id*

no system-id

Context

[Tree] (config>router>isis system-id)

Full Context

configure router isis system-id

Description

This command configures the IS-IS system ID. The system ID has a fixed length of 6 octets; it is determined using the following preference:

1. **config>router>isis>system-id**
2. **config>router>isis>router-id**
3. **config>router>router-id**
4. **config>router>interface>system> address**

5. The default system ID 2550.0000.0000, based on the default router ID 255.0.0.0

The system ID is integral to IS-IS; therefore, for the **system-id** command to take effect, the IS-IS instance must be **shutdown** and then **no shutdown**. This will ensure that the configured and operational system ID are always the same.

The **no** form of this command removes the system ID from the configuration. The router ID is used when no system ID is specified.

Parameters

isis-system-id

Specifies 12 hexadecimal characters in dotted-quad notation.

Values aaaa.bbbb.cccc, where aaaa, bbbb, and cccc are hexadecimal numbers

Platforms

7705 SAR Gen 2

29.66 system-mac

system-mac

Syntax

system-mac *mac-address*

no system-mac

Context

[Tree] (config>system>ned>profile system-mac)

Full Context

configure system network-element-discovery profile system-mac

Description

This command configures the MAC address to be advertised.

The **no** form of this command removes any explicitly defined MAC address and chassis MAC address will be advertised.

Default

no system-mac

Parameters

mac-address

Specifies the MAC address to be associated with the profile in *xx:xx:xx:xx:xx:xx* or *xx-xx-xx-xx-xx-xx* format.

Platforms

7705 SAR Gen 2

29.67 system-password

system-password

Syntax

system-password admin-password

system-password dynsvc-password

Context

[Tree] (admin>system>security system-password)

Full Context

admin system security system-password

Description

This operational command changes a local system password.

Parameters

admin-password

Specifies to change the administrative password.

dynsvc-password

Specifies to change the dynamic services password.

Platforms

7705 SAR Gen 2

29.68 system-priority

system-priority

Syntax

system-priority *value*

no system-priority

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ep system-priority)

Full Context

configure redundancy multi-chassis peer mc-endpoint system-priority

Description

This command allows the operator to set the system priority. The peer configured with the lowest value is chosen to be the master. If system-priority are equal then the one with the highest system-id (chassis MAC address) is chosen as the master.

The **no** form of this command sets the system priority to default.

Default

no system-priority

Parameters

value

Specifies the priority assigned to the local MC-EP peer.

Values 1 to 255

Platforms

7705 SAR Gen 2

29.69 system-profile

system-profile

Syntax

system-profile {**profile-a** | **profile-b**}

no system-profile

Context

[Tree] (bof system-profile)

Full Context

bof system-profile

Description

This command configures the system profile in the BOF.

See "System profiles" in the *7705 SAR Gen 2 Basic System Configuration Guide* for more information.

The **no** form of this command removes the **system-profile** parameter from the BOF.

Parameters

profile-a

Specifies that the system profile is for generic deployment scenarios, IP forwarding and MPLS switching use-cases.

profile-b

Specifies that the system profile is primarily intended for applications requiring high packet manipulation and processing (for example, NAT and IPsec).

Platforms

7705 SAR Gen 2

30 t Commands

30.1 tab

tab

Syntax

[**no**] **tab**

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>command-completion tab)

Full Context

configure system management-interface cli md-cli environment command-completion tab

Description

This command enables completion on the tab character.

The **no** form of this command reverts to the default value.

Default

tab

Platforms

7705 SAR Gen 2

30.2 table-size

table-size

Syntax

table-size *table-size*

Context

[Tree] (config>service>vpls>proxy-arp table-size)

[Tree] (config>service>vpls>proxy-nd table-size)

Full Context

configure service vpls proxy-arp table-size

configure service vpls proxy-nd table-size

Description

This command adds a table-size limit per service. By default, the table-size limit is 250; it can be set up to 16k entries per service. A non-configurable implicit high watermark of 95% and low watermark of 90% exists, per service and per system. When those watermarks are reached, a syslog/trap is triggered. When the system/service limit is reached, entries for a specified IP can be replaced (a different MAC can be learned and added) but no new IP entries will be added, regardless of the type (Static, evpn, dynamic). If the user attempts to change the **table-size** value to a value that cannot accommodate the number of existing entries, the attempt will fail.

Default

table-size 250

Parameters

table-size

Specifies the table-size as number of entries for the service.

Values 1 to 16384

Platforms

7705 SAR Gen 2

30.3 tacplus

tacplus

Syntax

no tacplus

tacplus [**create**]

Context

[Tree] (config>service>vprn>aaa>rmt-srv tacplus)

Full Context

configure service vprn aaa remote-servers tacplus

Description

This command creates the context to configure TACACS+ authentication on the VPRN.

Configure multiple server addresses for each router for redundancy.

The **no** form of this command removes the TACACS+ configuration.

Parameters

create

Keyword to create the TACACS+ configuration.

Platforms

7705 SAR Gen 2

tacplus

Syntax

[**no**] **tacplus**

Context

[Tree] (config>system>security tacplus)

Full Context

configure system security tacplus

Description

This command creates the context to configure TACACS+ authentication on the router.

Configure multiple server addresses for each router for redundancy.

The **no** form of this command removes the TACACS+ configuration.

Platforms

7705 SAR Gen 2

30.4 tacplus-map-to-priv-lvl

tacplus-map-to-priv-lvl

Syntax

tacplus-map-to-priv-lvl [*admin-priv-lvl*]

no tacplus-map-to-priv-lvl

Context

[Tree] (config>system>security>password>enable-admin-control tacplus-map-to-priv-lvl)

Full Context

configure system security password enable-admin-control tacplus-map-to-priv-lvl

Description

When **tacplus-map-to-priv-lvl** is enabled, and tacplus authorization is enabled with the *use-priv-lvl* option, typing **enable-admin** starts an interactive authentication exchange from the node to the TACACS+ server. The start message (service=enable) contains the user-id and the requested *admin-priv-lvl*. Successful authentication results in the use of a new profile (as configured under **config>system>security>tacplus>priv-lvl-map**).

Platforms

7705 SAR Gen 2

30.5 tag

tag

Syntax

tag *tag*

no tag [*tag*]

Context

[Tree] (config>service>vprn>static-route-entry tag)

Full Context

configure service vprn static-route-entry tag

Description

This command associates a 4-byte route-tag with the static route. The tag value can be used in route policies to control distribution of the static route into other protocols.

The tag specified at this level of the static route causes tag values configured under the next-hop, black-hole, and indirect contexts of the static route to be ignored.

The **no** form of this command removes the tag association.

Default

no tag

Parameters

tag

Specifies an integer value.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

tag

Syntax

tag *tag-value*

no tag [*tag-value*]

Context

[Tree] (config>service>vprn>static-route-entry>ipsec-tunnel tag)

[Tree] (config>service>vprn>static-route-entry>next-hop tag)

[Tree] (config>service>vprn>static-route-entry>indirect tag)

Full Context

configure service vprn static-route-entry ipsec-tunnel tag

configure service vprn static-route-entry next-hop tag

configure service vprn static-route-entry indirect tag

Description

This command adds a 32-bit integer tag to the associated static route.

The tag value can be used in route policies to control distribution of the route into other protocols.

Default

no tag

Parameters

tag-value

Specifies an integer tag value.

Values 32 bit integer

Platforms

7705 SAR Gen 2

tag

Syntax

tag *tag*

no tag

Context

[Tree] (config>service>vprn>isis>if tag)

Full Context

configure service vprn isis interface tag

Description

This command configures a route tag to the specified IP address of an interface.

Parameters

tag

Specifies the tag value.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

tag

Syntax

tag *tag*

no tag [*tag*]

Context

[Tree] (config>router>static-route-entry>next-hop tag)

[Tree] (config>router>static-route-entry tag)

[Tree] (config>router>static-route-entry>black-hole tag)

[Tree] (config>router>static-route-entry>indirect tag)

Full Context

configure router static-route-entry next-hop tag

configure router static-route-entry tag

configure router static-route-entry black-hole tag

configure router static-route-entry indirect tag

Description

This command associates a 4-byte route-tag with the static route. The tag value can be used in route policies to control distribution of the static route into other protocols.

The tag specified at this level of the static route causes tag values configured under the next-hop, black-hole and indirect contexts of the static route to be ignored.

The **no** form of this command removes the tag association.

Default

no tag

Parameters

tag

Specifies an integer tag value.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

tag

Syntax

tag *tag*

no tag

Context

[Tree] (config>router>isis>interface tag)

Full Context

configure router isis interface tag

Description

This command configures a route tag to the specified IP address of an interface.

The **no** form of this command removes the tag value from the configuration.

Parameters

tag

Specifies a route tag.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

tag

Syntax

tag *tag*

no tag

Context

[Tree] (config>router>isis>interface tag)

Full Context

configure router isis interface tag

Description

This command configures a route tag to the specified IP address of an interface.

The **no** form of this command removes the tag value from the configuration.

Default

no tag

Parameters

tag

Specifies a route tag.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

tag

Syntax

tag {**no-tag** | *tag*}

no tag

Context

[Tree] (config>router>policy-options>policy-statement>entry>from tag)

Full Context

configure router policy-options policy-statement entry from tag

Description

This command matches the tag value in static or IGP routes. A decimal or hexadecimal value of 4 octets can be entered. For IS-IS, OSPF, and static routes, all four octets can be used. For RIP and RIPng, only the two most significant octets are used if more than two octets are configured.

The **no** form of this command removes the tag field match criterion.

Default

no tag

Parameters

tag

Matches the configured tag value.

Values Accepts decimal or hexadecimal formats:

- IS-IS, OSPF and static routes: 0x0 – 0xFFFFFFFF or 1 – 4294967295
- RIP and RIPng: 0x0 – 0xFFFF or 1 – 65535

no-tag

Specifies that no tag value is set.

Platforms

7705 SAR Gen 2

tag

Syntax

tag *tag*

no tag

Context

[Tree] (config>router>policy-options>policy-statement>default-action tag)

[Tree] (config>router>policy-options>policy-statement>entry>action tag)

Full Context

configure router policy-options policy-statement default-action tag

configure router policy-options policy-statement entry action tag

Description

This command assigns a tag to routes matching the entry, which is then applied to IGP routes. A decimal or hexadecimal value of 4 octets can be entered.

For IS-IS and OSPF, all four octets can be used.

For RIP and RIPng, only the two most significant octets are used if more than two octets are configured.

The **no** form of this command removes the tag.

Default

no tag

Parameters

tag

Assigns an IS-IS, OSPF, RIP or RIPng tag to routes matching the entry.

Values Accepts decimal or hexadecimal formats:

IS-IS and OSPF: 0x0–0xFFFFFFFF or 1–4294967295

RIP and RIPng: 0x0–0xFFFF or 1–65535

name — The tag parameter variable name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@”.

Platforms

7705 SAR Gen 2

30.6 taii-type2

taii-type2

Syntax

taii-type2 *global-id:prefix:ac-id*

no taii-type2

Context

[Tree] (config>service>epipe>spoke-sdp-fec taii-type2)

Full Context

configure service epipe spoke-sdp-fec taii-type2

Description

taii-type2 configures the target attachment individual identifier for the SDP SDP. This is only applicable to FEC129 AII type 2.

This command is blocked in CLI if this end of the spoke SDP is configured for single-sided auto configuration (using the **auto-config** command).

Parameters

global-id

Specifies a global ID of this router T-PE. This value must correspond to one of the global_id values configured for a local-prefix under **config>service>pw-routing>local-prefix** context.

Values 1 to 4294967295

prefix

Speecifies prefix on this router T-PE that the spoke SDP SDP is associated with. This value must correspond to one of the prefixes configured under **config>service>pw-routing>local-prefix** context.

Values an IPv4-formatted address a.b.c.d or 1 to 4294967295

ac-id

Specifies an unsigned integer representing a locally unique identifier for the spoke SDP.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

30.7 target-name

target-name

Syntax

target-name {node-name | user-agent | custom-string *name***}**

no target-name

Context

[Tree] (config>system>grpc-tunnel>tunnel target-name)

Full Context

configure system grpc-tunnel tunnel target-name

Description

This command assigns a target name that the node will register with.

The **no** form of this command removes the target name.

Default

no target-name

Parameters

node-name

Keyword to register the tunnel with the node name configured using the **configure system name** command.

user-agent

Keyword to register the tunnel with the user agent name string defined as *node-name:vendor:model:software-version*.

custom-string

Assigns an arbitrary string as the target name.

name

Specifies a string, up to 64 characters, that defines the target name.

Platforms

7705 SAR Gen 2

30.8 target-power

target-power

Syntax

target-power *power*

Context

[Tree] (config>port>dwdm>coherent target-power)

Full Context

configure port dwdm coherent target-power

Description

This command configures the target transmit optical power for the port.

Default

target-power 1.00

Parameters

power

Specifies the desired average output power in dBm.

Values -20.00 to 3.00

Platforms

7705 SAR Gen 2

30.9 target-type

target-type

Syntax

target-type {grpc-server | ssh-server | custom-type *type***}**

no target-type

Context

[Tree] (config>system>grpc-tunnel>tunnel>handler target-type)

Full Context

configure system grpc-tunnel tunnel handler target-type

Description

This command assigns a server as a handler for all tunnel sessions.

The **no** form of this command disables the tunnel handler server.

Default

no target-type

Parameters

grpc-server

Keyword that assigns the gRPC server as a handler for all tunnels sessions. The gRPC-tunnel protocol value corresponds to "GNMI_GNOI”.

ssh-server

Keyword that assigns the SSH server as a handler for all tunnels sessions. The gRPC-tunnel protocol value corresponds to "SSH”.

custom-type

Keyword that assigns an arbitrary string as the target type.

type

Specifies a string, up to 255 characters, defining the client to serve as a handler for all tunnel sessions. Values used by gRPC tunnel protocol, such as "GNMI_GNOI” or "SSH” can also be used.

Platforms

7705 SAR Gen 2

30.10 targeted-session

targeted-session

Syntax

targeted-session

Context

[Tree] (config>router>ldp targeted-session)

Full Context

configure router ldp targeted-session

Description

This command configures targeted LDP sessions. Targeted sessions are LDP sessions between non-directly connected peers. Hello messages are sent directly to the peer platform instead of to all the routers on this subnet multicast address. The user can configure different default parameters for IPv4 and IPv6 LDP targeted hello adjacencies.

The discovery messages for an indirect LDP session are addressed to the specified peer and not to the multicast address.

Platforms

7705 SAR Gen 2

30.11 task

task

Syntax

task [**detail**]

no task

Context

[Tree] (debug>router>pcep>pcc task)

[Tree] (debug>router>pcep>pcc>conn task)

Full Context

debug router pcep pcc task

debug router pcep pcc connection task

Description

This command enables debugging for PCC or connection task events.

The **no** form of this command disables debugging.

Parameters

detail

Keyword used to specify detailed information about PCC or connection task events.

Platforms

7705 SAR Gen 2

30.12 tcp-ack

tcp-ack

Syntax

tcp-ack {**true** | **false**}

no tcp-ack

Context

[Tree] (config>filter>ipv6-filter>entry>match tcp-ack)

[Tree] (config>filter>ip-filter>entry>match tcp-ack)

Full Context

configure filter ipv6-filter entry match tcp-ack

configure filter ip-filter entry match tcp-ack

Description

This command configures an IP filter match criterion based on the Acknowledgment (ACK) TCP Flag bit, defined in RFC 793, as being set or not in the TCP header of an IP packet.

The **no** form of the command removes the criterion from the match entry.

Default

no tcp-ack

Parameters

true

Specifies matching on IP packets that have the selected TCP flag bit set in the TCP header.

false

Specifies matching on IP packets that do not have the selected TCP flag bit set in the TCP header.

Platforms

7705 SAR Gen 2

30.13 tcp-cwr

tcp-cwr

Syntax

tcp-cwr {**true** | **false**}

no tcp-cwr

Context

[Tree] (config>filter>ipv6-filter>entry>match tcp-cwr)

[Tree] (config>filter>ip-filter>entry>match tcp-cwr)

Full Context

configure filter ipv6-filter entry match tcp-cwr

configure filter ip-filter entry match tcp-cwr

Description

This command configures an IP filter match criterion based on the Congestion Window Reduced (CWR) TCP Flag bit, defined in RFC 3168, as being set or not in the TCP header of an IP packet.

The **no** form of the command removes the criterion from the match entry.

Default

no tcp-cwr

Parameters

true

Specifies matching on IP packets that have the selected TCP flag bit set in the TCP header.

false

Specifies matching on IP packets that do not have the selected TCP flag bit set in the TCP header.

Platforms

7705 SAR Gen 2

30.14 tcp-ece

tcp-ece

Syntax

tcp-ece {**true** | **false**}

no tcp-ece

Context

[Tree] (config>filter>ip-filter>entry>match tcp-ece)

[Tree] (config>filter>ipv6-filter>entry>match tcp-ece)

Full Context

configure filter ip-filter entry match tcp-ece

configure filter ipv6-filter entry match tcp-ece

Description

This command configures an IP filter match criterion based on the ECN-Echo (ECE) TCP Flag bit, defined in RFC 3168, as being set or not in the TCP header of an IP packet.

The **no** form of the command removes the criterion from the match entry.

Default

no tcp-ece

Parameters

true

Specifies matching on IP packets that have the selected TCP flag bit set in the TCP header.

false

Specifies matching on IP packets that do not have the selected TCP flag bit set in the TCP header.

Platforms

7705 SAR Gen 2

30.15 tcp-established

tcp-established

Syntax

tcp-established [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no tcp-established

Context

[Tree] (config>service>nat>nat-policy>timeouts tcp-established)

Full Context

configure service nat nat-policy timeouts tcp-established

Description

This command configures the idle timeout applied to a TCP session in the established state.

Default

tcp-established hrs 2 min 4

Parameters

hours

Specifies the timeout hours field.

Values 1 to 24

minutes

Specifies the timeout minutes field.

Values 1 to 59

seconds

Specifies the timeout seconds field.

Values 1 to 59

Platforms

7705 SAR Gen 2

tcp-established

Syntax

[**no**] **tcp-established**

Context

[Tree] (config>filter>ipv6-filter>entry>match tcp-established)

[Tree] (config>filter>ip-filter>entry>match tcp-established)

Full Context

configure filter ipv6-filter entry match tcp-established

configure filter ip-filter entry match tcp-established

Description

This command matches packets with the TCP flag ACK or RST.

Default

tcp-established

Platforms

7705 SAR Gen 2

30.16 tcp-fin

tcp-fin

Syntax

tcp-fin {**true** | **false**}

no tcp-fin

Context

[Tree] (config>filter>ip-filter>entry>match tcp-fin)

[Tree] (config>filter>ipv6-filter>entry>match tcp-fin)

Full Context

configure filter ip-filter entry match tcp-fin

configure filter ipv6-filter entry match tcp-fin

Description

This command configures an IP filter match criterion based on the FIN TCP Flag bit, defined in RFC 793, as being set or not in the TCP header of an IP packet.

The **no** form of the command removes the criterion from the match entry.

Default

no tcp-fin

Parameters

true

Specifies matching on IP packets that have the selected TCP flag bit set in the TCP header.

false

Specifies matching on IP packets that do not have the selected TCP flag bit set in the TCP header.

Platforms

7705 SAR Gen 2

30.17 tcp-keepalive

tcp-keepalive

Syntax

tcp-keepalive

Context

[Tree] (config>system>grpc tcp-keepalive)

Full Context

configure system grpc tcp-keepalive

Description

Commands in this context configure the sending of TCP keepalives by the router towards all gRPC clients.

Enabling TCP keepalive speeds up the detection of certain failures. The TCP keepalives sent by the router are controlled by three commands: **idle-time**, **interval**, and **retries**. The router starts sending TCP keepalives when the connection has been idle (no TCP segments sent or received) for more than **idle-time** seconds. At that point, the router sends a probe (TCP ACK with a sequence number = current sequence number - 1) and expects a TCP ACK. It repeats this probe every **interval** seconds for the configured number of **retries**. If no response is received to any of the probes, the connection is immediately closed, which starts the purge timer if the TCP connection is currently supporting the RibApi service.

Platforms

7705 SAR Gen 2

tcp-keepalive

Syntax

tcp-keepalive

Context

[Tree] (config>system>grpc-tunnel>destination-group tcp-keepalive)

[Tree] (config>system>telemetry>destination-group tcp-keepalive)

Full Context

configure system grpc-tunnel destination-group tcp-keepalive

configure system telemetry destination-group tcp-keepalive

Description

Commands in this context configure TCP keepalive commands.

Platforms

7705 SAR Gen 2

30.18 tcp-mss

tcp-mss

Syntax

tcp-mss *mss-value*

no tcp-mss

Context

[Tree] (config>service>ies>if>ipv6 tcp-mss)

[Tree] (config>service>ies>if tcp-mss)

Full Context

configure service ies interface ipv6 tcp-mss

configure service ies interface tcp-mss

Description

This command statically sets the TCP maximum segment size (MSS) for TCP connections originated from the associated IP interface to the specified value.

The **no** form of this command removes the static value and allows the TCP MSS value to be calculated based on the IP MTU value by subtracting the base IP and TCP header lengths from the IP MTU value (tcp_mss = ip_mtu – 40).

Default

no tcp-mss

Parameters

mss-value

The TCP MSS value that should be used in the TCP SYN packet during the three-way handshake negotiation of a TCP connection.

Note: 9158 = max-IP_MTU (9198)-40

Values 536 to 9746 (IPv4) 1220 to 9726 (IPv6)

Platforms

7705 SAR Gen 2

tcp-mss

Syntax

tcp-mss *mss-value*

no tcp-mss

Context

[Tree] (config>service>vprn>if tcp-mss)

[Tree] (config>service>vprn>if>ipv6 tcp-mss)

[Tree] (config>service>vprn>nw-if tcp-mss)

Full Context

configure service vprn interface tcp-mss

configure service vprn interface ipv6 tcp-mss

configure service vprn network-interface tcp-mss

Description

This command statically sets the TCP maximum segment size (MSS) for TCP connections originated from the associated IP or network interface to the specified value.

The **no** form of this command removes the static value and allows the TCP MSS value to be calculated based on the IP MTU value by subtracting the base IP and TCP header lengths from the IP MTU value (tcp_mss = ip_mtu – 40).

Default

no tcp-mss

Parameters

mss-value

Specifies the TCP MSS value that should be used in the TCP SYN packet during the three-way handshake negotiation of a TCP connection.

Note: 9746 = max-IP_MTU (9786)-40

Values 384 to 9746 (IPv4 or network)

1220 to 9726(IPv6)

Platforms

7705 SAR Gen 2

tcp-mss

Syntax

tcp-mss *mss-value*

no tcp-mss

Context

[Tree] (config>router>if>ipv6 tcp-mss)

[Tree] (config>router>if tcp-mss)

Full Context

configure router interface ipv6 tcp-mss

configure router interface tcp-mss

Description

This command statically sets the TCP maximum segment size (MSS) for TCP connections originated from the associated IP interface to the specified value.

The **no** form of this command removes the static value and allows the TCP MSS value to be calculated based on the IP MTU value by subtracting the base IP and TCP header lengths from the IP MTU value (tcp_mss = ip_mtu – 40).

Default

no tcp-mss

Parameters

mss-value

Specifies the TCP MSS value that should be used in the TCP SYN packet during the three-way handshake negotiation of a TCP connection.

9158 = max-IP_MTU (9198)-40

Values 536 to 9746 (IPv4) 1220 to 9726 (IPv6)

Platforms

7705 SAR Gen 2

tcp-mss

Syntax

tcp-mss *mss-value*

no tcp-mss

Context

[Tree] (config>service>vprn>bgp tcp-mss)

[Tree] (config>router>bgp tcp-mss)

Full Context

configure service vprn bgp tcp-mss

configure router bgp tcp-mss

Description

This command configures an override for the TCP maximum segment size to use with a specific peer or set of peers (depending on the scope of the command).

The configured value controls two properties of the TCP connection as follows:

- TCP MSS option — The router advertises the TCP MSS option value in the TCP SYN packet it sends as part of the 3-way handshake. The advertised value may be lower than the configured value, depending on the IP MTU of the first hop IP interface. The peers are asked to abide by this value when sending TCP segments to the local router.
- TCP maximum segment size — The actual transmitted size may be lower than the configured value, depending on the TCP MSS option value signaled by the peers, the effect of path MTU discovery, or other factors.

The **no** form of this command removes the TCP MSS override values from the configuration.

Default

no tcp-mss

Parameters

mss-value

Specifies the The router uses the TCP SYN to advertise the TCP MSS option value towards its peer. MSS value, in bytes, to use with the peers that fall within the scope of the command.

Values 384 to 9746

Platforms

7705 SAR Gen 2

tcp-mss

Syntax

tcp-mss ip-stack

tcp-mss *mss-value*

no tcp-mss

Context

[Tree] (config>service>vprn>bgp>group>neighbor tcp-mss)

[Tree] (config>service>vprn>bgp>group tcp-mss)

[Tree] (config>router>bgp>group tcp-mss)

[Tree] (config>router>bgp>group>neighbor tcp-mss)

Full Context

configure service vprn bgp group neighbor tcp-mss

configure service vprn bgp group tcp-mss

configure router bgp group tcp-mss

configure router bgp group neighbor tcp-mss

Description

This command configures an override for the TCP maximum segment size to use with a specific peer or set of peers (depending on the scope of the command).

The configured value controls two properties of the TCP connection as follows:

- TCP MSS option — The router advertises the TCP MSS option value in the TCP SYN packet it sends as part of the 3-way handshake. The advertised value may be lower than the configured value, depending on the IP MTU of the first hop IP interface. The peers are asked to abide by this value when sending TCP segments to the local router.
- TCP maximum segment size — The actual transmitted size may be lower than the configured value, depending on the TCP MSS option value signaled by the peers, the effect of path MTU discovery, or other factors.

The **no** form of this command removes the TCP MSS override values from the configuration.

Default

no tcp-mss

Parameters

mss-value

Specifies the TCP MSS value, in bytes, to use with the peers that fall within the scope of the command.

Values 384 to 9746

ip-stack

This keyword requests that TCP MSS be derived from mechanisms and configurations outside of BGP, including the configuration of **tcp-mss** at the IP interface level. It provides a method to override inheritance within the BGP configuration.

Platforms

7705 SAR Gen 2

30.19 tcp-mss-adjust

tcp-mss-adjust

Syntax

tcp-mss-adjust *segment-size*

no tcp-mss-adjust

Context

[Tree] (config>service>nat>nat-policy tcp-mss-adjust)

Full Context

configure service nat nat-policy tcp-mss-adjust

Description

This command configures the value to adjust the TCP Maximum Segment Size (MSS) option.

The **no** form of the command returns the segment size to the default.

Default

no tcp-mss-adjust

Parameters

segment-size

Specifies the value to put into the TCP Maximum Segment Size (MSS) option if not already present, or if the present value is higher.

Values 160 to 10240

Platforms

7705 SAR Gen 2

tcp-mss-adjust

Syntax

tcp-mss-adjust

Context

[Tree] (config>filter>ip-filter>entry>action tcp-mss-adjust)

[Tree] (config>filter>ipv6-filter>entry>action tcp-mss-adjust)

Full Context

configure filter ip-filter entry action tcp-mss-adjust

configure filter ipv6-filter entry action tcp-mss-adjust

Description

This command activates the adjustment of the TCP Maximum Segment Size (MSS) option of TCP packets matching the entry.

Platforms

7705 SAR Gen 2

30.20 tcp-ns

tcp-ns

Syntax

tcp-ns {**true** | **false**}

no tcp-ns

Context

[Tree] (config>filter>ipv6-filter>entry>match tcp-ns)

[Tree] (config>filter>ip-filter>entry>match tcp-ns)

Full Context

configure filter ipv6-filter entry match tcp-ns

configure filter ip-filter entry match tcp-ns

Description

This command configures an IP filter match criterion based on the Nonce Sum (NS) TCP Flag bit, defined in RFC 3540, as being set or not in the TCP header of an IP packet.

The **no** form of the command removes the criterion from the match entry.

Default

no tcp-ns

Parameters

true

Specifies matching on IP packets that have the selected TCP flag bit set in the TCP header.

false

Specifies matching on IP packets that do not have the selected TCP flag bit set in the TCP header.

Platforms

7705 SAR Gen 2

30.21 tcp-option-number

tcp-option-number

Syntax

tcp-option-number

Context

[Tree] (config>system>security>keychain tcp-option-number)

Full Context

configure system security keychain tcp-option-number

Description

Commands in this context configure the TCP option number to be placed in the TCP packet header.

Platforms

7705 SAR Gen 2

30.22 tcp-psh

tcp-psh

Syntax

tcp-psh {**true** | **false**}

no tcp-psh

Context

[Tree] (config>filter>ip-filter>entry>match tcp-psh)

[Tree] (config>filter>ipv6-filter>entry>match tcp-psh)

Full Context

configure filter ip-filter entry match tcp-psh

configure filter ipv6-filter entry match tcp-psh

Description

This command configures an IP filter match criterion based on the Push (PSH) TCP Flag bit, defined in RFC 793, as being set or not in the TCP header of an IP packet.

The **no** form of the command removes the criterion from the match entry.

Default

no tcp-psh

Parameters

true

Specifies matching on IP packets that have the selected TCP flag bit set in the TCP header.

false

Specifies matching on IP packets that do not have the selected TCP flag bit set in the TCP header.

Platforms

7705 SAR Gen 2

30.23 tcp-rst

tcp-rst

Syntax

tcp-rst {**true** | **false**}

no tcp-rst

Context

[Tree] (config>filter>ip-filter>entry>match tcp-rst)

[Tree] (config>filter>ipv6-filter>entry>match tcp-rst)

Full Context

configure filter ip-filter entry match tcp-rst

configure filter ipv6-filter entry match tcp-rst

Description

This command configures an IP filter match criterion based on the Reset (RST) TCP Flag bit, defined in RFC 793, as being set or not in the TCP header of an IP packet.

The **no** form of the command removes the criterion from the match entry.

Default

no tcp-rst

Parameters

true

Specifies matching on IP packets that have the selected TCP flag bit set in the TCP header.

false

Specifies matching on IP packets that do not have the selected TCP flag bit set in the TCP header.

Platforms

7705 SAR Gen 2

30.24 tcp-session-parameters

tcp-session-parameters

Syntax

tcp-session-parameters

Context

[Tree] (config>router>ldp tcp-session-parameters)

Full Context

configure router ldp tcp-session-parameters

Description

Commands in this context configure parameters applicable to TCP transport session of an LDP session to remote peer.

Platforms

7705 SAR Gen 2

30.25 tcp-syn

tcp-syn

Syntax

tcp-syn [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no tcp-syn

Context

[Tree] (config>service>nat>nat-policy>timeouts tcp-syn)

Full Context

configure service nat nat-policy timeouts tcp-syn

Description

This command configures the timeout applied to a TCP session in the SYN state.

Default

tcp-syn sec 15

Parameters

hours

Specifies the timeout hours field.

Values 1 to 24

minutes

Specifies the timeout minutes field.

Values 1 to 59

seconds

Specifies the timeout seconds field.

Values 1 to 59

Platforms

7705 SAR Gen 2

tcp-syn

Syntax

tcp-syn {**true** | **false**}

no tcp-syn

Context

[Tree] (config>filter>ipv6-filter>entry>match tcp-syn)

[Tree] (config>filter>ip-filter>entry>match tcp-syn)

Full Context

configure filter ipv6-filter entry match tcp-syn

configure filter ip-filter entry match tcp-syn

Description

This command configures an IP filter match criterion based on the Synchronize (SYN) TCP Flag bit, defined in RFC 793, as being set or not in the TCP header of an IP packet.

The **no** form of the command removes the criterion from the match entry.

Default

no tcp-syn

Parameters

true

Specifies matching on IP packets that have the selected TCP flag bit set in the TCP header.

false

Specifies matching on IP packets that do not have the selected TCP flag bit set in the TCP header.

Platforms

7705 SAR Gen 2

30.26 tcp-time-wait

tcp-time-wait

Syntax

tcp-time-wait [**min** *minutes*] [**sec** *seconds*]

no tcp-time-wait

Context

[Tree] (config>service>nat>nat-policy>timeouts tcp-time-wait)

Full Context

configure service nat nat-policy timeouts tcp-time-wait

Description

This command configures the timeout applied to a TCP session in a time-wait state.

Default

no tcp-time-wait

Parameters

minutes

Specifies the timeout minutes field.

Values 1 to 4

seconds

Specifies the timeout seconds field.

Values 1 to 59

Platforms

7705 SAR Gen 2

30.27 tcp-transitory

tcp-transitory

Syntax

tcp-transitory [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no tcp-transitory

Context

[Tree] (config>service>nat>nat-policy>timeouts tcp-transitory)

Full Context

configure service nat nat-policy timeouts tcp-transitory

Description

This command configures the idle timeout applied to a TCP session in a transitory state.

Default

tcp-transitory min 4

Parameters

hours

Specifies the timeout hours field.

Values 1 to 24

minutes

Specifies the timeout minutes field.

Values 1 to 59

seconds

Specifies the timeout seconds field.

Values 1 to 59

Platforms

7705 SAR Gen 2

30.28 tcp-urg

tcp-urg

Syntax

tcp-urg {**true** | **false**}

no tcp-urg

Context

[Tree] (config>filter>ipv6-filter>entry>match tcp-urg)

[Tree] (config>filter>ip-filter>entry>match tcp-urg)

Full Context

configure filter ipv6-filter entry match tcp-urg

configure filter ip-filter entry match tcp-urg

Description

This command configures an IP filter match criterion based on the Urgent (URG) TCP Flag bit, defined in RFC 793, as being set or not in the TCP header of an IP packet.

The **no** form of the command removes the criterion from the match entry.

Default

no tcp-urg

Parameters

true

Specifies matching on IP packets that have the selected TCP flag bit set in the TCP header.

false

Specifies matching on IP packets that do not have the selected TCP flag bit set in the TCP header.

Platforms

7705 SAR Gen 2

30.29 te

te

Syntax

[**no**] **te**

Context

[Tree] (debug>router>mpls>event te)

Full Context

debug router mpls event te

Description

This command debugs te events.

The **no** form of the command disables the debugging.

Platforms

7705 SAR Gen 2

30.30 te-class

te-class

Syntax

te-class *te-class-number* **class-type** *ct-number* **priority** *priority*

no te-class te-class-number

Context

[Tree] (config>router>rsvp>diffserv-te te-class)

Full Context

configure router rsvp diffserv-te te-class

Description

This command configures a TE class. A TE class is defined as:

TE Class = {Class Type (CT), LSP priority}

Eight TE classes are supported. There is no default TE class once Diff-Serv is enabled. The user has to explicitly define each TE class.

When Diff-Serv is disabled, there will be an internal use of the default CT (CT0) and eight pre-emption priorities as shown in Table 138: Default Class Type.

Table 138: Default Class Type

Class Type (CT internal)	LSP Priority
0	7
0	6
0	5
0	4
0	3
0	2
0	1
0	0

The **no** form of this command deletes the TE class.

Parameters

te-class *te-class-number*

Specifies the TE class number.

Values 0 to 7

class-type *ct-number*

Specifies the Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.

Values 0 to 7

priority *priority*

Specifies the LSP priority.

Values 0 to 7

Platforms

7705 SAR Gen 2

30.31 te-down-threshold

te-down-threshold

Syntax

te-down-threshold *threshold-level* [*threshold-level*]

no te-down-threshold

Context

[Tree] (config>router>rsvp>interface te-down-threshold)

[Tree] (config>router>rsvp te-down-threshold)

Full Context

configure router rsvp interface te-down-threshold

configure router rsvp te-down-threshold

Description

This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The **te-threshold-update** command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates must be supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels is supported.

Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface is used to trigger IGP updates if non-default threshold levels are configured.

The **no** form of this command resets te-down-threshold to its default value.

Default

no te-down-threshold (equals following values 100 99 98 97 96 95 90 85 80 75 60 45 30 15 0)

Parameters

threshold-level

Specifies the threshold level.

Values 0 to 100

Platforms

7705 SAR Gen 2

30.32 te-metric

te-metric

Syntax

te-metric *value*

no te-metric

Context

[Tree] (config>router>mpls>interface te-metric)

Full Context

configure router mpls interface te-metric

Description

This command configures the TE metric used on the interface. This metric is in addition to the interface metric used by IGP for the shortest path computation.

This metric is flooded as part of the TE parameters for the interface using an opaque LSA or an LSP. The IS-IS TE metric is encoded as sub-TLV 18 as part of the extended IS reachability TLV. The metric value is encoded as a 24-bit unsigned integer. The OSPF TE metric is encoded as a sub-TLV Type 5 in the Link TLV. The metric value is encoded as a 32-bit unsigned integer.

When the use of the TE metric is enabled for an LSP, CSPF will first prune all links in the network topology which do not meet the constraints specified for the LSP path. Such constraints include bandwidth, admin-groups, and hop limit. Then, CSPF will run an SPF on the remaining links. The shortest path among the all SPF paths will be selected based on the TE metric instead of the IGP metric which is used by default.

The TE metric in CSPF LSP path computation can be configured by entering the command **config>router>mpls>lsp>metric-type te**.

Note that the TE metric is only used in CSPF computations for MPLS paths and not in the regular SPF computation for IP reachability.

The **no** form of this command reverts to the default value.

Default

no te-metric

The value of the IGP metric is advertised in the TE metric sub-TLV by IS-IS and OSPF.

Parameters

value

Specifies the metric value.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

30.33 te-threshold-update

te-threshold-update

Syntax

[**no**] **te-threshold-update**

Context

[Tree] (config>router>rsvp te-threshold-update)

Full Context

configure router rsvp te-threshold-update

Description

This command is used to control threshold-based IGP TE updates. The **te-threshold-update** command must enable IGP TE update based only on bandwidth reservation thresholds per interface and must block IGP TE update on bandwidth changes for each reservation. Threshold levels can be defined using the **te-up-threshold** and **te-down-threshold** commands at the global RSVP or per-interface level.

The **no** form of this command should reset te-threshold-update to the default value and disable threshold based update.

Default

no te-threshold-update

Platforms

7705 SAR Gen 2

te-threshold-update

Syntax

te-threshold-update

no te-threshold-update

Context

[Tree] (debug>router>rsvp>event te-threshold-update)

Full Context

debug router rsvp event te-threshold-update

Description

This command debugs the TE threshold update and the dark bandwidth threshold events.

The **no** form of this command disables the debugging.

Platforms

7705 SAR Gen 2

30.34 te-up-threshold

te-up-threshold

Syntax

te-up-threshold *threshold-level* [*threshold-level*]

no te-up-threshold

Context

[Tree] (config>router>rsvp te-up-threshold)

[Tree] (config>router>rsvp>interface te-up-threshold)

Full Context

configure router rsvp te-up-threshold

configure router rsvp interface te-up-threshold

Description

This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The **te-threshold-update** command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates must be supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels must be supported.

Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface is used to trigger IGP updates if non-default threshold levels are configured.

The **no** form of this command resets te-up-threshold to its default value.

Default

no te-up-threshold (equals values of 0 15 30 45 60 75 80 85 90 95 96 97 98 99 100)

Parameters

threshold-level

Specifies the threshold level.

Values 0 to 100

Platforms

7705 SAR Gen 2

30.35 tech-support

tech-support

Syntax

tech-support [*file-url*]

Context

[Tree] (admin tech-support)

Full Context

admin tech-support

Description

This command creates a system core dump. If the *file-url* is omitted, and a *ts-location* is defined, then the **tech support** file will have an automatic SR OS generated file name based on the system name and the date and time and will be saved to the directory indicated by the configured *ts-location*.

The format of the auto-generated filename is ts-XXXXX.YYYYMMDD.HHMMUTC.dat where:

- XXXXX: system name with special characters expanded to avoid problems with file systems (for example, a '.' is expanded to %2E.)
- YYYYMMDD: Date with leading zeros on year, month and day
- HHMM: Hours and Minutes in UTC time (24hr format, always 4 chars, with leading zeros on hours and minutes)

Note:
This command should only be used with authorized direction of Nokia support.

Parameters

file-url

Specifies the file URL location to save the binary file.

Values *local-url* | *remote-url*

local-url	[*cflash-id/*][*file-path*] 200 chars max, including cflash-id
	directory length 99 chars max each
remote-url	[{ftp:// | tftp://}*login:pswd@remote-locn/*][*file-path*]
	199 chars max
remote-locn	[*hostname* | *ipv4-address* | *ipv6-address*]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[*-interface*]
	x:x:x:x:x:x:d.d.d.d[*-interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D
	interface - 32 chars max, for link local addresses
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

30.36 telemetry

telemetry

Syntax

telemetry

Context

[Tree] (config>system telemetry)

[Tree] (admin>system telemetry)

Full Context

configure system telemetry

admin system telemetry

Description

Commands in this context configure the dial-out telemetry commands.

Platforms

7705 SAR Gen 2

30.37 telemetry-data

telemetry-data

Syntax

[**no**] **telemetry-data**

Context

[Tree] (config>system>security>management-interface>output-authorization telemetry-data)

Full Context

configure system security management-interface output-authorization telemetry-data

Description

This command controls output authorization of telemetry configuration and state data in gNMI Subscribe RPC responses.

When enabled, telemetry data output authorization is performed, which may significantly increase the system response time with command authorization requests, especially when remote AAA servers are used.

By default, authorization checks are not performed for telemetry data.

The **no** form of this command reverts to the default value.

Default

no telemetry-data

Platforms

7705 SAR Gen 2

30.38 telnet

telnet

Syntax

telnet {*ip-address* | *dns-name*} [*port*] **service-name** *service-name* [**source** *ip-address*]

telnet {*ip-address* | *dns-name*} [*port*] [**router** *router-instance*] [**source** *ip-address*]

Context

[Tree] (telnet)

Full Context

telnet

Description

This command opens a Telnet session to a remote host. In 7705 SAR Gen 2 networks, the Telnet servers limit Telnet clients to three login attempts; if unsuccessful, the Telnet client session is disconnected. The number is not user configurable.

If a source address is specified, it is used for the source IP address in the originated IP packets for the Telnet session.

Parameters

ip-address

Specifies the IP address or the DNS name (if DNS name resolution is configured).

Values

ipv4-address	*a.b.c.d*	
ipv6-address	*x:x:x:x:x:x:x:x*[-interface]	*x*: [0 to FFFF]H
	x:x:x:x:x:x:d.d.d.d[-interface]	*d*: [0 to 255]D *ipv6-address*
		interface: up to 32 characters, mandatory for link local addresses
dns-name	up to 128 characters	

dns-name

Specifies the DNS name (if DNS name resolution is configured), up to 128 characters.

port

Specifies the TCP port number to use Telnet to the remote host, expressed as a decimal integer.

Values 1 to 65535

Default 23

router-instance

Specifies the router name or service ID used to identify the router instance.

Values

router-instance: *router-name* or *vprn-svc-id*

router-name	"Base", "management", vpls-management"
vprn-svc-id	1 to 2147483647

Default Base

service-name

Specifies the service name, up to 64 characters.

source ip-address

Specifies the source IP address to use as the source of the Telnet packets.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]

Platforms

7705 SAR Gen 2

telnet

Syntax

telnet

Context

[Tree] (config>system>security telnet)

[Tree] (config>system>login-control telnet)

Full Context

configure system security telnet

configure system login-control telnet

Description

Commands in this context configure the Telnet parameters.

Platforms

7705 SAR Gen 2

30.39 telnet-max-sessions

telnet-max-sessions

Syntax

telnet-max-sessions *number-of-sessions*

no telnet-max-sessions

Context

[Tree] (config>system>security>profile telnet-max-sessions)

[Tree] (config>system>security>cli-session-group telnet-max-sessions)

Full Context

configure system security profile telnet-max-sessions

configure system security cli-session-group telnet-max-sessions

Description

This command is used to limit the number of Telnet-based CLI sessions available to all users that are part of a particular profile, or to all users of all profiles that are part of the same cli-session-group.

The **no** form of this command disables the command and the profile/group limit is not applied on the number of sessions.

Default

no telnet-max-sessions

Parameters

number-of-sessions

Specifies the maximum number of allowed Telnet-based CLI sessions.

Values 0 to 50

Platforms

7705 SAR Gen 2

30.40 telnet-reply

telnet-reply

Syntax

[**no**] **telnet-reply**

Context

[Tree] (config>service>ies>if>ipv6>vrrp telnet-reply)

Full Context

configure service ies interface ipv6 vrrp telnet-reply

Description

This command enables the non-owner master to reply to TCP port 23 Telnet requests directed at the virtual router instances IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Proper login and CLI command authentication is still enforced.

When telnet-reply is not enabled, TCP port 23 Telnet packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to Telnet requests regardless of the telnet-reply configuration.

The **telnet-reply** command is only available in non-owner VRRP nodal context. If the telnet-reply command is not executed, Telnet packets to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all Telnet packets destined to the non-owner virtual router instance IP addresses.

Default

no telnet-reply

Platforms

7705 SAR Gen 2

telnet-reply

Syntax

[**no**] **telnet-reply**

Context

[Tree] (config>service>ies>if>vrrp telnet-reply)

Full Context

configure service ies interface vrrp telnet-reply

Description

The telnet-reply command enables the non-owner master to reply to TCP port 23 Telnet Requests directed at the virtual router instances IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Proper login and CLI command authentication is still enforced.

When telnet-reply is not enabled, TCP port 23 Telnet packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to Telnet Requests regardless of the telnet-reply configuration.

The telnet-reply command is only available in non-owner VRRP nodal context. If the telnet-reply command is not executed, Telnet packets to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all Telnet packets destined to the non-owner virtual router instance IP addresses.

Default

no telnet-reply

Platforms

7705 SAR Gen 2

telnet-reply

Syntax

[**no**] **telnet-reply**

Context

[Tree] (config>service>vprn>if>vrrp telnet-reply)

[Tree] (config>service>vprn>if>ipv6>vrrp telnet-reply)

Full Context

configure service vprn interface vrrp telnet-reply

configure service vprn interface ipv6 vrrp telnet-reply

Description

This command enables the non-owner master to reply to TCP port 23 Telnet Requests directed at the virtual router instance's IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Proper login and CLI command authentication is still enforced.

When telnet-reply is not enabled, TCP port 23 Telnet packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to Telnet Requests regardless of the telnet-reply configuration.

The telnet-reply command is only available in non-owner **VRRP** nodal context. If the telnet-reply command is not executed, Telnet packets to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all Telnet packets destined to the non-owner virtual router instance IP addresses.

Default

no telnet-reply

Platforms

7705 SAR Gen 2

telnet-reply

Syntax

[**no**] **telnet-reply**

Context

[Tree] (config>router>if>vrrp telnet-reply)

[Tree] (config>router>if>ipv6>vrrp telnet-reply)

Full Context

configure router interface vrrp telnet-reply

configure router interface ipv6 vrrp telnet-reply

Description

This command enables the non-owner master to reply to TCP port 23 Telnet requests directed at the virtual router instances' IP addresses.

Non-owner virtual router instances are limited by the VRRP specifications to responding to ARP requests destined to the virtual router IP addresses and routing IP packets not addressed to the virtual router IP addresses. Many network administrators find this limitation frustrating when troubleshooting VRRP connectivity issues.

This limitation can be disregarded for certain applications. Ping, SSH and Telnet can each be individually enabled or disabled on a per-virtual-router-instance basis.

The **telnet-reply** command enables the non-owner master to reply to Telnet requests directed at the virtual router instances' IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Correct login and CLI command authentication is still enforced.

When **telnet-reply** is not enabled, Telnet requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to Telnet requests regardless of the **telnet-reply** setting.

The **telnet-reply** command is only available in non-owner **vrrp** nodal context.

By default, Telnet requests to the virtual router instance IP addresses will be silently discarded.

The **no** form of the command configures discarding all Telnet request messages destined to the non-owner virtual router instance IP addresses.

Default

no telnet-reply — Telnet requests to the virtual router instance IP addresses are discarded.

Platforms

7705 SAR Gen 2

30.41 telnet-server

telnet-server

Syntax

[**no**] **telnet-server**

Context

[Tree] (config>system>security telnet-server)

Full Context

configure system security telnet-server

Description

This command enables Telnet servers running on the system.

Telnet servers are shut down by default. At system startup, only SSH servers are enabled.

Telnet servers in networks limit a Telnet clients to three retries to login. The Telnet server disconnects the Telnet client session after three retries.

The **no** form of this command disables Telnet servers running on the system.

Platforms

7705 SAR Gen 2

30.42 telnet6-server

telnet6-server

Syntax

[**no**] **telnet6-server**

Context

[Tree] (config>system>security telnet6-server)

Full Context

configure system security telnet6-server

Description

This command enables Telnet IPv6 servers running on the system.

Telnet servers are shut down by default. At system startup, only SSH servers are enabled.

The **no** form of this command disables Telnet IPv6 servers running on the system.

Platforms

7705 SAR Gen 2

30.43 temp-flooding

temp-flooding

Syntax

temp-flooding flood-time

no temp-flooding

Context

[Tree] (config>service>vpls temp-flooding)

[Tree] (config>service>template>vpls-template temp-flooding)

Full Context

configure service vpls temp-flooding

configure service template vpls-template temp-flooding

Description

The temporary flooding is designed to minimize failover times by eliminating the time it takes to flush the MAC tables and if MVRP is enabled the time it takes for MVRP registration. Temporary flooding is initiated only upon xSTP TCN reception. During this procedure while the MAC flush takes place the frames received on one of the VPLS SAPs/pseudowires are flooded in a VPLS context which for MVRP case includes also the unregistered MVRP trunk ports. The MAC Flush action is initiated by the STP TCN reception or if MVRP is enabled for the data VPLS, by the reception of a MVRP New message for the SVLAN ID associated with the data VPLS. As soon as the MAC Flush is done, regardless of whether the temp-flooding timer expired or not, traffic will be delivered according to the regular FDB content which may be built from MAC Learning or based on MVRP registrations. This command provides a flood-time value that configures a fixed amount of time, in seconds, during which all traffic is flooded (BUM or known unicast) as a safety mechanism. Once the flood-time expires, traffic will be delivered according to the regular FDB content which may be built from MAC Learning or based on MVRP registrations. The temporary flooding timer should be configured in such a way to allow auxiliary processes like MAC Flush, MMRP and/or MVRP to complete/converge. The temporary flooding behavior applies to regular VPLS, VPLS instantiated with VPLS-template, IVPLS and BVPLS when MMRP is disabled.

The **no** form of this command disables the temporary flooding behavior.

Default

no temp-flooding

Parameters

flood-time

Specifies the flood time, in seconds

Values 3 to 600

Platforms

7705 SAR Gen 2

30.44 template

template

Syntax

template

Context

[Tree] (config>service template)

Full Context

configure service template

Description

This is the node for service templates.

Platforms

7705 SAR Gen 2

template

Syntax

[**no**] **template** *name*

Context

[Tree] (config>router>route-next-hop-policy template)

Full Context

configure router route-next-hop-policy template

Description

This command creates a template to configure the attributes of a Loop-Free Alternate (LFA) Shortest Path First (SPF) policy. An LFA SPF policy allows the user to apply specific criteria, such as admin group and SRLG constraints, to the selection of an LFA backup next-hop for a subset of prefixes that resolve to a specific primary next-hop.

The user first creates a route next-hop policy template under the global router context and then applies it to a specific OSPF or IS-IS interface in the global routing instance or in a VPRN instance.

A policy template can be used in both IS-IS and OSPF to apply the specific criteria to prefixes protected by LFA. Each instance of IS-IS or OSPF can apply the same policy template to one or more interface.

The commands within the route next-hop policy template use the **begin-commit-abort** model. The following are the steps to create and modify the template:

To create a template, the user enters the name of the new template directly under the route-next-hop-policy context.

1. To delete a template that is not in use, the user enters the **no** form for the template name under the route-next-hop-policy context.
2. The user enters the editing mode by executing the begin command under the route-next-hop-policy context. The user can then edit and change any number of route next-hop policy templates. However, the parameter value will still be stored temporarily in the template module until the commit is executed under the route-next-hop-policy context. Any temporary parameter changes will be lost if the user enters the abort command before the commit command.
3. The user is allowed to create or delete a template instantly once in the editing mode without the need to enter the commit command. Furthermore, the abort command, if entered, will have no effect on the prior deletion or creation of a template.

Once the commit command is issued, IS-IS or OSPF will re-evaluate the templates and if there are any net changes, it will schedule a new LFA SPF to re-compute the LFA next-hop for the prefixes associated with these templates.

Parameters

name

Specifies the name of the template, up to 32 characters.

Platforms

7705 SAR Gen 2

30.45 terminal

terminal

Syntax

terminal

no terminal

Context

[Tree] (environment terminal)

Full Context

environment terminal

Description

Commands in this context configure the terminal screen length for the current CLI session.

Platforms

7705 SAR Gen 2

30.46 tertiary-config

tertiary-config

Syntax

tertiary-config *file-url*

no tertiary-config

Context

[Tree] (bof tertiary-config)

Full Context

bof tertiary-config

Description

This command specifies the name and location of the tertiary configuration file.

The system attempts to use the configuration specified in **tertiary-config** if both the primary and secondary config files cannot be located. If this file cannot be located, the system boots with the factory default configuration.

Note that if an error in the configuration file is encountered, the boot process aborts.

The **no** form of this command removes the **tertiary-config** configuration.

Parameters

file-url

Specifies the tertiary configuration file location, expressed as a file URL.

Values

file-url	{*local-url* | *remote-url*} (up to 180 characters)
local-url	[*cflash-id*/][*file-path*]
remote-url	[{ftp://| tftp://} *login:pswd@remote-locn*/][*file-path*]
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

30.47 tertiary-dns

tertiary-dns

Syntax

tertiary-dns *ip-address*

no tertiary-dns

Context

[Tree] (config>service>vprn>dns tertiary-dns)

Full Context

configure service vprn dns tertiary-dns

Description

This command configures the tertiary DNS server for DNS name resolution. The tertiary DNS server is used only if the primary DNS server and the secondary DNS server do not respond.

DNS name resolution can be used when executing ping, traceroute, and service-ping, and also when defining file URLs. DNS name resolution is not supported when DNS names are embedded in configuration files.

The **no** form of this command removes the tertiary DNS server from the configuration.

Default

no tertiary-dns — No tertiary DNS server is configured.

Parameters

ip-address

The IP or IPv6 address of the tertiary DNS server.

Values

ipv4-address -a.b.c.d

ipv6-address: x:x:x:x:x:x:x:x[-interface]

x:x:x:x:x:x:d.d.d.d[-interface]

x: [0 to FFFF]H

d: [0 to 255]D

interface - 32 characters max, for link local addresses.

Platforms

7705 SAR Gen 2

tertiary-dns

Syntax

tertiary-dns *ip-address*

no tertiary-dns [*ip-address*]

Context

[Tree] (bof tertiary-dns)

Full Context

bof tertiary-dns

Description

This command configures the tertiary DNS server for DNS name resolution. The tertiary DNS server is used only if the primary DNS server and the secondary DNS server do not respond.

DNS name resolution can be used when executing ping, traceroute, and service-ping, and also when defining file URLs. DNS name resolution is not supported when DNS names are embedded in configuration files.

The **no** form of this command removes the tertiary DNS server from the configuration.

Default

no tertiary-dns

Parameters

ip-address

Specifies the IP or IPv6 address of the tertiary DNS server.

Values

ipv4-address	a.b.c.d
ipv6-address	*x:x:x:x:x:x:x:x*[-interface]
	x:x:x:x:x:x:d.d.d.d[-interface]
	x: [0 to FFFF]H
	d: [0 to 255]D
interface	32 chars max, for link local addresses

Platforms

7705 SAR Gen 2

30.48 tertiary-image

tertiary-image

Syntax

tertiary-image *file-url*

no tertiary-image

Context

[Tree] (bof tertiary-image)

Full Context

bof tertiary-image

Description

This command specifies the tertiary directory location for runtime image file loading.

The system attempts to load all runtime image files configured in the **primary-image** first. If this fails, the system attempts to load the runtime images from the location configured in the **secondary-image**. If the secondary image load fails, the tertiary image specified in **tertiary-image** is used.

All runtime image files (*.tim files) must be located in the same directory.

The **no** form of this command removes the **tertiary-image** configuration.

Parameters

file-url

Specifies the file URL; can be either local (this CPM) or a remote FTP server.

Values

file-url	{*local-url* | *remote-url*} (up to 180 characters)
local-url	[*cflash-id*/][*file-path*]
remote-url	[{ftp://| tftp://} *login:pswd@remote-locn*/][*file-path*]
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

30.49 tertiary-ip-address

tertiary-ip-address

Syntax

tertiary-ip-address *ipv4-address*

no tertiary-ip-address

Context

[Tree] (config>router>bgp>orr>location tertiary-ip-address)

Full Context

configure router bgp optimal-route-reflection location tertiary-ip-address

Description

This command specifies the tertiary IP address of a reference location used for BGP optimal route reflection. Up to three IPv4 addresses and three IPv6 addresses can be specified per location.

If the TE DB is unable to find a node in its topology database that matches the primary address, then the TE DB tries to find a node with the matching secondary address. If this attempt also fails, the TE DB then tries to find a node with the matching tertiary address.

The IP addresses specified for a location should be topologically "close” to a set of clients that should all receive the same optimal path for that location.

The **no** form of this command removes the tertiary IP address information.

Default

no tertiary-ip-address

Parameters

ipv4-address

Specifies the tertiary IPv4 address of a location, expressed in dotted decimal notation.

Values a.b.c.d

Platforms

7705 SAR Gen 2

30.50 tertiary-ipv6-address

tertiary-ipv6-address

Syntax

tertiary-ipv6-address *ipv6-address*

no tertiary-ipv6-address

Context

[Tree] (config>router>bgp>orr>location tertiary-ipv6-address)

Full Context

configure router bgp optimal-route-reflection location tertiary-ipv6-address

Description

This command specifies the tertiary IPv6 address of a reference location used for BGP optimal route reflection. Up to three IPv4 addresses and three IPv6 addresses can be specified per location.

If the TE DB is unable find a node in its topology database that matches a primary address of the location, then it tries to find a node matching a secondary address. If this attempt also fails, the TE DB tries to find a node matching a tertiary address.

The IP addresses specified for a location should be topologically "close” to a set of clients that should all receive the same optimal path for that location.

The **no** form of this command removes the tertiary IPv6 address information.

Default

no tertiary-ipv6-address

Parameters

ipv6-address

Specifies the tertiary IPv6 address of a location.

Values ipv6-address:

- x:x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Platforms

7705 SAR Gen 2

30.51 test

test

Syntax

[**no**] **test** *test-name* [**owner** *test-owner*]

Context

[Tree] (config>saa test)

Full Context

configure saa test

Description

This command identifies a test and enables the context to provide the test parameters for the named test. After the creation of the test instance, the test can be started in the OAM context.

A test can only be modified while it is shut down.

The **no** form of this command removes the test from the configuration. To remove a test, it cannot be active at the time.

Parameters

test-name

Identifies the SAA test name, up to 32 characters.

test-owner

Specifies the owner, up to 32 characters, of an SAA operation. If a value is not specified, the default owner is used.

Default "TiMOS CLI"

Platforms

7705 SAR Gen 2

30.52 test-completion-enable

test-completion-enable

Syntax

[**no**] **test-completion-enable**

Context

[Tree] (config>saa>test>trap-gen test-completion-enable)

Full Context

configure saa test trap-gen test-completion-enable

Description

This command enables the generation of a trap when an SAA test completes.

The **no** form of this command disables the trap generation.

Platforms

7705 SAR Gen 2

30.53 test-duration

test-duration

Syntax

test-duration *seconds*

no test-duration

Context

[Tree] (config>oam-pm>session>ip>twamp-light test-duration)

Full Context

configure oam-pm session ip twamp-light test-duration

Description

This command defines the length of time the test runs before stopping automatically. This optional command is only valid when a session has been configured with a **session-type** of **on-demand**. This is not an option when the **session-type** is configured as **proactive**. On-demand tests do not start until the **config>oam-pm>session>start** command has been issued and they stop when the **config>oam-pm>session>stop** command is issued.

The **no** form of this command removes a previously configured test-duration value and allows the TWAMP Light test to execute until it is stopped manually.

Parameters

seconds

Specifies the length of time, in seconds, that the TWAMP Light test runs.

Values 1 to 86400

Platforms

7705 SAR Gen 2

30.54 test-fail-enable

test-fail-enable

Syntax

[**no**] **test-fail-enable**

Context

[Tree] (config>saa>test>trap-gen test-fail-enable)

Full Context

configure saa test trap-gen test-fail-enable

Description

This command enables the generation of a trap when a test fails. In the case of a ping test, the test is considered failed (for trap generation) if the number of failed probes is at least the value of the **test-fail-threshold** parameter.

The **no** form of this command disables the trap generation.

Platforms

7705 SAR Gen 2

30.55 test-fail-threshold

test-fail-threshold

Syntax

test-fail-threshold *threshold*

no test-fail-threshold

Context

[Tree] (config>saa>test>trap-gen test-fail-threshold)

Full Context

configure saa test trap-gen test-fail-threshold

Description

This command configures the threshold for trap generation on test failure.

This command has no effect when **test-fail-enable** is disabled. This command is not applicable to SAA trace route tests.

The **no** form of this command returns the threshold value to the default.

Default

test-fail-threshold 1

Parameters

threshold

Specifies the number of consecutive test failures required to generate a trap.

Values 0 to 15

Platforms

7705 SAR Gen 2

30.56 test-oam

test-oam

Syntax

test-oam

Context

[Tree] (config test-oam)

Full Context

configure test-oam

Description

Commands in this context configure operations, administration, and maintenance (OAM) test parameters.

Platforms

7705 SAR Gen 2

30.57 third-party-nexthop

third-party-nexthop

Syntax

third-party-nexthop

no third-party-nexthop

Context

[Tree] (config>service>vprn>bgp>group third-party-nexthop)

[Tree] (config>service>vprn>bgp>group>neighbor third-party-nexthop)

[Tree] (config>service>vprn>bgp third-party-nexthop)

Full Context

configure service vprn bgp group third-party-nexthop

configure service vprn bgp group neighbor third-party-nexthop

configure service vprn bgp third-party-nexthop

Description

Use this command to enable the router to send third-party next-hop to EBGP peers in the same subnet as the source peer, as described in RFC 4271. If enabled when an IPv4 or IPv6 route is received from one EBGP peer and advertised to another EBGP peer in the same IP subnet, the BGP next-hop is left unchanged. Third-party next-hop is not done if the address family of the transport does not match the address family of the route.

The **no** form of this command prevents BGP from performing any third party next-hop processing toward any single-hop EBGP peers within the scope of the command. No third-party next-hop means the next-hop will always carry the IP address of the interface used to establish the TCP connection to the peer.

Default

no third-party-nexthop

Platforms

7705 SAR Gen 2

third-party-nexthop

Syntax

third-party-nexthop

no third-party-nexthop

Context

[Tree] (config>router>bgp third-party-nexthop)

[Tree] (config>router>bgp>group>neighbor third-party-nexthop)

Full Context

configure router bgp third-party-nexthop

configure router bgp group neighbor third-party-nexthop

Description

Use this command to enable the router to send third-party next-hop to EBGP peers in the same subnet as the source peer, as described in RFC 4271. If enabled when an IPv4 or IPv6 route is received from one EBGP peer and advertised to another EBGP peer in the same IP subnet, the BGP next-hop is left unchanged. Third-party next-hop is not done if the address family of the transport does not match the address family of the route.

The **no** form of this command prevents BGP from performing any third party next-hop processing toward any single-hop EBGP peers within the scope of the command. No third-party next-hop means the next-hop will always carry the IP address of the interface used to establish the TCP connection to the peer.

Default

no third-party-nexthop

Platforms

7705 SAR Gen 2

30.58 three-way-hello

three-way-hello

Syntax

[**no**] **three-way-hello**

Context

[Tree] (config>service>vprn>pim>if three-way-hello)

Full Context

configure service vprn pim interface three-way-hello

Description

This command configures the compatibility mode for enabling the three way hello.

Platforms

7705 SAR Gen 2

three-way-hello

Syntax

three-way-hello [**compatibility-mode**]

no three-way-hello

Context

[Tree] (config>router>pim>interface three-way-hello)

Full Context

configure router pim interface three-way-hello

Description

This command sets the compatibility mode to enable three-way hello. By default, the value is disabled on all interface which specifies that the standard two-way hello is supported. When enabled, the three-way hello is supported.

The **no** form of this command disables three-way hello.

Default

no three-way-hello

Platforms

7705 SAR Gen 2

30.59 threshold

threshold

Syntax

threshold *threshold*

no threshold

Context

[Tree] (config>router>segment-routing>maintenance-policy threshold)

Full Context

configure router segment-routing maintenance-policy threshold

Description

This command configures the minimum number of S-BFD sessions that must be up in order to consider the SR policy candidate path to which the maintenance template is bound to be up. If it is below this number, then the policy candidate path is marked as BFD degraded by the system. This command is only valid in the **ecmp-protected** mode.

The **no** form of this command reverts to the default.

Default

threshold 1

Parameters

threshold

Specifies the minimum number of S-BFD sessions that must be up.

Values 1 to 32

Platforms

7705 SAR Gen 2

30.60 thresholds

thresholds

Syntax

thresholds

Context

[Tree] (config>service>vprn>dhcp6>server>pool thresholds)

[Tree] (config>router>dhcp6>server>pool thresholds)

[Tree] (config>router>dhcp6>server>pool>prefix thresholds)

[Tree] (config>service>vprn>dhcp6>server>pool>prefix thresholds)

Full Context

configure service vprn dhcp6 local-dhcp-server pool thresholds

configure router dhcp6 local-dhcp-server pool thresholds

configure router dhcp6 local-dhcp-server pool prefix thresholds

configure service vprn dhcp6 local-dhcp-server pool prefix thresholds

Description

Commands in this context configure pool level thresholds.

Default

thresholds

Platforms

7705 SAR Gen 2

thresholds

Syntax

thresholds

Context

[Tree] (config>system thresholds)

Full Context

configure system thresholds

Description

Commands in this context configure monitoring thresholds.

Platforms

7705 SAR Gen 2

30.61 throttle-rate

throttle-rate

Syntax

throttle-rate *events* [**interval** *seconds*]

no throttle-rate

Context

[Tree] (config>log throttle-rate)

Full Context

configure log throttle-rate

Description

This command configures the number of events and interval length to be applied to all event types that have throttling enabled by the **event-control** command and do not have a **specific-throttle-rate** configured.

The **no** form of this command reverts to the default values.

Default

throttle-rate 2000 interval 1

Parameters

events

Specifies the number of log events that can be logged within the specified interval for a specific event. Once the limit has been reached, any additional events of that type will be dropped, for example, the event drop count will be incremented. At the end of the throttle interval if any events have been dropped a trap notification will be sent.

Values 1 to 20000

Default 2000

seconds

Specifies the number of seconds that an event throttling interval lasts.

Values 1 to 1200

Default 1

Platforms

7705 SAR Gen 2

30.62 ti-lfa

ti-lfa

Syntax

ti-lfa [**max-sr-frr-labels** *value*] [**max-srv6-frr-sids** *sids-value*]

no ti-lfa

Context

[Tree] (config>router>isis>lfa ti-lfa)

Full Context

configure router isis loopfree-alternates ti-lfa

Description

This command enables the use of the Topology-Independent LFA (TI-LFA) algorithm in the LFA SPF calculation for this IS-IS instance.

The **no** form of this command disables the use of the TI-LFA algorithm in the LFA SPF calculation for this IS-IS instance.

Default

no ti-lfa

Parameters

value

Specifies the maximum number of labels allowed in the segment list of the TI-LFA repair tunnel. A higher value results in better coverage by TI-LFA at the expense of increased packet encapsulation overhead. The TI-LFA algorithm uses this value to limit the search for the Q-node from the P-node on the post-convergence path.

Values 0 to 3

Default 2

sids-value

Specifies the maximum number of SRv6 SIDs allowed in the segment list of the TI-LFA repair tunnel. A higher value results in better coverage by TI-LFA at the expense of increased packet encapsulation overhead. The TI-LFA algorithm uses this value to limit the search for the Q-node from the P-node on the post-convergence path.

Values 0 to 3

Default 1

Platforms

7705 SAR Gen 2

ti-lfa

Syntax

ti-lfa [**max-sr-frr-labels** *value*]

no ti-lfa

Context

[Tree] (config>router>ospf>loopfree-alternates ti-lfa)

Full Context

configure router ospf loopfree-alternates ti-lfa

Description

This command enables the use of the Topology Independent Loop-Free Alternate (TI-LFA) algorithm in the LFA SPF calculation for this OSPF or OSPFv3 instance.

The **no** form of this command disables the use of the TI-LFA algorithm in the LFA SPF calculation in this OSPF or OSPFv3 instance.

Default

no ti-lfa

Parameters

max-sr-frr-labels [*value*]]

Specifies the maximum number of labels allowed in the segment list of the TI-LFA repair tunnel. A higher value results in better coverage by TI-LFA at the expense of increased packet encapsulation overhead. The TI-LFA algorithm uses this value to limit the search for the Q-node from the P-node on the post-convergence path.

Values 0 to 3

Default 2

Platforms

7705 SAR Gen 2

30.63 tier

tier

Syntax

tier {**1** | **2**}

Context

[Tree] (config>qos>policer-control-policy tier)

Full Context

configure qos policer-control-policy tier

Description

This command is used to create, configure, and delete tiered arbiters. Two tiers are supported that always exist, specified as tier 1 and tier 2. Tiered arbiters enable the creation of a bandwidth control hierarchy for managing child policers in an arbitrary fashion. Each arbiter enables parenting of child policers within eight strict levels of priority and a maximum aggregate rate may be defined for the children that the arbiter will enforce. Arbiters created on tier 1 are automatically parented to the root arbiter that is always present. Arbiters created on tier 2 default to the root arbiter as parent but can also be explicitly parented to a tier 2 arbiter. Child policers associated with an instance of the **policer-control-policy** can be parented to any tiered arbiter or to the root arbiter.

Platforms

7705 SAR Gen 2

tier

Syntax

[**no**] **tier** *tier*

Context

[Tree] (config>qos>scheduler-policy tier)

Full Context

configure qos scheduler-policy tier

Description

This command identifies the level of hierarchy that a group of schedulers are associated with. Within a tier level, a scheduler can be created or edited. Schedulers created within a tier can only be a child (take bandwidth from a scheduler in a higher tier). Tier levels increase sequentially with 1 being the highest tier. All tier 1 schedulers are considered to be root and cannot be a child of another scheduler. Schedulers defined in tiers other than 1 can also be root (parentless).

3 tiers (levels 1, 2, and 3) are supported.

The **save config** and **show config** commands only display information on scheduler tiers that contain defined schedulers. When all schedulers have been removed from a level, that level ceases to be included in output from these commands.

Parameters

tier

This parameter is required to indicate the group of schedulers to create or be edited. Tier levels cannot be created or deleted. If a value for level is given that is out-of-range, an error will occur and the current context of the CLI session will not change.

Values 1 to 3

Platforms

7705 SAR Gen 2

30.64 time

time

Syntax

time

Context

[Tree] (config>system time)

Full Context

configure system time

Description

Commands in this context configure the system time zone and time synchronization parameters.

Platforms

7705 SAR Gen 2

30.65 time-display

time-display

Syntax

time-display {**local** | **utc**}

Context

[Tree] (environment time-display)

Full Context

environment time-display

Description

This command displays time stamps in the CLI session based on local time or Coordinated Universal Time (UTC).

The system keeps time internally in UTC and is capable of displaying the time in either UTC or local time based on the time zone configured.

This environment command only applies to times displayed in the current CLI session. This includes displays of event logs and all other places where a time stamp is displayed.

In event logs, the selected time is used to control the timestamps in the CLI output of **show log log-id** and in YANG state in the /state/log/log-id branch (for logs such as session, cli, memory, SNMP and NETCONF).

Also see the **configure log log-id time-format** command.

Default

time-display local

Parameters

local

Indicates that local time should be used.

utc

Indicates that UTC time should be used.

Platforms

7705 SAR Gen 2

time-display

Syntax

time-display {**local** | **utc**}

Context

[Tree] (config>system>management-interface>cli>md-cli>environment time-display)

Full Context

configure system management-interface cli md-cli environment time-display

Description

This command configures whether the time is displayed in coordinated Universal Time (UTC) or local time (as configured in **config**>**system**>**time**).

Default

time-display local

Parameters

local

Specifies that the local time zone is used.

utc

Specifies that UTC is used.

Platforms

7705 SAR Gen 2

30.66 time-exceeded

time-exceeded

Syntax

time-exceeded [*number seconds*]

no time-exceeded

Context

[Tree] (config>service>ies>if>ipv6>icmp6 time-exceeded)

Full Context

configure service ies interface ipv6 icmp6 time-exceeded

Description

This command specifies whether time-exceeded ICMP messages should be sent. When enabled, ICMPv6 time-exceeded messages are generated by this interface.

When disabled, ICMPv6 time-exceeded messages are not sent.

The **no** form of this command reverts to the default.

Default

time-exceeded 100 10

Parameters

number

Specifies the number of time-exceeded ICMP messages are to be issued in the time frame specified by the *seconds* parameter.

Values 10 to 2000

seconds

Specifies the time frame, in seconds, that is used to limit the number of time-exceeded ICMP message to be issued.

Values 1 to 60

Platforms

7705 SAR Gen 2

time-exceeded

Syntax

time-exceeded [*number seconds*]

no time-exceeded

Context

[Tree] (config>router>if>ipv6>icmp6 time-exceeded)

[Tree] (config>service>vprn>if>ipv6>icmp6 time-exceeded)

Full Context

configure router interface ipv6 icmp6 time-exceeded

configure service vprn interface ipv6 icmp6 time-exceeded

Description

This command configures rate for ICMPv6 time-exceeded messages.

Parameters

number

Limits the number of time-exceeded messages issued per the time frame specified in *seconds* parameter.

Values 10 to 2000

seconds

Determines the time frame, in seconds, that is used to limit the number of time-exceeded messages issued per time frame.

Values 1 to 60

Platforms

7705 SAR Gen 2

30.67 time-format

time-format

Syntax

time-format {**local** | **utc**}

Context

[Tree] (config>service>vprn>log>log-id time-format)

Full Context

configure service vprn log log-id time-format

Description

This command specifies whether the time should be displayed in local or Coordinated Universal Time (UTC) format.

Default

time-format utc

Parameters

local

Specifies that timestamps are written in the system's local time.

utc

Specifies that timestamps are written using the UTC value. This was formerly called Greenwich Mean Time (GMT) and Zulu time.

Platforms

7705 SAR Gen 2

time-format

Syntax

time-format {**local** | **utc**}

Context

[Tree] (config>log>log-id time-format)

Full Context

configure log log-id time-format

Description

This command specifies whether the time should be output in local or Coordinated Universal Time (UTC) format in the following event log locations:

- in the syslog TIMESTAMP field
- in the timestamp of log events inside log files on local storage devices

The timestamp in the filename of event log files is not affected by this command.

The output of **show log log-id** and the output of YANG state under /state/log/log-id are not affected by this command. See the **environment time-display** command.

Default

time-format utc

Parameters

local

Specifies that timestamps are written in the system's local time.

utc

Specifies that timestamps are written using the UTC value. This was formerly called Greenwich Mean Time (GMT) and Zulu time.

Platforms

7705 SAR Gen 2

30.68 time-stamp

time-stamp

Syntax

[**no**] **time-stamp**

Context

[Tree] (environment time-stamp)

Full Context

environment time-stamp

Description

This command specifies whether the time-stamp should be displayed before the prompt.

Platforms

7705 SAR Gen 2

30.69 timeout

timeout

Syntax

timeout [**sec** *seconds*] [**min** *minutes*]

no timeout

Context

[Tree] (config>aaa>radius-srv-plcy>servers timeout)

Full Context

configure aaa radius-server-policy servers timeout

Description

This command configures the time the router waits for a response from a RADIUS server.

The no form of this command reverts to the default value.

Default

timeout sec 5

Parameters

seconds

Specifies the number of seconds for the timeout.

Values 1 to 59

minutes

Specifies the number of minutes for the timeout.

Values 1 to 5

Values Max. value = 5 min 40 sec

Platforms

7705 SAR Gen 2

timeout

Syntax

timeout *seconds*

no timeout

Context

[Tree] (config>service>vprn>aaa>rmt-srv>radius timeout)

[Tree] (config>system>security>radius timeout)

Full Context

configure service vprn aaa remote-servers radius timeout

configure system security radius timeout

Description

This command configures the number of seconds the router waits for a response from a RADIUS server.

The **no** form of this command reverts to the default value.

Default

timeout 3

Parameters

seconds

Specifies the number of seconds the router waits for a response from a RADIUS server, expressed as a decimal integer.

Values 1 to 90

Platforms

7705 SAR Gen 2

timeout

Syntax

timeout *seconds*

no timeout

Context

[Tree] (config>service>vprn>aaa>rmt-srv>tacplus timeout)

[Tree] (config>system>security>tacplus timeout)

Full Context

configure service vprn aaa remote-servers tacplus timeout

configure system security tacplus timeout

Description

This command configures the number of seconds the router waits for a response from a TACACS+ server.

The **no** form of this command reverts to the default value.

Default

timeout 3

Parameters

seconds

Specifies the number of seconds the router waits for a response from a TACACS+ server, expressed as a decimal integer.

Values 1 to 90

Platforms

7705 SAR Gen 2

timeout

Syntax

timeout *seconds*

Context

[Tree] (config>system>file-trans-prof timeout)

Full Context

configure system file-transmission-profile timeout

Description

This command specifies timeout value in seconds for transport protocol. The timeout is the maximum waiting time to receive any data from the server (e.g., FTP or HTTP server).

Default

timeout 60

Parameters

seconds

Specifies the connection timeout (in seconds) for the file transmission.

Values 1 to 3600

Platforms

7705 SAR Gen 2

timeout

Syntax

timeout *timeout*

no timeout

Context

[Tree] (config>saa>test>type-multi-line>lsp-ping>sr-policy timeout)

[Tree] (config>saa>test>type-multi-line>lsp-ping timeout)

Full Context

configure saa test type-multi-line lsp-ping sr-policy timeout

configure saa test type-multi-line lsp-ping timeout

Description

This command configures the number, in seconds, used to override the default *timeout* value and is the amount of time that the router waits for a message reply after sending the last probe for a specific test. Upon the expiration of the time out, the test is marked complete and no more packets are processed for any of the request probes.

The **no** form of this command reverts to the default value.

Default

timeout 5

Parameters

timeout

Specifies the timeout value in seconds.

Values 1 to 10

Default 5

Platforms

7705 SAR Gen 2

timeout

Syntax

timeout [*seconds*]

no timeout

Context

[Tree] (config>filter>redirect-policy>dest>ping-test timeout)

Full Context

configure filter redirect-policy destination ping-test timeout

Description

Specifies the amount of time, in seconds, that is allowed for receiving a response from the far-end host. If a reply is not received within this time the far-end host is considered unresponsive.

Default

timeout 1

Parameters

seconds

Specifies the amount of time, in seconds, that is allowed for receiving a response from the far end host.

Values 1 to 60

Platforms

7705 SAR Gen 2

timeout

Syntax

timeout *seconds*

no timeout

Context

[Tree] (config>vrrp>policy>priority-event>host-unreachable timeout)

Full Context

configure vrrp policy priority-event host-unreachable timeout

Description

This command defines the time, in seconds, that must pass before considering the far-end IP host unresponsive to an outstanding ICMP echo request message.

The **timeout** value is not directly related to the configured **interval** parameter. The **timeout** value may be larger, equal, or smaller, relative to the **interval** value.

If the **timeout** value is larger than the **interval** value, multiple ICMP echo request messages may be outstanding. Every ICMP echo request message transmitted to the far end host is tracked individually according to the message identifier and sequence number.

With each consecutive attempt to send an ICMP echo request message, the timeout timer is loaded with the **timeout** value. The timer decrements until:

- an internal error occurs preventing message sending (request unsuccessful)
- an internal error occurs preventing message reply receiving (request unsuccessful)
- a required route table entry does not exist to reach the IP address (request unsuccessful)
- a required ARP entry does not exist and ARP request timed out (request unsuccessful)
- a valid reply is received (request successful)

It is possible for a required ARP request to succeed or timeout after the message timeout timer expires. In this case, the message request is unsuccessful.

If an ICMP echo reply message is not received prior to the **timeout** period for a given ICMP echo request, that request is considered to be dropped and increments the consecutive message drop counter for the priority event.

If an ICMP echo reply message with the same sequence number as an outstanding ICMP echo request message is received prior to that message timing out, the request is considered successful. The consecutive message drop counter is cleared and the request message no longer is outstanding.

If an ICMP Echo Reply message with a sequence number equal to an ICMP echo request sequence number that had previously timed out is received, that reply is silently discarded while incrementing the priority event reply discard counter.

The **no** form of the command reverts to the default value.

Default

timeout 1

Parameters

seconds

The number of seconds before an ICMP echo request message is timed out. Once a message is timed out, a reply with the same identifier and sequence number is discarded.

Values 1 to 60

Platforms

7705 SAR Gen 2

timeout

Syntax

timeout *timeout*

no timeout

Context

[Tree] (config>service>sdp>keep-alive timeout)

Full Context

configure service sdp keep-alive timeout

Description

This command configures the time interval that the SDP waits before tearing down the session.

Default

timeout 5

Parameters

timeout

Specifies the timeout time, in seconds.

Values 1 to 10

Platforms

7705 SAR Gen 2

timeout

Syntax

timeout *seconds*

no timeout

Context

[Tree] (config>system>security>ldap timeout)

Full Context

configure system security ldap timeout

Description

The **timeout** value is the number of seconds that the SR OS will wait for a response from the current server that it is trying to establish a connection with. If the server does not reply within the configured **timeout** value, the SR OS will increment the retry counter by 1. The SR OS attempts to establish the connection to the current server up to the configured **retry** value before it moves to the next configured server.

The **no** form of this command reverts to the default value.

Default

timeout 3

Parameters

seconds

The length of time that the SR OS waits for a response from the server.

Values 1 to 90

Default 3

Platforms

7705 SAR Gen 2

30.70 timeouts

timeouts

Syntax

[**no**] **timeouts**

Context

[Tree] (config>service>nat>nat-policy timeouts)

Full Context

configure service nat nat-policy timeouts

Description

This command configures session idle timeouts for this policy.

Platforms

7705 SAR Gen 2

30.71 timers

timers

Syntax

[**no**] **timers**

Context

[Tree] (config>service>vprn>isis timers)

Full Context

configure service vprn isis timers

Description

Commands in this context configure the IS-IS timer values.

Default

n/a

Platforms

7705 SAR Gen 2

timers

Syntax

timers

Context

[Tree] (config>service>vprn>ospf timers)

[Tree] (config>service>vprn>ospf3 timers)

Full Context

configure service vprn ospf timers

configure service vprn ospf3 timers

Description

Commands in this context configure OSPF timers. Timers control the delay between receipt of a LSA requiring a Dijkstra (Shortest Path First (SPF)) calculation and the minimum time between successive SPF calculations.

Changing the timers affect CPU utilization and network reconvergence times. Lower values reduce convergence time but increase CPU utilization. Higher values reduce CPU utilization but increase reconvergence time.

Platforms

7705 SAR Gen 2

timers

Syntax

timers *update timeout flush*

no timers

Context

[Tree] (config>service>vprn>ripng>group>neighbor timers)

[Tree] (config>service>vprn>rip>group timers)

[Tree] (config>service>vprn>ripng timers)

[Tree] (config>service>vprn>rip>group>neighbor timers)

[Tree] (config>service>vprn>ripng>group timers)

[Tree] (config>service>vprn>rip timers)

Full Context

configure service vprn ripng group neighbor timers

configure service vprn rip group timers

configure service vprn ripng timers

configure service vprn rip group neighbor timers

configure service vprn ripng group timers

configure service vprn rip timers

Description

This command configures the values for the update, timeout, and flush timers:

- **update timer**

  Determines how often RIP updates are sent.

- **timeout timer**

  If a router is not updated by the time the timer expires, the route is declared invalid, but maintained in the RIP database.

- **flush timer**

  Determines how long a route is maintained in the RIP database, after it has been declared invalid. Once this timer expires it is flushed from the RIP database completely.

The **no** form of this command resets all timers to their default values of 30, 180, and 120 seconds respectively.

Default

no timers

Parameters

update

The RIP update timer value in seconds.

Values 1 to 600

Default 30

timeout

The RIP timeout timer value in seconds.

Values 1 to 1200

Default 180

flush

The RIP flush timer value in seconds.

Values 1 to 1200

Default 120

Platforms

7705 SAR Gen 2

timers

Syntax

timers [**neighbor** *ip-address* | **group** *name*]

no timers

Context

[Tree] (debug>router>bgp timers)

Full Context

debug router bgp timers

Description

This command logs all BGP timer events to the debug log.

The **no** form of this command disables debugging.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x [-interface] (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H
- d: [0 to 255]D
- interface: up to 32 characters for link local addresses

group *name*

Debugs only events affecting the specified peer group name, up to 64 characters, and associated neighbors.

Platforms

7705 SAR Gen 2

timers

Syntax

[**no**] **timers**

Context

[Tree] (config>router>isis timers)

Full Context

configure router isis timers

Description

This command configures the IS-IS timer values.

Platforms

7705 SAR Gen 2

timers

Syntax

timers

Context

[Tree] (config>router>ospf timers)

[Tree] (config>router>ospf3 timers)

Full Context

configure router ospf timers

configure router ospf3 timers

Description

Commands in this context configure OSPF timers. Timers control the delay between receipt of a link state advertisement (LSA) requiring a Dijkstra (Shortest Path First (SPF)) calculation and the minimum time between successive SPF calculations.

Changing the timers affects CPU utilization and network re-convergence times. Lower values reduce convergence time but increase CPU utilization. Higher values reduce CPU utilization but increase re-convergence time.

Platforms

7705 SAR Gen 2

timers

Syntax

timers *update timeout flush*

no timers

Context

[Tree] (config>router>rip>group>neighbor timers)

[Tree] (config>router>rip timers)

[Tree] (config>router>ripng>group timers)

[Tree] (config>router>ripng>group>neighbor timers)

[Tree] (config>router>ripng timers)

[Tree] (config>router>rip>group timers)

Full Context

configure router rip group neighbor timers

configure router rip timers

configure router ripng group timers

configure router ripng group neighbor timers

configure router ripng timers

configure router rip group timers

Description

This command configures values for the update, timeout and flush RIP timers.

The RIP update timer determines how often RIP updates are sent.

If the route is not updated by the time the RIP timeout timer expires, the route is declared invalid but is maintained in the RIP database.

The RIP flush timer determines how long a route is maintained in the RIP database after it has been declared invalid. After the flush timer expires, the route is removed from the RIP database.

The **no** form of the command reverts to the default values.

Default

timers 30 180 120

Parameters

update

Specifies the RIP update timer value in seconds expressed as a decimal integer.

Values 1 to 600

timeout

Specifies the RIP timeout timer value in seconds expressed as a decimal integer.

Values 1 to 1200

flush

Specifies the RIP flush timer value in seconds expressed as a decimal integer.

Values 1 to 1200

Platforms

7705 SAR Gen 2

30.72 timestamp

timestamp

Syntax

[**no**] **timestamp**

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>prompt timestamp)

Full Context

configure system management-interface cli md-cli environment prompt timestamp

Description

This command displays the timestamp before the first prompt line.

The **no** form of this command suppresses the timestamp before the first prompt line.

Default

timestamp

Platforms

7705 SAR Gen 2

30.73 timestamp-format

timestamp-format

Syntax

timestamp-format millisecond

no timestamp-format

Context

[Tree] (config>log>syslog timestamp-format)

Full Context

configure log syslog timestamp-format

Description

This command controls the format of the syslog timestamp.

The **no** form of this command reverts to the default.

Default

no timestamp-format

Parameters

millisecond

Keyword to set the timestamp format to milliseconds.

Platforms

7705 SAR Gen 2

30.74 timing

timing

Syntax

timing frames-per-delta-t *frames* **consec-delta-t** *deltas* **chli-threshold** *threshold*

no timing

Context

[Tree] (config>oam-pm>session>ip>twamp-light>loss timing)

Full Context

configure oam-pm session ip twamp-light loss timing

Description

This command defines various availability parameters but not the probe interval. A single TWAMP-Light frame is used to collect both delay and loss metrics; the interval is common to both and as such not unique per metric type. Any TWAMP light test that is attempting to become active validates the configuration of the timing parameter regardless of which statistics are being recorded.

The **no** form of this command restores the default values for all timing parameters and use those values to compute availability and set the loss frequency.

Default

timing frames-per-delta-t 1 consec-delta-t 10 chli-threshold 5

Parameters

frames

Defines the size of the small measurement window. Each delta-t is marked as available of unavailable based on the flr-threshold. The size of the delta-t measurement is the product of the number of frames and the interval. This value defaults to a different value thank single probe per metric approaches.

Values 1 to 50

Default 1

deltas

Specifies the number of consecutive delta-t small measurement intervals that make up the sliding window over which availability and unavailability are determined. Transitions from one state to another occurs when the consec-delta-t are now in a new state. The sliding window cannot exceed 100 seconds.

Values 2 to 10

Default 10

threshold

Specifies the number of consecutive high loss intervals (unavailable delta-t) that when equal to or exceeded increments the CHLI counter. A CHLI counter is an indication that the sliding window is available but has crossed a threshold consecutive of unavailable delta-t intervals. A CHLI can only be incremented once during a sliding window and, by default, is only incremented during times of availability.

Values 1 to 9

Default 5

Platforms

7705 SAR Gen 2

30.75 tls

tls

Syntax

tls

Context

[Tree] (config>system>security tls)

Full Context

configure system security tls

Description

This command configures TLS parameters.

Platforms

7705 SAR Gen 2

30.76 tls-client-profile

tls-client-profile

Syntax

tls-client-profile *name*

no tls-client-profile

Context

[Tree] (config>system>telemetry>destination-group tls-client-profile)

[Tree] (config>system>grpc-tunnel>destination-group tls-client-profile)

Full Context

configure system telemetry destination-group tls-client-profile

configure system grpc-tunnel destination-group tls-client-profile

Description

This command configures a TLS client profile to a destination group.

This command is mutually exclusive with the **allow-unsecured-connection** command.

The **no** form of this command removes the TLS client profile.

Default

no tls-client-profile

Parameters

name

Specifies the TLS client profile name, up to 32 characters.

Platforms

7705 SAR Gen 2

tls-client-profile

Syntax

tls-client-profile *tls-client-profile*

no tls-client-profile

Context

[Tree] (config>service>vprn>log>syslog tls-client-profile)

[Tree] (config>log>syslog tls-client-profile)

Full Context

configure service vprn log syslog tls-client-profile

configure log syslog tls-client-profile

Description

This command specifies the Transport Layer Security (TLS) client profile used to encrypt syslog communications. When configured, syslog messages are sent using TLS.

Any change to this command results in a brief interruption of the event log, which may cause the loss of a few syslog messages.

The **no** form of this command removes TLS encryption of syslog communications and sends syslog messages over UDP.

Parameters

tls-client-profile

Specifies the name of a TLS profile configured in the **config**>**system**>**security**>**tls** context, up to 32 characters.

Platforms

7705 SAR Gen 2

tls-client-profile

Syntax

tls-client-profile *profile-name*

no tls-client-profile

Context

[Tree] (config>router>pcep>pcc>peer tls-client-profile)

Full Context

configure router pcep pcc peer tls-client-profile

Description

This command configures a TLS client profile on the PCC. When the TLS profile is configured, the PCC tries to establish a PCEP connection with the PCE over TLS. Because SR OS supports a strict TLS-only mode, both the PCE and PCC must support TLS. If a TLS failure occurs, the connection over TLS is closed and a new connection is retried within 60 seconds.

The **no** form of this command removes TLS encryption from the communication between this PCC and the PCE.

Default

no tls-client-profile

Parameters

profile-name

Specifies the TLS client profile name, up to 32 characters.

Platforms

7705 SAR Gen 2

30.77 tls-profile

tls-profile

Syntax

tls-profile *tls-profile-name*

no tls-profile

Context

[Tree] (config>system>security>ldap>server tls-profile)

Full Context

configure system security ldap server tls-profile

Description

This command attaches a TLS client profile to the LDAP client. The parameter in the TLS profile is used to encrypt the LDAP connection to the server. Each LDAP server can use its own TLS profile.

When a TLS profile is assigned, the LDAP application will send encrypted PDUs from the client to the LDAP server. If TLS is operationally down, the LDAP application should not send any PDUs.

The **no** form of this command removes the TLS profile from LDAP and disables the TLS encryption from LDAP.

Parameters

tls-profile-name

Specifies the TLD profile for encryption.

Platforms

7705 SAR Gen 2

30.78 tls-re-negotiate-timer

tls-re-negotiate-timer

Syntax

tls-re-negotiate-timer *timer-min*

no tls-re-negotiate-timer

Context

[Tree] (config>system>security>tls>server-tls-profile tls-re-negotiate-timer)

Full Context

configure system security tls server-tls-profile tls-re-negotiate-timer

Description

This command configures the timed interval after which the server is triggered to send a Hello request message to all clients and force a renegotiation of the symmetric encryption key. When an interval of 0 is configured, the server will never send a hello request message.

Default

tls-re-negotiate-timer 0

Parameters

timer-min

Specifies the interval, in minutes, after which the server is triggered to send a Hello request message.

Values 0 to 65000

Platforms

7705 SAR Gen 2

30.79 tls-server-profile

tls-server-profile

Syntax

tls-server-profile *name*

no tls-server-profile

Context

[Tree] (config>system>grpc tls-server-profile)

Full Context

configure system grpc tls-server-profile

Description

This command adds a configured TLS server profile to the gRPC session. The TLS server is used for encryption of the gRPC session. gRPC will not transmit any PDUs if there is a TLS server profile assigned to it and the TLS connection is down.

The **no** form of this command removes the specified TLS server profile from the gRPC session.

Parameters

name

Specifies the name of the TLS server profile configured under the **config**>**system**>**security**>**tls** context.

Platforms

7705 SAR Gen 2

30.80 tls-wait-timer

tls-wait-timer

Syntax

tls-wait-timer *tls-wait-timer*

no tls-wait-timer

Context

[Tree] (config>router>pcep>pcc>peer tls-wait-timer)

Full Context

configure router pcep pcc peer tls-wait-timer

Description

This command configures the time that the PCC waits before declaring a TLS handshake failure if the handshake is not established.

The **no** form of this command reverts to the default.

Default

tls-wait-timer 60

Parameters

tls-wait-timer

Specifies the time, in seconds.

Values 60 to 255

Platforms

7705 SAR Gen 2

30.81 tls13-cipher

tls13-cipher

Syntax

tls13-cipher *index* **name** *cipher-suite-code*

no tls13-cipher *index*

Context

[Tree] (config>system>security>tls>server-cipher-list tls13-cipher)

[Tree] (config>system>security>tls>client-cipher-list tls13-cipher)

Full Context

configure system security tls server-cipher-list tls13-cipher

configure system security tls client-cipher-list tls13-cipher

Description

This command configures the TLS 1.3-supported ciphers that are used by the client and server.

The **no** form of this command removes the cipher suite.

Parameters

index

Specifies the index number, which provides the location of the cipher in the negotiation list. The lower index numbers are higher in the negotiation list, and the higher index numbers are at the bottom of the list.

Values 1 to 255

cipher-suite-code

Specifies the cipher suite code.

Values tls-aes128-gcm-sha256
tls-aes256-gcm-sha384
tls-chacha20-poly1305-sha256
tls-aes128-ccm-sha256
tls-aes128-ccm8-sha256

Platforms

7705 SAR Gen 2

30.82 tls13-group

tls13-group

Syntax

tls13-group *index* **name** *group-suite-code*

no tls13-group *index*

Context

[Tree] (config>system>security>tls>server-group-list tls13-group)

[Tree] (config>system>security>tls>client-group-list tls13-group)

Full Context

configure system security tls server-group-list tls13-group

configure system security tls client-group-list tls13-group

Description

This command configures the TLS 1.3-supported group suite codes sent by the client or server in their respective Hello messages.

SR OS supports the use of Elliptic-curve Diffie-Hellman Ephemeral (ECDHE) groups.

The **no** form of this command removes the group suite code.

Parameters

index

Specifies the index number , which provides the location of the group suite code in the client or server group list. The lower index numbers are higher in the list and the higher index numbers are at the bottom of the list.

Values 1 to 255

group-suite-code

Specifies the group suite code.

Values tls-ecdhe-256
tls-ecdhe-384
tls-ecdhe-521
tls-x25519
tls-x448

Platforms

7705 SAR Gen 2

30.83 tls13-signature

tls13-signature

Syntax

tls13-signature *index* **name** *signature-suite-code*

no tls13-signature *index*

Context

[Tree] (config>system>security>tls>server-signature-list tls13-signature)

[Tree] (config>system>security>tls>client-signature-list tls13-signature)

Full Context

configure system security tls server-signature-list tls13-signature

configure system security tls client-signature-list tls13-signature

Description

This command configures the TLS 1.3-supported signature suite codes sent by the client or server in their respective Hello messages.

The **no** form of this command removes the signature suite code.

Parameters

index

Specifies the index number, which provides the location of the signature suite code in the client or server group list. The lower index numbers are higher in the list, and the higher index numbers are at the bottom of the list.

Values 1 to 255

signature-suite-code

Specifies the signature suite code.

Values
tls-rsa-pkcs1-sha256
tls-rsa-pkcs1-sha384
tls-rsa-pkcs1-sha512
tls-ecdsa-secp256r1-sha256
tls-ecdsa-secp384r1-sha384
tls-ecdsa-secp521r1-sha512
tls-rsa-pss-rsae-sha256
tls-rsa-pss-rsae-sha384
tls-rsa-pss-rsae-sha512
tls-rsa-pss-pss-sha256
tls-rsa-pss-pss-sha384
tls-rsa-pss-pss-sha512
tls-ed25519
tls-ed448

Platforms

7705 SAR Gen 2

30.84 to

to

Syntax

to [*ip-address* | **node-id** [*a.b.c.d* | *1...4294967295*]]

Context

[Tree] (config>router>mpls>lsp to)

Full Context

configure router mpls lsp to

Description

This command specifies the IP address or MPLS-TP node-id of the egress router for the LSP. This command is mandatory to create an LSP.

An IP address for which a route does not exist is allowed in the configuration. If the LSP signaling fails because the destination is not reachable, an error is logged and the LSP operational status is set to down.

For a non MPLS-TP LSP, the **to** *ip-address* can be an IP address of a network IP interface, the system interface, or a loopback interface of the egress router. When used in a SDP, if the LSP **to** address does not match the SDP address, the LSP is not included in the SDP definition.

For an MPLS-TP LSP, the **to node-id** may be either in 4-octet IPv4 address format, or a 32-bit unsigned integer. This command is mandatory to create an MPLS-TP LSP. A value of zero is invalid. This **to** address is used in the MPLS-TP LSP ID, and the MPLS-TP MEP ID for the LSP.

Default

no default

Parameters

ip-address

Specifies the IP address of the egress router. When the LSP type is **sr-te**, then an IPv6 address can be used.

Values ipv4-address — a.b.c.d

ipv6-address — x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d

x — 0 to FFFF (hexadecimal)

d — 0 to 255 (decimal)

node-id ***a.b.c.d.*** **|** ***1...4294967295***

4-octet IPv4 formatted or unsigned 32-bit integer MPLS-TP node-id of the egress router.

Platforms

7705 SAR Gen 2

to

Syntax

to *ip-address*

Context

[Tree] (config>router>mpls>static-lsp to)

Full Context

configure router mpls static-lsp to

Description

This command specifies the IP address of the egress router for the static LSP. When creating an LSP this command is required. The **to** IP address may be the address of a local interface, the system IP interface, or of a loopback interface of the egress router. When used in a SDP and the **to** address does not match the far-end SDP address, the LSP is not included in the SDP definition.

Parameters

ip-address

Specifies the system IP address of the egress router.

Platforms

7705 SAR Gen 2

to

Syntax

to file *file-id*

Context

[Tree] (config>log>accounting-policy to)

Full Context

configure log accounting-policy to

Description

This command specifies the destination for the accounting records selected for the accounting policy.

Parameters

file-id

Specifies the destination for the accounting records selected for this destination. The characteristics of the file ID must have already been defined in the **config>log>file** context. A file ID can only be used once.

The file is generated when the file policy is referenced. This command identifies the type of accounting file to be created. The file definition defines its characteristics.

If the **to** command is executed while the accounting policy is in operation, then it becomes active during the next collection interval.

Values 1 to 99

Platforms

7705 SAR Gen 2

to

Syntax

[**no**] **to**

Context

[Tree] (config>router>policy-options>policy-statement>entry to)

Full Context

configure router policy-options policy-statement entry to

Description

This command creates the context to configure export policy match criteria based on a route's destination or the protocol into which the route is being advertised.

If no condition is specified, all route destinations are considered to match.

The **to** command context only applies to export policies. If it is used for an import policy, match criteria is ignored.

The **no** form of this command deletes export match criteria for the route policy statement entry.

Platforms

7705 SAR Gen 2

to

Syntax

to cli [*size*]

to console

to file *log-file-id*

to memory [*size*]

to netconf [*size*]

to session

to snmp [*size*]

to syslog *syslog-id*

Context

[Tree] (config>log>log-id to)

Full Context

configure log log-id to

Description

This command specifies a destination for the log event data.

The source of the data stream must be specified in the **from** command before configuring the destination with the **to** command.

The **to** command cannot be modified or re-entered. If the destination or maximum size of an SNMP or memory log needs to be modified, the log ID must be removed and then recreated.

Parameters

cli

Specifies that log events are directed to any subscribed CLI session. Subscribe to a CLI log from within a CLI session using the **tools**>**perform**>**log**>**subscribe-to log-id** *log-id* command. Events are sent to the CLI session for the duration of that CLI session, or until an **unsubscribe-from** command is issued. A local circular memory log is maintained for CLI logs.

console

Specifies that log events events are directed to the console port. If the console is not connected, all the entries are dropped.

file *log-file-id*

Specifies that log events are directed to a file with the specified *log-file-id*. The characteristics of the *log-file-id* referenced in this parameter must have already been defined in the **config**>**log**>**file** *file-id* context. When the *file-id* location parameter is modified, log files are not written to the new location until a rollover occurs or the log is manually cleared. A rollover can be forced by using the **clear**>**log** command. Subsequent log entries are then written to the new location. If a rollover does not occur or the log is not cleared, the old location continues to be used.

Values 1 to 99, *name* (up to 64 characters max)

memory

Specifies that log events are directed to a memory file. A memory file is a circular buffer; when the file is full, each new entry replaces the oldest entry in the log. If the optional size parameter is not configured, the default value is used.

Default 100

netconf

Specifies that log events are directed to a NETCONF session as notifications. A NETCONF client can subscribe to a NETCONF log using the configured **netconf-stream** *stream-name* for the log in a subscription request. One or more NETCONF sessions can subscribe to a NETCONF log or stream.

session

Specifies that log events are directed to the current console or telnet session. This command is only valid for the duration of the session. When the session is terminated, the

to session configuration is removed. A log ID with a **session** destination is saved in the configuration file but the **to session** part is not stored.

size

Specifies the maximum size of the log data destination, in bytes.

Values 50 to 3000

snmp

Specifies that log events are directed to the **snmp-trap-group** associated with the log ID. A local circular memory log is maintained for SNMP logs.

syslog *syslog-id*

Specifies that log events are directed to the specified syslog collector. To remain consistent with the standards governing syslog, messages to syslog are truncated to 1024 bytes. The characteristics of the *syslog-id* referenced in this parameter must have already been defined in the **config**>**log**>**syslog** *syslog-id* context.

Values 1 to 10

Platforms

7705 SAR Gen 2

to

Syntax

to *ipv4-address*

no to

Context

[Tree] (config>oam-pm>session>ip>tunnel>mpls>rsvp-te-auto to)

Full Context

configure oam-pm session ip tunnel mpls rsvp-te-auto to

Description

This command configures the termination point of the RSV LSP. Configure the following three commands to identify an RSVP-TE Auto LSP: **from**, **to**, and **lsp-template**.When all three of these values are configured, the specific RSVP LSP can be identified and the test packets can be carried across the tunnel

The **no** form of this command removes the IPv4 address.

Parameters

ipv4-address

Specifies IPv4 address.

Values ipv4-address: a.b.c.d (host bits must be 0)

Platforms

7705 SAR Gen 2

30.85 tolerance

tolerance

Syntax

tolerance [*seconds* | **forever**]

no tolerance

Context

[Tree] (config>system>security>keychain>direction>bi>entry tolerance)

[Tree] (config>system>security>keychain>direction>uni>receive>entry tolerance)

Full Context

configure system security keychain direction bi entry tolerance

configure system security keychain direction uni receive entry tolerance

Description

This command configures the amount of time that an eligible receive key should overlap with the active send key or to never expire.

Parameters

seconds

Specifies the duration that an eligible receive key overlaps with the active send key.

Values 0 to 4294967294 seconds

forever

Specifies that an eligible receive key overlap with the active send key forever.

Platforms

7705 SAR Gen 2

30.86 tos-marking-state

tos-marking-state

Syntax

tos-marking-state {**trusted** | **untrusted**}

no tos-marking-state

Context

[Tree] (config>service>vprn>interface tos-marking-state)

[Tree] (config>service>ies>if tos-marking-state)

Full Context

configure service vprn interface tos-marking-state

configure service ies interface tos-marking-state

Description

This command is used to alter the default trusted state to a non-trusted state. When unset or reverted to the trusted default, the ToS field are not remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set, in which case the egress network interface treats all VPRN and network IP interface as untrusted.

When the ingress interface is set to untrusted, all egress network IP interfaces remark IP packets received on the network interface according to the egress marking definitions on each network interface. The egress network remarking rules also apply to the ToS field of IP packets routed using IGP shortcuts (tunneled to a remote next-hop). However, the tunnel QoS markings are always derived from the egress network QoS definitions.

Egress marking and remarking is based on the internal forwarding class and profile state of the packet once it reaches the egress interface. The forwarding class is derived from ingress classification functions. The profile of a packet is either derived from ingress classification or ingress policing.

The default marking state for network IP interfaces is trusted. This is equivalent to declaring no tos-marking-state on the network IP interface. When undefined or set to tos-marking-state trusted, the trusted state of the interface will not be displayed when using show config or show info unless the detail parameter is given. The **save config** command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified.

The **no** form of this command restores the trusted state to a network IP interface. This is equivalent to executing the tos-marking-state trusted command.

Default

tos-marking-state trusted

Parameters

trusted

The default prevents the ToS field to not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set.

untrusted

Specifies that all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface.

Platforms

7705 SAR Gen 2

tos-marking-state

Syntax

tos-marking-state {**trusted** | **untrusted**}

no tos-marking-state

Context

[Tree] (config>service>vprn>nw-if tos-marking-state)

Full Context

configure service vprn network-interface tos-marking-state

Description

This command is used to alter the default trusted state to a non-trusted state. When unset or reverted to the trusted default, the ToS field will not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set, in which case the egress network interface treats all VPRN and network IP interface as untrusted.

When the ingress interface is set to untrusted, all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface. The egress network remarking rules also apply to the ToS field of IP packets routed using IGP shortcuts (tunneled to a remote next-hop). However, the tunnel QoS markings are always derived from the egress network QoS definitions. Egress marking and remarking is based on the internal forwarding class and profile state of the packet once it reaches the egress interface. The forwarding class is derived from ingress classification functions. The profile of a packet is either derived from ingress classification or ingress policing.

The default marking state for network IP interfaces is trusted. This is equivalent to declaring no tos-marking-state on the network IP interface. When undefined or set to tos-marking-state trusted, the trusted state of the interface will not be displayed when using show config or show info unless the detail parameter is given. The **save config** command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified.

The **no** tos-marking-state command is used to restore the trusted state to a network IP interface. This is equivalent to executing the tos-marking-state trusted command.

Default

tos-marking-state trusted

Parameters

trusted

The default prevents the ToS field to not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set.

untrusted

Specifies that all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface.

Platforms

7705 SAR Gen 2

tos-marking-state

Syntax

tos-marking-state {**trusted** | **untrusted**}

no tos-marking-state

Context

[Tree] (config>router>if tos-marking-state)

Full Context

configure router interface tos-marking-state

Description

This command is used on a network IP interface to alter the default trusted state to a non-trusted state. When unset or reverted to the trusted default, the ToS field will not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set, in which case the egress network interface treats all IES and network IP interface as untrusted. When the ingress network IP interface is set to untrusted, all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface. The egress network remarking rules also apply to the ToS field of IP packets routed using IGP shortcuts (tunneled to a remote next-hop). However, the tunnel QoS markings are always derived from the egress network QoS definitions. Egress marking and remarking is based on the internal forwarding class and profile state of the packet once it reaches the egress interface. The forwarding class is derived from ingress classification functions. The profile of a packet is either derived from ingress classification or ingress policing. The default marking state for network IP interfaces is trusted. This is equivalent to declaring no tos-marking-state on the network IP interface. When undefined or set to tos-marking-state trusted, the trusted state of the interface will not be displayed when using show config or show info unless the detail parameter is given. The **save config** command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified.

The **no** form of this command is used to restore the trusted state to a network IP interface. This is equivalent to executing the tos-marking-state trusted command.

Default

tos-marking-state trusted

Parameters

trusted

Specifies that the default prevents the ToS field to not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set

untrusted

Specifies that all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface.

Platforms

7705 SAR Gen 2

30.87 traceroute

traceroute

Syntax

Context

[Tree] (traceroute)

Full Context

traceroute

Description

This command determines the route to a destination address. DNS lookups for the responding hosts are enabled by default.

Parameters

candidate-path

Specifies a candidate path of the SRv6 policy to traceroute. The candidate path does not need to be the currently active candidate path.

dest-port-udp-fixed

Specifies that the destination UDP port number should not increment with each packet transmitted. By default, the UDP traceroute starts with destination UDP port 33434 and each subsequent packet sent to this destination UDP port increases by 1. The next packet uses UDP seat port 33435, the next 33436, and so on.

For a UDP test, this parameter prevents the per-transmitted packet increment of the destination UDP port number. The TCP protocol does not increment the destination TCP port, using a single destination TCP port for all traceroute packets for the test.

decode

Perform additional original datagram parsing functions. This parameter must be used with the **detail** parameter.

detail

Specifies to display additional information about the resulting packet.

distinguisher

Specifies the distinguisher of the SRv6 policy candidate path to send the traceroute probe on. This parameter must be configured if **protocol-owner** is configured to **bgp**.

Values 1 to 4294967295

dns-name

Specifies the DNS name, up to 63 characters, of the far-end device on which to send the traceroute request message.

endpoint *ipv6-address*

Specifies an SRv6 policy for a specific endpoint as the target of the traceroute.

Values

ipv6-address:	x:x:x:x:x:x:x:x	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 to FFFF]H
	d:	[0 to 255]D

ip-address

Specifies the far-end IP address on which to send the traceroute request message in dotted decimal notation.

Values

ipv4-address:	a.b.c.d	
ipv6-address:	x:x:x:x:x:x:x:x	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 to FFFF]H
	d:	[0 to 255]D

max-ttl

Specifies the maximum Time-To-Live (TTL) value to include in the traceroute request, expressed as a decimal integer.

Values 1 to 255

Default 30

milliseconds

Specifies the time in milliseconds to wait for a response to a probe, expressed as a decimal integer.

Values 1 to 60000

Default 5000

min-ttl

Specifies the IP TTL in the initial traceroute packet to target a specific node or starting node along the path.

Values 1 to 255

Default 1

no-dns

Specifies that, when the **no-dns** keyword is specified, DNS lookups of the responding hosts are not performed, and only the IP addresses are printed.

original-datagram

Parse the returned original datagram including any IPv6 and SRH header information.

pad-size

Specifies the number of bytes added to the UDP or TCP payload.

Values 0 to 9786

Default 0

port-number

Specifies the transport protocol destination port number.

Values 1 to 65535

Default 33434

preference

Specifies the preference of the SRv6 policy candidate path to send the traceroute probe on.

Values 0 to 4294967295

Default 100

probes-per-hop

Specifies the number of probes per hop.

Values 1 to 10

Default 3

protocol-owner

Specifies the protocol owner of the SRv6 policy candidate path to traceroute.

Values bgp — Specifies a BGP SRv6 policy.

static — Specifies a locally configured static SRv6 policy.

protocol udp | tcp

Sets the transport protocol for the traceroute packet. The TCP protocol is silently discarded on a targeted VRPN service. VPRN services only respond to UDP traceroutes.

Default udp

router-or-service

Specifies the routing instance or service, by number. The *router-instance* parameter is the preferred parameter to specify the router or service.

Values router-name: Base, management, vpls-management

vprn-svc-id: 1 to 2147483647

Default Base

router-instance

Specifies the preferred method for entering a service name. Stored as the service name, this is the only service-linking function allowed for both mixed-mode and model-driven configuration modes.

Values router-name: Base, management, vpls-management

vprn-svc-name: up to 64 characters

service-name

Specifies the alias function that allows the service name to be used, converted, and stored as service ID.

source *ip-address*

Specifies the source IP address to use as the source of the probe packets, in dotted decimal notation. If the IP address is not one of the device's interfaces, an error is returned.

Values

ipv4-address:	a.b.c.d	
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 to FFFF]H
	d:	[0 to 255]D

type-of-service

Specifies the Type-of-Service (ToS) bits in the IP header of the probe packets, expressed as a decimal integer.

Values 0 to 255

Default 0

srv6-policy

Keyword to specify that the traceroute probe is applied to an SRv6 policy matching a specific color and endpoint. The traceroute probe may optionally be targeted at a specific segment list of the SRv6 policy. When the segment list is not specified, the traceroute probe is sent on the lowest available segment list.

color-id

Specifies the SRv6 policy color ID.

Values 0 to 4294967295

segment-list

Specifies the SRv6 policy segment list to trace.

Values 1 to 32

Platforms

7705 SAR Gen 2

Output

Table 139: ICMPv4 Type 3 symbols in CLI, Table 140: ICMPv6 Type 1 symbols in CLI, and Table 141: ICMPv6 Type 2 symbols in CLI describe the ICMPv4 Type 3, and the ICMPv6 Type 1 and 2 symbols in the CLI outputs. For references without a symbol in the form !<code>, see www.iana.org/assignments/icmp-parameters/icmp-parameters.xhtml.

The following output is an example of traceroute for an IPv4 prefix.

Output Example

```
A:node-2# traceroute 192.168.xx.xx4
traceroute to 192.168.xx.xx4, 30 hops max, 40 byte packets
 1  192.168.xx.xx4 0.000 ms  0.000 ms  0.000 ms
```

The following output is an example of traceroute for an IPv4 prefix resolved to an IPv4 SR policy with ICMP tunneling enabled.

Output Example

```
A:node-2# traceroute 11.21.1.6 detail no-dns
traceroute to 11.21.1.6, 30 hops max, 40 byte packets
  1   1  10.10.11.3  3.36 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28303, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28306, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =  524283, Exp = 7, TTL =   1, S = 1
  1   2  10.10.11.3  3.68 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28303, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28306, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =  524283, Exp = 7, TTL =   1, S = 1
  1   3  10.10.11.3  4.18 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28303, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28306, Exp = 7, TTL =   1, S = 0
```

```
            entry  3:  MPLS Label =  524283, Exp = 7, TTL =   1, S = 1
  2   1  10.10.10.5  3.77 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28506, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =  524283, Exp = 7, TTL =   2, S = 1
  2   2  10.10.10.5  8.02 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28506, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =  524283, Exp = 7, TTL =   2, S = 1
  2   3  10.10.10.5  4.72 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28506, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =  524283, Exp = 7, TTL =   2, S = 1
  3   1  11.21.1.6  5.33 ms
  3   2  11.21.1.6  4.77 ms
  3   3  11.21.1.6  4.07 ms
```

The following output is an example of traceroute for an IPv6 prefix resolved to an IPv4 SR policy with ICMP tunneling enabled.

Output Example

```
A:node-2# traceroute fc00::b15:106 detail no-dns
traceroute to fc00::b15:106, 30 hops max, 60 byte packets
  1   1  fc00::a0a:b03  3.41 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28303, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28306, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =       2, Exp = 7, TTL =   1, S = 1
  1   2  fc00::a0a:b03  2.58 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28303, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28306, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =       2, Exp = 7, TTL =   1, S = 1
  1   3  fc00::a0a:b03  3.90 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28303, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28306, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =       2, Exp = 7, TTL =   1, S = 1
  2   1  fc00::a0a:a05  4.65 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28506, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =       2, Exp = 7, TTL =   2, S = 1
  2   2  fc00::a0a:a05  4.85 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28506, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =       2, Exp = 7, TTL =   2, S = 1
  2   3  fc00::a0a:a05  4.78 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28506, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =       2, Exp = 7, TTL =   2, S = 1
  3   1  fc00::b15:106  2.89 ms
  3   2  fc00::b15:106  3.58 ms
  3   3  fc00::b15:106  4.15 ms
```

The following output is an example of traceroute for an IPv6 prefix resolved to an IPv6 SR-OSPF3 tunnel with ICMP tunneling enabled.

Output Example

```
A:node-2# traceroute fc00::b14:106 detail
traceroute to fc00::b14:106, 30 hops max, 60 byte packets
  1   1  fc00::a0a:402  (fc00::a0a:402)  4.38 ms
```

```
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   29266, Exp = 7, TTL =   1, S = 1
  1   2  fc00::a0a:402  (fc00::a0a:402)  3.42 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   29266, Exp = 7, TTL =   1, S = 1
  1   3  fc00::a0a:402  (fc00::a0a:402)  4.19 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   29266, Exp = 7, TTL =   1, S = 1
  2   1  fc00::a0a:904  (fc00::a0a:904)  4.05 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   29466, Exp = 7, TTL =   1, S = 1
  2   2  fc00::a0a:904  (fc00::a0a:904)  3.62 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   29466, Exp = 7, TTL =   1, S = 1
  2   3  fc00::a0a:904  (fc00::a0a:904)  4.64 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   29466, Exp = 7, TTL =   1, S = 1
  3   1  fc00::b14:106  (fc00::b14:106)  3.35 ms
  3   2  fc00::b14:106  (fc00::b14:106)  4.02 ms
  3   3  fc00::b14:106  (fc00::b14:106)  3.30 ms
```

The following output is an example of traceroute for a label-ipv4 prefix resolved to an IPv6 SR-TE LSP with ICMP tunneling enabled (requires IPv4 system address).

Output Example

```
A:node-2# traceroute 11.21.1.1 source 11.21.1.6 detail
traceroute to 11.21.1.1 from 11.21.1.6, 30 hops max, 40 byte packets
  1   1  10.20.1.4  (10.20.1.4)  4.96 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =  524270, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =  524285, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =  524236, Exp = 7, TTL =   1, S = 1
  1   2  10.20.1.4  (10.20.1.4)  5.35 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =  524270, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =  524285, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =  524236, Exp = 7, TTL =   1, S = 1
  1   3  10.20.1.4  (10.20.1.4)  5.43 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =  524270, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =  524285, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =  524236, Exp = 7, TTL =   1, S = 1
  2   1  10.20.1.2  (10.20.1.2)  4.72 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =  524285, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =  524236, Exp = 7, TTL =   2, S = 1
  2   2  10.20.1.2  (10.20.1.2)  5.71 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =  524285, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =  524236, Exp = 7, TTL =   2, S = 1
  2   3  10.20.1.2  (10.20.1.2)  5.03 ms
         returned MPLS Label Stack Object
            entry  1:  MPLS Label =  524285, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =  524236, Exp = 7, TTL =   2, S = 1
  3   1  11.21.1.1  (11.21.1.1)  3.51 ms
  3   2  11.21.1.1  (11.21.1.1)  3.91 ms
  3   3  11.21.1.1  (11.21.1.1)  3.09 ms
```

The following output is an example of traceroute for a label-ipv6 prefix resolved to an IPv4 SR-TE LSP with ICMP tunneling enabled.

Output Example

```
A:node-2# traceroute fc00::b15:101 detail
traceroute to fc00::b15:101, 30 hops max, 60 byte packets
  1   1  fc00::a0a:404  (fc00::a0a:404)  3.36 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =  524270, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =  524285, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =       2, Exp = 7, TTL =   1, S = 1
  1   2  fc00::a0a:404  (fc00::a0a:404)  3.46 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =  524270, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =  524285, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =       2, Exp = 7, TTL =   1, S = 1
  1   3  fc00::a0a:404  (fc00::a0a:404)  3.77 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =  524270, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =  524285, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =       2, Exp = 7, TTL =   1, S = 1
  2   1  fc00::a0a:102  (fc00::a0a:102)  4.54 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =  524285, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =       2, Exp = 7, TTL =   2, S = 1
  2   2  fc00::a0a:102  (fc00::a0a:102)  4.70 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =  524285, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =       2, Exp = 7, TTL =   2, S = 1
  2   3  fc00::a0a:102  (fc00::a0a:102)  3.63 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =  524285, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =       2, Exp = 7, TTL =   2, S = 1
  3   1  fc00::b15:101  (fc00::b15:101)  3.40 ms
  3   2  fc00::b15:101  (fc00::b15:101)  3.15 ms
  3   3  fc00::b15:101  (fc00::b15:101)  3.23 ms
```

The following output is an example of traceroute for a vpn-ipv4 prefix resolved to an IPv6 SR-TE LSP with ICMP tunneling enabled (requires IPv4 system address).

Output Example

```
A:node-2# traceroute router-instance "vprn.sr-te.4" 1.0.4.1 source 6.0.4.1 detail
traceroute to 1.0.4.1 from 6.0.4.1, 30 hops max, 40 byte packets
  1   1  10.20.1.4  (10.20.1.4)  5.03 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28462, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28261, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =  524241, Exp = 7, TTL =   1, S = 1
  1   2  10.20.1.4  (10.20.1.4)  4.52 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28462, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28261, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =  524241, Exp = 7, TTL =   1, S = 1
  1   3  10.20.1.4  (10.20.1.4)  5.61 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28462, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28261, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =  524241, Exp = 7, TTL =   1, S = 1
  2   1  10.20.1.2  (10.20.1.2)  5.38 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28262, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28261, Exp = 7, TTL =   2, S = 0
             entry  3:  MPLS Label =  524241, Exp = 7, TTL =   2, S = 1
  2   2  10.20.1.2  (10.20.1.2)  5.39 ms
```

```
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28262, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28261, Exp = 7, TTL =   2, S = 0
             entry  3:  MPLS Label =  524241, Exp = 7, TTL =   2, S = 1
  2   3  10.20.1.2  (10.20.1.2)  5.27 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28262, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28261, Exp = 7, TTL =   2, S = 0
             entry  3:  MPLS Label =  524241, Exp = 7, TTL =   2, S = 1
  3   1  1.0.4.1  (1.0.4.1)  4.09 ms
  3   2  1.0.4.1  (1.0.4.1)  4.47 ms
  3   3  1.0.4.1  (1.0.4.1)  4.13 ms
```

The following output is an example of traceroute for a vpn-ipv6 prefix resolved to an IPv6 SR-TE LSP with ICMP tunneling enabled.

Output Example

```
A:node-2# traceroute router 5004 fc00::100:401 detail
traceroute to fc00::100:401, 30 hops max, 60 byte packets
  1   1  fc00::a0a:404  (fc00::a0a:404)  5.45 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28462, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28261, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =  524241, Exp = 7, TTL =   1, S = 1
  1   2  fc00::a0a:404  (fc00::a0a:404)  5.14 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28462, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28261, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =  524241, Exp = 7, TTL =   1, S = 1
  1   3  fc00::a0a:404  (fc00::a0a:404)  5.31 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28462, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28261, Exp = 7, TTL =   1, S = 0
             entry  3:  MPLS Label =  524241, Exp = 7, TTL =   1, S = 1
  2   1  fc00::a0a:102  (fc00::a0a:102)  4.70 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28262, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28261, Exp = 7, TTL =   2, S = 0
             entry  3:  MPLS Label =  524241, Exp = 7, TTL =   2, S = 1
  2   2  fc00::a0a:102  (fc00::a0a:102)  5.20 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28262, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28261, Exp = 7, TTL =   2, S = 0
             entry  3:  MPLS Label =  524241, Exp = 7, TTL =   2, S = 1
  2   3  fc00::a0a:102  (fc00::a0a:102)  5.16 ms
         returned MPLS Label Stack Object
             entry  1:  MPLS Label =   28262, Exp = 7, TTL =   1, S = 0
             entry  2:  MPLS Label =   28261, Exp = 7, TTL =   2, S = 0
             entry  3:  MPLS Label =  524241, Exp = 7, TTL =   2, S = 1
  3   1  fc00::100:401  (fc00::100:401)  5.38 ms
  3   2  fc00::100:401  (fc00::100:401)  4.48 ms
  3   3  fc00::100:401  (fc00::100:401)  4.39 ms
```

The following output is an example of traceroute for an IPv4 prefix using the tcp and the detail options.

Output Example

Note: Reaching the destination and the port is closed on the destination.

```
A:node-2# traceroute 192.168.34.2 protocol tcp detail
```

```
traceroute to 192.168.34.2, 30 hops max, 40 byte packets
  1   1  192.168.13.2  (192.168.13.2)  0.755 ms
  1   2  192.168.13.2  (192.168.13.2)  0.913 ms
  1   3  192.168.13.2  (192.168.13.2)  0.928 ms
  2   1  192.168.34.2  (192.168.34.2)  1.19 ms (port closed)
  2   2  192.168.34.2  (192.168.34.2)  1.29 ms (port closed)
  2   3  192.168.34.2  (192.168.34.2)  1.59 ms (port closed)
```

The following output is an example of traceroute for an IPv4 prefix using the tcp and the detail options.

Output Example

Note: Reaching the destination and the port is open on the destination.

```
A:node-2# traceroute 192.168.34.2 protocol tcp dest-port 862 detail
traceroute to 192.168.34.2, 30 hops max, 40 byte packets
  1   1  192.168.13.2  (192.168.13.2)  0.915 ms
  1   2  192.168.13.2  (192.168.13.2)  0.861 ms
  1   3  192.168.13.2  (192.168.13.2)  0.825 ms
  2   1  192.168.34.2  (192.168.34.2)  1.42 ms (port open)
  2   2  192.168.34.2  (192.168.34.2)  1.27 ms (port open)
  2   3  192.168.34.2  (192.168.34.2)  1.52 ms (port open)
```

The following output is an example of traceroute of an SRv6 SID using the **decode original-datagram** option.

Output Example

```
A:node-2# traceroute 2002:abcd:1100:102:1:: detail decode original-datagram probe-count 1
traceroute to 2002:abcd:1100:102:1::, 30 hops max, 60 byte packets
  1   1  2001:100:4:12::4  (2001:100:4:12::4)  1.23 ms
         Original Datagram
             IPv6 Header, Hop Limit 1, DSCP be
                 SA = 2001:1:1:1::112, DA = 2002:abcd:1100:102:1::
  2   1  2001:100:3:4::3  (2001:100:3:4::3)  2.25 ms
         Original Datagram
             IPv6 Header, Hop Limit 1, DSCP be
                 SA = 2001:1:1:1::112, DA = 2002:abcd:1100:101:1::
             Segment Routing Header SRv6, Segments Left 1
                 Segment_List[0] = 2002:abcd:1100:102:1::
  3   1  2001:100:1:3::1  (2001:100:1:3::1)  3.21 ms
         Original Datagram
             IPv6 Header, Hop Limit 1, DSCP be
                 SA = 2001:1:1:1::112, DA = 2002:abcd:1100:101:1::
             Segment Routing Header SRv6, Segments Left 1
                 Segment_List[0] = 2002:abcd:1100:102:1::
  4   1  2001:1:1:1::102  (2001:1:1:1::102)  9.16 ms
         Original Datagram
             IPv6 Header, Hop Limit 1, DSCP be
                 SA = 2001:1:1:1::112, DA = 2002:abcd:1100:102:1::
             Segment Routing Header SRv6, Segments Left 0
                 Segment_List[0] = 2002:abcd:1100:102:1::
```

The following output is an example of traceroute of an SRv6 policy.

Output Example

```
A:node-2# traceroute srv6-policy color 10 endpoint 6:6:6:6::86 probe-count 1
```

```
traceroute srv6-policy color 10 endpoint 6:6:6:6::86, 30 hops max, 60 byte packets (excluding
 SRH)
  1  fc00::a0a:203 (fc00::a0a:203)    2.76 ms
  2  fc00::a0a:505 (fc00::a0a:505)    5.11 ms
  3  6:6:6:6::86 (6:6:6:6::86)    6.18 ms
```

The following output is an example of traceroute of an SRv6 policy using the **decode original-datagram** option.

Output Example

```
A:node-2# traceroute srv6-policy color 10 endpoint 6:6:6:6::86 probe-count 1 detail decode
 original-datagram
traceroute srv6-policy color 10 endpoint 6:6:6:6::86, 30 hops max, 60 byte packets (excluding
 SRH)
  1   1  fc00::a0a:203  (fc00::a0a:203)  2.70 ms
         Original Datagram
             IPv6 Header, Hop Limit 1, DSCP be
                 SA = 1:1:1:1::61, DA = 3:3:3:3:0:a::
             Segment Routing Header SRv6, Segments Left 2
                 Segment_List[0] = 6:6:6:6::86
                 Segment_List[1] = 5:5:5:5:0:a::
  2   1  fc00::a0a:505  (fc00::a0a:505)  4.88 ms
         Original Datagram
             IPv6 Header, Hop Limit 1, DSCP be
                 SA = 1:1:1:1::61, DA = 5:5:5:5:0:a::
             Segment Routing Header SRv6, Segments Left 1
                 Segment_List[0] = 6:6:6:6::86
                 Segment_List[1] = 5:5:5:5:0:a::
  3   1  6:6:6:6::86  (6:6:6:6::86)  5.51 ms
         Original Datagram
             IPv6 Header, Hop Limit 1, DSCP be
                 SA = 1:1:1:1::61, DA = 6:6:6:6::86
```

The following output is an example of traceroute for a candidate path of an SRv6 policy.

Output Example

```
A:node-2# traceroute srv6-policy color 20 endpoint fc00::a14:106 probe-count 1 detail
 candidate-path protocol-owner static distinguisher 126 preference 100
traceroute srv6-policy color 20 endpoint fc00::a14:106 candidate-path protocol-owner static
 preference 100 distinguisher 126, 30 hops max, 60 byte packets (excluding SRH)
  1   1  fc00::a0a:203  (fc00::a0a:203)  2.87 ms
  2   1  fc00::a0a:505  (fc00::a0a:505)  4.58 ms
  3   1  fc00::a14:106  (fc00::a14:106)  6.28 ms
```

Table 139: ICMPv4 Type 3 symbols in CLI

Symbol	Description	Code
!N	Destination Network Unreachable	0
!P	Destination Protocol Unreachable	2
!	Destination Port Unreachable	3
!F-mtu	Fragmentation Needed and Don't Fragment was Set	4
!S	Source Route Failed	5

Symbol	Description	Code
!X	Communication Administratively Prohibited	13
!V	Host Precedence Violation	14
!C	Precedence Cutoff In Effect	15

Table 140: ICMPv6 Type 1 symbols in CLI

Symbol	Description	Code
!N	No Route to Destination	0
!H	Destination Address Unreachable	3
!	Destination Port Unreachable	4

Table 141: ICMPv6 Type 2 symbols in CLI

Symbol	Description	Code
!F-mtu	MTU Exceeded - Fragmentation Required	0

30.88 traceroute-reply

traceroute-reply

Syntax

[**no**] **traceroute-reply**

Context

[Tree] (config>service>ies>if>ipv6>vrrp traceroute-reply)

Full Context

configure service ies interface ipv6 vrrp traceroute-reply

Description

This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.

When this command is enabled, a non-owner master can reply to traceroute requests directed to the virtual router instance IP addresses.

A non-owner backup virtual router never responds to such traceroute requests regardless of the **traceroute-reply** status.

Default

no traceroute-reply

Platforms

7705 SAR Gen 2

traceroute-reply

Syntax

[**no**] **traceroute-reply**

Context

[Tree] (config>service>ies>if>vrrp traceroute-reply)

Full Context

configure service ies interface vrrp traceroute-reply

Description

This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.

When this command is enabled, a non-owner master can reply to traceroute requests directed to the virtual router instance IP addresses.

A non-owner backup virtual router never responds to such traceroute requests regardless of the **traceroute-reply** status.

Default

no traceroute-reply

Platforms

7705 SAR Gen 2

traceroute-reply

Syntax

[**no**] **traceroute-reply**

Context

[Tree] (config>service>vprn>if>ipv6>vrrp traceroute-reply)

[Tree] (config>service>vprn>if>vrrp traceroute-reply)

Full Context

configure service vprn interface ipv6 vrrp traceroute-reply

configure service vprn interface vrrp traceroute-reply

Description

This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.

When this command is enabled, a non-owner master can reply to traceroute requests directed to the virtual router instance IP addresses.

A non-owner backup virtual router never responds to such traceroute requests regardless of the **trace-route-reply** status.

Default

no traceroute-reply

Platforms

7705 SAR Gen 2

traceroute-reply

Syntax

[**no**] **traceroute-reply**

Context

[Tree] (config>router>if>ipv6>vrrp traceroute-reply)

[Tree] (config>router>if>vrrp traceroute-reply)

Full Context

configure router interface ipv6 vrrp traceroute-reply

configure router interface vrrp traceroute-reply

Description

This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.

When this command is enabled, a non-owner master can reply to traceroute requests directed to the virtual router instance IP addresses.

A non-owner backup virtual router never responds to such traceroute requests regardless of the **trace-route-reply** status.

Traceroute must not have been disabled at the management security level (either on the parental IP interface or the source host address).

Default

no traceroute-reply

Platforms

7705 SAR Gen 2

30.89 tracking-support

tracking-support

Syntax

[**no**] **tracking-support**

Context

[Tree] (config>service>vprn>pim>if tracking-support)

Full Context

configure service vprn pim interface tracking-support

Description

This command sets the T bit in the LAN Prune Delay option of the Hello Message. This indicates the router's capability to disable Join message suppression.

Default

no tracking-support

Platforms

7705 SAR Gen 2

tracking-support

Syntax

[**no**] **tracking-support**

Context

[Tree] (config>router>pim>interface tracking-support)

Full Context

configure router pim interface tracking-support

Description

This command sets the T bit in the LAN Prune Delay option of the Hello Message. This indicates the router's capability to enable join message suppression. This capability allows for upstream routers to explicitly track join membership.

The **no** form of this command disables tracking support.

Default
no tracking-support

Platforms
7705 SAR Gen 2

30.90 traffic-engineering

traffic-engineering

Syntax
[**no**] **traffic-engineering**

Context
[Tree] (config>router>isis traffic-engineering)

Full Context
configure router isis traffic-engineering

Description
This command enables this IS-IS instance to advertise TE link attributes for RSVP-TE and SR-TE enabled interfaces.

Default
no traffic-engineering

Platforms
7705 SAR Gen 2

traffic-engineering

Syntax
[**no**] **traffic-engineering**

Context
[Tree] (config>router>ospf traffic-engineering)

Full Context
configure router ospf traffic-engineering

Description

This command enables the advertisement of the traffic engineering information for the router and its links.

Traffic engineering enables the router to perform route calculations constrained by nodes or links. The traffic engineering of this router are limited to calculations based on link and nodal constraints.

The **no** form of this command disables the advertisement of the traffic engineering information.

Default

no traffic-engineering

Platforms

7705 SAR Gen 2

30.91 traffic-engineering-options

traffic-engineering-options

Syntax

[**no**] **traffic-engineering-options**

Context

[Tree] (config>router>isis traffic-engineering-options)

Full Context

configure router isis traffic-engineering-options

Description

Commands in this context configure advanced traffic-engineering options.

The **no** form of this command deletes the context.

Default

no traffic-engineering-options

Platforms

7705 SAR Gen 2

traffic-engineering-options

Syntax

[**no**] **traffic-engineering-options**

Context

[Tree] (config>router>ospf traffic-engineering-options)

Full Context

configure router ospf traffic-engineering-options

Description

Commands in this context configure the advanced traffic-engineering options.

The **no** form of this command removes the context to configure the advanced traffic-engineering options.

Default

no traffic-engineering-options

Platforms

7705 SAR Gen 2

30.92 transceiver

transceiver

Syntax

transceiver

Context

[Tree] (config>port transceiver)

Full Context

configure port transceiver

Description

Commands in this context configure transceiver parameters.

Platforms

7705 SAR Gen 2

30.93 transform

transform

Syntax

transform *transform-id* [*transform-id*]

no transform

Context

[Tree] (config>ipsec>tnl-temp transform)

[Tree] (config>ipsec>trans-mode-prof>dyn transform)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>dyn transform)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>dyn transform)

[Tree] (config>router>if>ipsec>ipsec-tunnel>dyn transform)

Full Context

configure ipsec tunnel-template transform

configure ipsec ipsec-transport-mode-profile dynamic-keying transform

configure service vprn interface ipsec ipsec-tunnel dynamic-keying transform

configure service ies interface ipsec ipsec-tunnel dynamic-keying transform

configure router interface ipsec ipsec-tunnel dynamic-keying transform

Description

This command associates the IPsec transform sets allowed for this the CHILD_SA. A maximum of four transforms can be specified. The transforms are listed in decreasing order of preference (the first one specified is the most preferred).

The **no** form of this command removes the transform ID from the configuration.

Default

no transform

Parameters

transform-id

Specifies a number to identify a tranform used for CHILD_SA negotiation. Up to four transform ID can be specified.

Values 1 to 2048

Platforms

7705 SAR Gen 2

transform

Syntax

transform *transform-id* [*transform-id*]

no transform

Context

[Tree] (config>service>vprn>if>sap>ipsec-tun>dyn transform)

Full Context

configure service vprn interface sap ipsec-tunnel dynamic-keying transform

Description

This command associates the IPsec transform sets allowed for this tunnel. A maximum of four transforms can be specified. The transforms are listed in decreasing order of preference (the first one specified is the most preferred).

Default

no transform

Parameters

transform-id

Specifies the value used for transforms for dynamic keying.

Values 1 to 2048

Platforms

7705 SAR Gen 2

30.94 transit

transit

Syntax

transit [**inherit** | **all** | **vc-only** | **none**]

Context

[Tree] (config>service>vprn>ttl-propagate transit)

Full Context

configure service vprn ttl-propagate transit

Description

This command overrides the global configuration of the TTL propagation for in transit packets which are forwarded over a MPLS LSPs in a given VPRN service context.

The global configuration is performed under config>router>ttl-propagate>vprn-transit.

The default behavior for a given VPRN instance is to inherit the global configuration for the same command. The user can explicitly set the default behavior by configuring the inherit value.

Default

transit inherit

Parameters

inherit

specifies the TTL propagation behavior is inherited from the global configuration under config>router>ttl-propagate>vprn-transit.

none

specifies the TTL of the IP packet is not propagated into the VC label or labels in the transport label stack.

vc-only

specifies the TTL of the IP packet is propagated into the VC label and not into the labels. in the transport label stack

all

specifies the TTL of the IP packet is propagated into the VC label and all labels in the transport label stack.

Platforms

7705 SAR Gen 2

30.95 transit-delay

transit-delay

Syntax

transit-delay *seconds*

no transit-delay

Context

[Tree] (config>service>vprn>ospf3>area>virtual-link transit-delay)

[Tree] (config>service>vprn>ospf>area>sham-link transit-delay)

[Tree] (config>service>vprn>ospf>area>virtual-link transit-delay)

[Tree] (config>service>vprn>ospf3>area>if transit-delay)

[Tree] (config>service>vprn>ospf>area>if transit-delay)

Full Context

configure service vprn ospf3 area virtual-link transit-delay

configure service vprn ospf area sham-link transit-delay

configure service vprn ospf area virtual-link transit-delay

configure service vprn ospf3 area interface transit-delay

configure service vprn ospf area interface transit-delay

Description

This command configures the estimated time, in seconds, that it takes to transmit a LSA on the interface or virtual link or sham-link.

The **no** form of this command reverts to the default delay time.

Default

transit-delay 1

Parameters

seconds

The transit delay in seconds expressed as a decimal integer.

Values 0 to 3600

Platforms

7705 SAR Gen 2

transit-delay

Syntax

transit-delay *seconds*

no transit-delay

Context

[Tree] (config>router>ospf3>area>virtual-link transit-delay)

[Tree] (config>router>ospf>area>interface transit-delay)

[Tree] (config>router>ospf3>area>interface transit-delay)

[Tree] (config>router>ospf>area>virtual-link transit-delay)

Full Context

configure router ospf3 area virtual-link transit-delay

configure router ospf area interface transit-delay

configure router ospf3 area interface transit-delay

configure router ospf area virtual-link transit-delay

Description

This command configures the estimated time, in seconds, that it takes to transmit a link state advertisement (LSA) on the interface or virtual link.

The **no** form of this command reverts to the default delay time.

Default

transit-delay 1

Parameters

seconds

Specifies the transit delay in seconds expressed as a decimal integer.

Values 1 to 1800

Platforms

7705 SAR Gen 2

30.96 transmission-profile

transmission-profile

Syntax

transmission-profile *name*

no transmission-profile

Context

[Tree] (config>system>security>pki>ca-profile>ocsp transmission-profile)

Full Context

configure system security pki ca-profile ocsp transmission-profile

Description

This command specifies the transmission-profile for OCSP. When specified, this configuration overrides the **service** *service-id* or **service** *service-name* configured in the **config>system>security>pki>ca-profile>ocsp** context.

The **no** form of the command removes the profile name from the configuration.

Default

no transmission-profile

Parameters

name

Specifies the file transmission profile name, up to 32 characters.

Platforms

7705 SAR Gen 2

transmission-profile

Syntax

transmission-profile *name*

no transmission-profile

Context

[Tree] (config>system>security>pki>est-profile transmission-profile)

Full Context

configure system security pki est-profile transmission-profile

Description

This command specifies the transmission profile name created in the **config>system file-transmission-profile** context for the EST profile.

The **no** form of the command removes the name from the EST profile configuration.

Default

no transmission-profile

Parameters

name

Specifies the file transmission profile name, up to 32 characters.

Platforms

7705 SAR Gen 2

30.97 transmit-interval

transmit-interval

Syntax

transmit-interval *transmit-interval*

no transmit-interval

Context

[Tree] (config>router>bfd>bfd-template transmit-interval)

Full Context

configure router bfd bfd-template transmit-interval

Description

This command specifies the transmit timer used for BFD packets. If the template is used for a BFD session on an MPLS-TP LSP, then this timer is used for CC packets.

The **no** form of this command reverts to the default value.

Default

transmit-interval 100

Parameters

transmit-interval

Specifies the transmit interval. The minimum interval that can be configured is hardware dependent.

Values 10 ms to 100,000 ms in 1 ms intervals

Default 10 ms for CPM3 or higher; 1 second for other hardware

Platforms

7705 SAR Gen 2

30.98 transmit-period

transmit-period

Syntax

transmit-period *seconds*

no transmit-period

Context

[Tree] (config>port>ethernet>dot1x transmit-period)

Full Context

configure port ethernet dot1x transmit-period

Description

This command configures the period after which the router sends a new EAPOL request message.

The **no** form of this command returns the value to the default.

Default

transmit-period 30

Parameters

seconds

Specifies the server transmit period in seconds.

Values 1 to 3600

Platforms

7705 SAR Gen 2

30.99 transport

transport

Syntax

transport *transport-protocol*

no transport

Context

[Tree] (config>system>snmp transport)

Full Context

configure system snmp transport

Description

This command configures the transport protocol used by the SNMP agent.

The **no** form of this command removes the transport protocol.

Default

no transport

Parameters

transport-protocol

Specifies the transport protocol.

Values **udp** — Keyword to specify UDP only.
tcp — Keyword to specify TCP only.
both — Keyword to specify TCP and UDP.

Default udp

Platforms

7705 SAR Gen 2

30.100 transport-address

transport-address

Syntax

transport-address {**interface** | **system**}

no transport-address

Context

[Tree] (config>router>ldp>if-params>ipv4 transport-address)

[Tree] (config>router>ldp>if-params>if>ipv4 transport-address)

[Tree] (config>router>ldp>if-params>ipv6 transport-address)

[Tree] (config>router>ldp>if-params>if>ipv6 transport-address)

Full Context

configure router ldp interface-parameters ipv4 transport-address

configure router ldp interface-parameters interface ipv4 transport-address

configure router ldp interface-parameters ipv6 transport-address

configure router ldp interface-parameters interface ipv6 transport-address

Description

This command configures the transport address to be used when setting up the LDP TCP sessions. The transport address can be configured as **interface** or **system**. The transport address can be configured globally (applies to all LDP interfaces) or per interface. The most specific value is used.

With the transport-address command, you can set up the LDP interface to the connection which can be set to the interface address or the system address. However, there can be an issue of which address to use when there are parallel adjacencies. This situation can not only happen with parallel links, it could be a link and a targeted adjacency since targeted adjacencies request the session to be set up only to the system IP address.

The **transport-address** value should not be **interface** if multiple interfaces exist between two LDP neighbors. Depending on the first adjacency to be formed, the TCP endpoint is chosen. In other words, if one LDP interface is set up as **transport-address interface** and another for **transport-address system**, then, depending on which adjacency was set up first, the TCP endpoint addresses are determined. After that, because the hello contains the LSR ID, the LDP session can be checked to verify that it is set up and then match the adjacency to the session.

For any iLDP interface, as the **local-lsr-id** parameters is changed to **interface**, the **transport-address** configuration loses effectiveness. Since it will be ignored and the iLDP session will always use the relevant interface IP address as transport-address even though system is chosen.

The **no** form of this command, at the global level, sets the transport address to the default value.

The **no** form of this command, at the interface level, sets the transport address to the value defined under the global level.

Default

system

Parameters

interface

Specifies the IP interface address is used to set up the LDP session between neighbors. The transport address interface cannot be used if multiple interfaces exist between two neighbors, since only one LDP session is set up between two neighbors.

system

Specifies the system IP address is used to set up the LDP session between neighbors.

Platforms

7705 SAR Gen 2

30.101 transport-encryption

transport-encryption

Syntax

transport-encryption

Context

[Tree] (config>redundancy>multi-chassis>peer>sync transport-encryption)

Full Context

configure redundancy multi-chassis peer sync transport-encryption

Description

Commands in this context configure MCS applications that need to encrypt synchronized states for transportation .

Platforms

7705 SAR Gen 2

30.102 transport-tunnel

transport-tunnel

Syntax

transport-tunnel

Context

[Tree] (config>router>bgp>next-hop-res>labeled-routes transport-tunnel)

Full Context

configure router bgp next-hop-resolution labeled-routes transport-tunnel

Description

Commands in this context configure options for the next-hop resolution of BGP labeled routes (VPN-IP and labeled-unicast) using tunnels in TTM. The context allows the selection of different tunnel resolution options for different types of BGP labeled routes: label-unicast IPv4, label-unicast IPv6, and VPN-IP routes (both VPN-IPv4 and VPN-IPv6).

By default (if this context and the resolution options are not configured), these routes resolve only to LDP tunnels.

If the **resolution** option is explicitly set to **disabled**, the default binding to LDP tunnel resumes. If **resolution** is set to **any**, then any supported tunnel type is allowed and the selection is based on the lowest numerical TTM preference value.

Platforms

7705 SAR Gen 2

30.103 trap-gen

trap-gen

Syntax

trap-gen

Context

[Tree] (config>saa>test trap-gen)

Full Context

configure saa test trap-gen

Description

Commands in this context configure trap generation for the SAA test.

Platforms

7705 SAR Gen 2

30.104 trap-target

trap-target

Syntax

trap-target name address *ip-address* [**port port**] [**snmpv1** | **snmpv2c** | **snmpv3**] **notify-community** *communityName* | *snmpv3SecurityName* [**security-level** {**no-auth-no-privacy** | **auth-no-privacy** | **privacy**}] [**replay**]

no trap-target *name*

Context

[Tree] (config>service>vprn>log>snmp-trap-group trap-target)

Full Context

configure service vprn log snmp-trap-group trap-target

Description

This command adds/modifies a trap receiver and configures the operational parameters for the trap receiver. A trap reports significant events that occur on a network device such as errors or failures.

Before an SNMP trap can be issued to a trap receiver, the **log-id**, **snmp-trap-group**, and at least one **snmp-trap-group** must be configured.

The **snmp-trap-group** command is used to add or remove a trap receiver from an **snmp-trap-group**. The operational parameters specified in the command include:

- The IP address of the trap receiver
- The UDP port used to send the SNMP trap
- SNMP version
- SNMP community name for SNMPv1 and SNMPv2c receivers.
- Security name and level for SNMPv3 trap receivers.

A single **snmp-trap-group** *log-id* can have multiple trap-receivers. Each trap receiver can have different operational parameters.

An address can be configured as a trap receiver more than once as long as a different port is used for each instance.

To prevent resource limitations, only configure a maximum of 10 trap receivers.

If the same **trap-target** *name* **port** *port* parameter value is specified in more than one SNMP trap group, each trap destination should be configured with a different *notify-community* value. This allows a trap receiving an application, such as NMS, to reconcile a separate event sequence number stream for each router event log when multiple event logs are directed to the same IP address and port destination.

The **no** form of this command removes the SNMP trap receiver from the SNMP trap group.

Default

No SNMP trap targets are defined.

Parameters

name

specifies the name of the trap target up to 28 characters in length

address ***ip-address***

The IP address of the trap receiver in dotted decimal notation. Only one IP address destination can be specified per trap destination group.

Values

ipv4-address	a.b.c.d (host bits must be 0)
ipv6-address	x:x:x:x:x:x:x:x[-interface]

x:x:x:x:x:x:d.d.d.d[-interface]

x: [0 to FFFF]H

d: [0 to 255]D

interface: 32 characters maximum, mandatory for link local addresses

port

Specifies the destination UDP port used to send traps to the destination, expressed as a decimal integer. Only one port can be specified per **trap-target** statement. If multiple traps need to be issued to the same address then multiple ports must be configured.

Values 1 to 65535

Default 162

snmpv1* | *snmpv2c* | *snmpv3

Specifies the SNMP version format to use for traps sent to the trap receiver.

The keyword **snmpv1** selects the SNMP version 1 format. When specifying **snmpv1**, the **notify-community** must be configured for the proper SNMP community string that the trap receiver expects to be present in alarms and traps messages. If the SNMP version is changed from **snmpv3** to **snmpv1,** then the **notify-community** parameter must be changed to reflect the community string rather than the *security-name* that is used by **snmpv3**.

The keyword **snmpv2c** selects the SNMP version 2c format. When specifying **snmpv2c**, the **notify-community** must be configured for the proper SNMP community string that the trap receiver expects to be present in alarms and traps messages. If the SNMP version is changed from **snmpv3** to **snmpv2c,** then the **notify-community** parameter must be changed to reflect the community string rather than the *security-name* that is used by **snmpv3**.

The keyword **snmpv3** selects the SNMP version 3 format. When specifying **snmpv3**, the **notify-community** must be configured for the SNMP *security-name*. If the SNMP version is changed from **snmpv1** or **snmpv2c** to **snmpv3**, then the **notify-community** parameter must be changed to reflect the *security-name* rather than the community string used by **snmpv1** or **snmpv2c**.

Pre-existing conditions are checked before the snmpv3SecurityName is accepted. These are:

- The username must be configured.
- The v3 access group must be configured.
- The v3 notification view must be configured.

Values snmpv1, snmpv2c, snmpv3

Default snmpv3

notify-community community | security-name

Specifies the community string for **snmpv1** or **snmpv2c** or the **snmpv3** *security-name*. If no **notify-community** is configured, then no alarms nor traps will be issued for the trap destination. If the SNMP version is modified, the **notify-community** must be changed to the proper form for the SNMP version.

community

The community string as required by the **snmpv1** or **snmpv2c** trap receiver. Allowed values are any string up to 31 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (for example, #, $, spaces), the entire string must be enclosed within double quotes.

security-name

The *security-name* as defined in the config>system>security>user context for SNMP v3. The *security-name* can be an ASCII string up to 31 characters in length.

security-level {no-auth-no-privacy | auth-no-privacy | privacy}

Specifies the required authentication and privacy levels required to access the views configured on this node when configuring an **snmpv3** trap receiver.

The keyword **no-auth-no-privacy** specifies no authentication and no privacy (encryption) are required.

The keyword **auth-no-privacy** specifies authentication is required but no privacy (encryption) is required. When this option is configured the *security-name* must be configured for **authentication**.

The keyword **privacy** specifies both authentication and privacy (encryption) is required. When this option is configured the *security-name* must be configured for **authentication** and **privacy**.

Values no-auth-no-privacy, auth-no-privacy, privacy

Default no-auth-no-privacy. This parameter can only be configured if SNMPv3 is also configured.

replay

Enable replay of missed events to target. If replay is applied to an SNMP trap target address, the address is monitored for reachability. Reachability is determined by whether or not there is a route in the routing table by which the target address can be reached. Before sending a trap to a target address, the SNMP module asks the PIP module if there is either an in-band or out-of-band route to the target address. If there is no route to the SNMP target address, the SNMP module saves the sequence-id of the first event that will be missed by the trap target. When the routing table changes again so that there is now a route by which the SNMP target address can be reached, the SNMP module replays (for example, retransmits) all events generated to the SNMP notification log while the target address was removed from the route table. Because of route table change convergence time, it is possible that one or more events may be lost at the beginning or end of a replay sequence. The cold-start-wait and route-recovery-wait timers under config>log>app-route-notifications can help reduce the probability of lost events.

Platforms

7705 SAR Gen 2

trap-target

Syntax

trap-target *name* [**address** *ip-address*] [**port** *port*] [**snmpv1** | **snmpv2c** | **snmpv3**] **notify-community** *communityName* | *snmpv3SecurityName* [**security-level** {**no-auth-no-privacy** | **auth-no-privacy** | **privacy**}] [**replay**]

no trap-target *name*

Context

[Tree] (config>log>snmp-trap-group trap-target)

Full Context

configure log snmp-trap-group trap-target

Description

This command configures a trap receiver and configures the operational parameters for the trap receiver. A trap reports significant events that occur on a network device such as errors or failures.

Before an SNMP trap can be issued to a trap receiver, the **log-id**, **snmp-trap-group** and at least one **trap-target** must be configured.

The **trap-target** command is used to add/remove a trap receiver from an **snmp-trap-group**. The operational parameters specified in the command include:

- The IP address of the trap receiver
- The UDP port used to send the SNMP trap
- SNMP version
- SNMP community name for SNMPv1 and SNMPv2c receivers.
- Security name and level for SNMPv3 trap receivers.

A single **snmp-trap-group** *log-id* can have multiple trap-receivers. Each trap receiver can have different operational parameters.

An address can be configured as a trap receiver more than once as long as a different port is used for each instance.

To prevent resource limitations, only configure a maximum of 10 trap receivers.

Note:
If the same **trap-target** *name* **port** *port* parameter value is specified in more than one SNMP trap group, each trap destination should be configured with a different *notify-community* value. This allows a trap receiving an application, such as NMS, to reconcile a separate event sequence number stream for each router event log when multiple event logs are directed to the same IP address and port destination.

The **no** form of this command removes the SNMP trap receiver from the SNMP trap group.

Parameters

name

Specifies the name of the trap target, up to 28 characters.

ip-address

Specifies the IP address of the trap receiver in dotted decimal notation. Only one IP address destination can be specified per trap destination group.

Values	ipv4-address	a.b.c.d (host bits must be 0)
	ipv6-address	x:x:x:x:x:x:x:x[-interface]
		x:x:x:x:x:x:d.d.d.d[-interface]
		x: [0..FFFF]H
		d: [0..255]D
		interface: 32 characters maximum, mandatory for link local addresses

port

Specifies the destination UDP port used for sending traps to the destination, expressed as a decimal integer. Only one port can be specified per **trap-target** statement. If multiple traps need to be issued to the same address then multiple ports must be configured.

Default 162

Values 1 **to** 65535

snmpv1* | *snmpv2c* | *snmpv3

Specifies the SNMP version format to use for traps sent to the trap receiver.

The keyword **snmpv1** selects the SNMP version 1 format. When specifying **snmpv1**, the **notify-community** must be configured for the proper SNMP community string that the trap receiver expects to be present in alarms and traps messages. If the SNMP version is changed from **snmpv3** to **snmpv1,** then the **notify-community** parameter must be changed to reflect the community string rather than the *security-name* that is used by **snmpv3**.

The keyword **snmpv2c** selects the SNMP version 2c format. When specifying **snmpv2c**, the **notify-community** must be configured for the proper SNMP community string that the trap receiver expects to be present in alarms and traps messages. If the SNMP version is changed from **snmpv3** to **snmpv2c,** then the **notify-community** parameter must be changed to reflect the community string rather than the *security-name* that is used by **snmpv3**.

The keyword **snmpv3** selects the SNMP version 3 format. When specifying **snmpv3**, the **notify-community** must be configured for the SNMP *security-name*. The security name is the name of a locally configured user. If the SNMP version is changed from **snmpv1** or **snmpv2c** to **snmpv3**, then the **notify-community** parameter must be changed to reflect the *security-name* rather than the community string used by **snmpv1** or **snmpv2c**.

The following conditions must all be met before traps will be issued using an SNMPv3 trap-target:

The user name must be configured, and must be configured with an snmp group that exists.

The v3 access group must be configured, or be one of the built-in SR OS views.

The v3 notification view must be configured, or be one of the built-in SR OS views.

Default snmpv3

Values snmpv1, snmpv2c, snmpv3

community* | *security-name

Specifies the community string for **snmpv1** or **snmpv2c** or the **snmpv3** *security-name*. If the **notify-community** is not configured, then no alarms or traps will be issued for the trap destination. If the SNMP version is modified, the **notify-community** must be changed to the proper form for the SNMP version.

community-name

Specifies the community string as required by the **snmpv1** or **snmpv2c** trap receiver. Allowed values are any string up to 31 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (for example, #, $, spaces), the entire string must be enclosed within double quotes.

security-name

For SNMPv3 trap targets, specifies the *security-name* as defined in the **config>system>security>user** context. The *security-name* can be an ASCII string up to 31 characters in length.

security-level {no-auth-no-privacy | auth-no-privacy | privacy}

Specifies the required authentication and privacy levels required to access the views configured on this node when configuring an **snmpv3** trap receiver.

The keyword **no-auth-no-privacy** specifies no authentication and no privacy (encryption) are required.

The keyword **auth-no-privacy** specifies authentication is required but no privacy (encryption) is required. When this option is configured the *security-name* must be configured for **authentication**.

The keyword **privacy** specifies both authentication and privacy (encryption) is required. When this option is configured the *security-name* must be configured for **authentication** and **privacy**.

Default **no-auth-no-privacy**. This parameter can only be configured if SNMPv3 is also configured.

Values no-auth-no-privacy, auth-no-privacy, privacy

replay

Enables the replay of missed events to target. If replay is applied to an SNMP trap target address, the address is monitored for reachability. Reachability is determined by whether or not there is a route in the routing table by which the target address can be reached. Before sending a trap to a target address, the SNMP module asks the PIP module if there

is either an in-band or out-of-band route to the target address. If there is no route to the SNMP target address, the SNMP module saves the sequence-id of the first event that will be missed by the trap target. When the routing table changes again so that there is now a route by which the SNMP target address can be reached, the SNMP module replays (for example, retransmits) all events generated to the SNMP notification log while the target address was removed from the route table.

Note:
Due to route table change convergence time, it is possible that one or more events may be lost at the beginning or end of a replay sequence. The cold-start-wait and route-recovery-wait timers under the **config>log>app-route-notifications** context can help reduce the probability of lost events.

Platforms

7705 SAR Gen 2

30.105 tree

tree

Syntax

tree [**detail**] [**flat**]

Context

[Tree] (tree)

Full Context

tree

Description

This command displays the command hierarchy structure of the current working context.

Parameters

detail

Displays parameter information for each command shown in the tree output.

flat

Displays the full context on each line.

Platforms

7705 SAR Gen 2

30.106 trigger

trigger

Syntax

[**no**] **trigger** [**neighbor** *ip-int-name* | *ip-address*]

Context

[Tree] (debug>router>rip trigger)

Full Context

debug router rip trigger

Description

This command enables debugging for RIP trigger updates.

Parameters

ip-int-name* | *ip-address

Debugs the RIP updates sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

trigger

Syntax

[**no**] **trigger** [**neighbor** *ip-int-name* | *ipv6-address*]

Context

[Tree] (debug>router>ripng trigger)

Full Context

debug router ripng trigger

Description

This command enables debugging for RIP trigger updates.

Parameters

ip-int-name*| *ipv6-address

Debugs the RIP updates sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

30.107 trigger-entry

trigger-entry

Syntax

[**no**] **trigger-entry** *entry-id*

Context

[Tree] (config>log>event-trigger>event trigger-entry)

Full Context

configure log event-trigger event trigger-entry

Description

This command configures an instance of a trigger for an EHS handler. A trigger entry binds a set of matching criteria for a log event to a particular handler. If the log event occurs in the system and matches the criteria configured in the associated log filter then the handler will be executed.

The **no** form of this command removes the specified trigger entry.

Parameters

entry-id

Specifies the identifier of the EHS event trigger entry.

Values 1 to 1500

Platforms

7705 SAR Gen 2

30.108 triggered-policy

triggered-policy

Syntax

[**no**] **triggered-policy**

Context

[Tree] (config>router triggered-policy)

Full Context

configure router triggered-policy

Description

This command triggers route policy re-evaluation.

By default, when a change is made to a policy in the **config router policy options** context and then committed, the change is effective immediately. There may be circumstances when the changes should or must be delayed; for example, if a policy change is implemented that would affect every BGP peer on a router, the consequences could be dramatic. It is more effective to control changes on a peer by peer basis.

If the **triggered-policy** command is enabled, and a given peer is established, and you want the peer to remain up, then, in order for a change to a route policy to take effect, a **clear** command with the *soft* or *soft-inbound* option must be used. In other words, when a **triggered-policy** is enabled, any routine policy change or policy assignment change within the protocol will not take effect until the protocol is reset or a clear command is issued to re-evaluate route policies; for example, **clear router bgp neighbor x.x.x.x soft**. This keeps the peer up and the change made to a route policy is applied only to that peer, or group of peers.

Default

no triggered-policy

Platforms

7705 SAR Gen 2

30.109 trust-anchor

trust-anchor

Syntax

[**no**] **trust-anchor** *ca-profile-name*

Context

[Tree] (config>ipsec>trust-anchor-profile trust-anchor)

Full Context

configure ipsec trust-anchor-profile trust-anchor

Description

This command specifies a CA profile as a trust anchor CA. Up to 8 multiple trust anchors can be specified in a single trust anchor profile.

The **no** form of this command removes the name from the configuration.

Parameters

ca-profile-name

Specifies the name of the trust anchor profile, up to 32 characters.

Platforms

7705 SAR Gen 2

trust-anchor

Syntax

[**no**] **trust-anchor** *ca-profile-name*

Context

[Tree] (config>system>security>tls>trust-anchor-profile trust-anchor)

Full Context

configure system security tls trust-anchor-profile trust-anchor

Description

This command configures a trust anchor with a CA profile used by the TLS profile. Up to eight CA profiles can be configured under the trust anchor. TLS will read the CA profiles one by one to try to authenticate the server certificate.

Parameters

ca-profile-name

Specifies the name of the TLS trust anchor, up to 32 characters.

Platforms

7705 SAR Gen 2

30.110 trust-anchor-profile

trust-anchor-profile

Syntax

trust-anchor-profile *name* [**create**]

no trust-anchor-profile *name*

Context

[Tree] (config>ipsec trust-anchor-profile)

Full Context

configure ipsec trust-anchor-profile

Description

This command specifies the trust anchor profile name for the IPsec tunnel or IPsec GW.

Default

no trust-anchor-profile

Parameters

name

Specifies the name of trust anchor profile up to 32 characters.

Platforms

7705 SAR Gen 2

trust-anchor-profile

Syntax

trust-anchor-profile *name*

no trust-anchor-profile

Context

[Tree] (config>service>ies>if>sap>ipsec-gw>cert trust-anchor-profile)

[Tree] (config>service>vprn>if>sap>ipsec-gw>cert trust-anchor-profile)

[Tree] (config>ipsec>trans-mode-prof>dyn>cert trust-anchor-profile)

[Tree] (config>router>if>ipsec>ipsec-tunnel>dyn>cert trust-anchor-profile)

[Tree] (config>service>ies>if>ipsec>ipsec-tunnel>dyn>cert trust-anchor-profile)

[Tree] (config>service>vprn>if>ipsec>ipsec-tunnel>dyn>cert trust-anchor-profile)

Full Context

configure service ies interface sap ipsec-gw cert trust-anchor-profile

configure service vprn interface sap ipsec-gw cert trust-anchor-profile

configure ipsec ipsec-transport-mode-profile dynamic-keying cert trust-anchor-profile

configure router interface ipsec ipsec-tunnel dynamic-keying cert trust-anchor-profile

configure service ies interface ipsec ipsec-tunnel dynamic-keying cert trust-anchor-profile

configure service vprn interface ipsec ipsec-tunnel dynamic-keying cert trust-anchor-profile

Description

This command specifies the name of trust anchor profile used for certificate authentication.

The **no** form of this command removes the name from the configuration.

Default

no trust-anchor-profile

Parameters

name

Specifies the name of trust anchor profile, up to 32 characters.

Platforms

7705 SAR Gen 2

trust-anchor-profile

Syntax

trust-anchor-profile *name*

no trust-anchor-profile

Context

[Tree] (config>system>security>tls>client-tls-profile trust-anchor-profile)

[Tree] (config>system>security>tls>server-tls-profile>authenticate-client trust-anchor-profile)

Full Context

configure system security tls client-tls-profile trust-anchor-profile

configure system security tls server-tls-profile authenticate-client trust-anchor-profile

Description

This command assigns the trust anchor used by this TLS profile to authenticate the server or client.

The **no** form of the command removes the configured trust anchor profile.

Parameters

name

Specifies the name of the trust anchor profile.

Platforms

7705 SAR Gen 2

trust-anchor-profile

Syntax

trust-anchor-profile *name* [**create**]

no trust-anchor-profile *name*

Context

[Tree] (config>system>security>tls trust-anchor-profile)

Full Context

configure system security tls trust-anchor-profile

Description

This command configures a trust anchor profile to be used in the TLS profile. The trust anchor is used for authentication of the server certificate.

Parameters

name

Specifies the name of the trust anchor profile, up to 32 characters.

create

Keyword used to create the trust anchor profile.

Platforms

7705 SAR Gen 2

30.111 trusted

trusted

Syntax

[**no**] **trusted**

Context

[Tree] (config>router>if>dhcp trusted)

[Tree] (config>service>vprn>if>dhcp trusted)

[Tree] (config>service>ies>if>dhcp trusted)

Full Context

configure router interface dhcp trusted

configure service vprn interface dhcp trusted

configure service ies interface dhcp trusted

Description

This command enables relaying untrusted packets. According to RFC 3046, *DHCP Relay Agent Information Option*, a DHCP request where the giaddr is 0.0.0.0 and which contains an Option 82 field in the packet, should be discarded, unless it arrives on a "trusted" circuit. If the **trusted** mode is enabled on an IP interface, the Relay Agent (the router) modifies the requested giaddr to be equal to the ingress interface and forward the request.

The **no** form of this command reverts to the default.

Default

no trusted

Platforms

7705 SAR Gen 2

30.112 trusted-mac-time

trusted-mac-time

Syntax

trusted-mac-time *range*

Context

[Tree] (config>service>vpls>bgp-evpn>mac-duplication trusted-mac-time)

Full Context

configure service vpls bgp-evpn mac-duplication trusted-mac-time

Description

This command determines how long a MAC address needs to stay in the FDB as type learned without being flushed or changed in its type so that the MAC is declared as trusted for the mac-duplication

procedures. If the MAC changes from SAP to SAP within the same VPLS service and node, the MAC does not reset its trusted MAC timer.

Default

trusted-mac-time 5

Parameters

range

Specifies the time, in minutes, before the MAC address can be flushed from the FDB.

Values 1 to 15

Platforms

7705 SAR Gen 2

30.113 ts-list

ts-list

Syntax

ts-list *list-name* [**create**]

no ts-list *list-name*

Context

[Tree] (config>ipsec ts-list)

Full Context

configure ipsec ts-list

Description

This command creates a new traffic selector (TS).

The **no** form of this command removes the list name from the configuration.

Parameters

list-name

Specifies the name of the TS-list.

Platforms

7705 SAR Gen 2

30.114 ts-location

ts-location

Syntax

ts-location *file-url*

no ts-location

Context

[Tree] (config>system>security>tech-support ts-location)

Full Context

configure system security tech-support ts-location

Description

The **ts-location** command is used (along with an automatic system generated file name) when no *file-url* parameter is provided for the **admin tech-support** command. If **no ts-location** is defined then the operator must provide a file-url with the **admin tech-support** command itself.

The directory specified for the ts-location is not auto-created by SR OS. The operator must ensure that it exists.

See the **admin tech-support** command for more details about the system generated file name.

Default

no ts-location

Parameters

file-url

Specifies the destination directory for auto-named tech-support files (when no *file-url* is specified with the **admin tech-support** command). The *file-url* for the **ts-location** must be a directory (no filename or extension). The root directory (for example, cf1:\) is blocked for local compact flash destinations. A sub-directory (for example, cf2:\tech-support) must be used if local cf is the location.

Values

local-url | *remote-url*	
local-url	[*cflash-id*/][*file-path*] 200 chars max, including cflash-id
	directory length 99 chars max each
remote-url	[ftp://*login:pswd@remote-locn*/][*file-path*]
	247 chars max

	directory length 99 chars max each
remote-locn	[*hostname* | *ipv4-address* | "["*ipv6-address*"]"]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[*-interface*]
	x:x:x:x:x:x:d.d.d.d[*-interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D
	interface - 32 chars max, for link local addresses
cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

Platforms

7705 SAR Gen 2

30.115 ts-negotiation

ts-negotiation

Syntax

ts-negotiation ts-list *list-name*

no ts-negotiation

Context

[Tree] (config>ipsec>client-db>client ts-negotiation)

Full Context

configure ipsec client-db client ts-negotiation

Description

This command specifies the traffic selector (TS) to be used for tunnel setup.

The no form of this command reverts to the default.

Default

no ts-negotiation

Parameters

list-name

Specifies the TS list used by this tunnel, up to 32 characters.

Platforms

7705 SAR Gen 2

ts-negotiation

Syntax

ts-negotiation ts-list *list-name*

no ts-negotiation

Context

[Tree] (config>service>ies>if>sap>ipsec-gw ts-negotiation)

[Tree] (config>service>vprn>if>sap>ipsec-gw ts-negotiation)

Full Context

configure service ies interface sap ipsec-gw ts-negotiation

configure service vprn interface sap ipsec-gw ts-negotiation

Description

This command enables the IKEv2 traffic selector negotiation with the specified ts-list.

Parameters

list-name

Specifies the ts-list name

Platforms

7705 SAR Gen 2

30.116 ttl

ttl

Syntax

ttl *label-ttl*

no ttl

Context

[Tree] (config>saa>test>type-multi-line>lsp-ping>sr-policy ttl)

[Tree] (config>saa>test>type-multi-line>lsp-ping ttl)

Full Context

configure saa test type-multi-line lsp-ping sr-policy ttl

configure saa test type-multi-line lsp-ping ttl

Description

This command configures a time-to-live value for the MPLS label.

The **no** form of this command reverts to the default value.

Default

ttl 255

Parameters

label-ttl

Specifies the time-to-live value.

Values 1 to 255

Default 255

Platforms

7705 SAR Gen 2

ttl

Syntax

ttl *time-to-live*

no ttl

Context

[Tree] (config>oam-pm>session>ip ttl)

Full Context

configure oam-pm session ip ttl

Description

This command defines the value of the TTL field of the packet header.

The **no** form of this command restores the default value.

Default

ttl 225

Parameters

time-to-live

Specifies the value to be used in the TTL field.

Values 1 to 255

Default 255

Platforms

7705 SAR Gen 2

30.117 ttl-expired

ttl-expired

Syntax

ttl-expired *number seconds*

no ttl-expired [*number seconds*]

Context

[Tree] (config>service>ies>if>icmp ttl-expired)

Full Context

configure service ies interface icmp ttl-expired

Description

This command configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface and reverts to the default values.

Default

ttl-expired 100 10

Parameters

number

The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

Values 10 to 2000

seconds

The time frame in seconds used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 to 60

Platforms

7705 SAR Gen 2

ttl-expired

Syntax

ttl-expired [*number seconds*]

no ttl-expired

Context

[Tree] (config>service>vprn>nw-if>icmp ttl-expired)

[Tree] (config>service>vprn>if>icmp ttl-expired)

Full Context

configure service vprn network-interface icmp ttl-expired

configure service vprn interface icmp ttl-expired

Description

This command configures the rate of Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default

ttl-expired 100 10

Parameters

number

Specifies the maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

Values 10 to 2000

seconds

Specifies the time frame in seconds used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 to 60

Platforms

7705 SAR Gen 2

ttl-expired

Syntax

ttl-expired [*number seconds*]

no ttl-expired

Context

[Tree] (config>router>if>icmp ttl-expired)

Full Context

configure router interface icmp ttl-expired

Description

This command configures the rate that Internet Control Message Protocol (ICMP) Time To Live (TTL) expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of TTL expired messages.

Default

ttl-expired 100 10 — Maximum of 100 TTL expired message in 10 seconds.

Parameters

number

The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. The *seconds* parameter must also be specified.

Values 10 to 2000

seconds

The time frame, in seconds, used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 to 60

Platforms

7705 SAR Gen 2

30.118 ttl-propagate

ttl-propagate

Syntax

ttl-propagate

Context

[Tree] (config>service>vprn ttl-propagate)

Full Context

configure service vprn ttl-propagate

Description

Commands in this context configure TTL propagation for transit and locally generated packets in a given VPRN routing context.

Platforms

7705 SAR Gen 2

ttl-propagate

Syntax

ttl-propagate

Context

[Tree] (config>router ttl-propagate)

Full Context

configure router ttl-propagate

Description

Commands in this context configure TTL propagation for transit and locally generated packets in the Global Routing Table (GRT) and VPRN routing contexts

Platforms

7705 SAR Gen 2

30.119 ttl-security

ttl-security

Syntax

ttl-security *min-ttl-value*

no ttl-security

Context

[Tree] (config>service>vprn>bgp>group>neighbor ttl-security)

[Tree] (config>service>vprn>bgp>group ttl-security)

Full Context

configure service vprn bgp group neighbor ttl-security

configure service vprn bgp group ttl-security

Description

Configure TTL security parameters for incoming packets.

Parameters

min-ttl-value

Specifies the minimum TTL value for an incoming BGP packet.

Values 1 to 255

Default 1

Platforms

7705 SAR Gen 2

ttl-security

Syntax

ttl-security *min-ttl-value*

no ttl-security

Context

[Tree] (config>system>login-control>telnet ttl-security)

[Tree] (config>system>login-control>ssh ttl-security)

[Tree] (config>router>ldp>tcp-session-params>peer-transport ttl-security)

[Tree] (config>router>bgp>group ttl-security)

[Tree] (config>router>bgp>group>neighbor ttl-security)

Full Context

configure system login-control telnet ttl-security

configure system login-control ssh ttl-security

configure router ldp tcp-session-parameters peer-transport ttl-security

configure router bgp group ttl-security

configure router bgp group neighbor ttl-security

Description

This command configures TTL security parameters for incoming packets. When the feature is enabled, LDP will accept incoming IP packets from a peer only if the TTL value in the packet is greater than or equal to the minimum TTL value configured for that peer. Per-peer-queueing must be enabled in order for TTL protection to operate.

The **no** form of this command disables TTL security.

Parameters

min-ttl-value

Specifies the minimum TTL value for an incoming BGP packet.

Values 1 to 255

Platforms

7705 SAR Gen 2

30.120 tunnel

tunnel

Syntax

tunnel *ipsec-tunnel-name* [**detail**] [**no-dpd-debug**] [**display-keys**]

no tunnel *ipsec-tunnel-name*

Context

[Tree] (debug>ipsec tunnel)

Full Context

debug ipsec tunnel

Description

This command enables debugging for specified IPsec tunnel.

Note:
Up to 16 IPsec tunnels are allowed, to enable debugging, at a time.

Parameters

ipsec-tunnel-name

Specifies the name of ipsec-tunnel, up to 32 characters.

detail

Displays detailed debug information.

no-dpd-debug

Stops logging IKEv1 and IKEv2 DPD events for less noise during debug.

display-keys

Specifies the IKE-SA and CHILD-SA keys for inclusion in the debug output.

Platforms

7705 SAR Gen 2

tunnel

Syntax

tunnel *name* **[create]**

no tunnel *name*

Context

[Tree] (config>system>grpc-tunnel tunnel)

Full Context

configure system grpc-tunnel tunnel

Description

Commands in this context configure gRPC tunnel parameters for the specified tunnel. There can be multiple tunnels to one or more destinations.

The **no** form of this command removes the specified gRPC tunnel.

Parameters

name

Specifies the tunnel name, up to 32 characters.

create

Keyword used to create a tunnel.

Platforms

7705 SAR Gen 2

tunnel

Syntax

tunnel

Context

[Tree] (config>oam-pm>session>ip tunnel)

Full Context

configure oam-pm session ip tunnel

Description

Commands in this context configure packet tunneling options for the session. This command and the **oam-pm session ip forwarding** command are mutually exclusive.

Platforms

7705 SAR Gen 2

30.121 tunnel-dot1q

tunnel-dot1q

Syntax

[**no**] **tunnel-dot1q**

Context

[Tree] (config>port>ethernet>dot1x tunnel-dot1q)

Full Context

configure port ethernet dot1x tunnel-dot1q

Description

This command configures the tunneling of single tagged (dot1q) dot1x packets arriving on the port. When enabled, the router extracts these packets to the CPM.

The **no** form of this command disables the tunelling of the dot1q dot 1x packets on the port.

Default

tunnel-dot1q

Platforms

7705 SAR Gen 2

30.122 tunnel-down-damp-time

tunnel-down-damp-time

Syntax

tunnel-down-damp-time *seconds*

no tunnel-down-damp-time

Context

[Tree] (config>router>ldp tunnel-down-damp-time)

Full Context

configure router ldp tunnel-down-damp-time

Description

This command specifies the time interval (in s), that LDP waits before posting a tunnel down event to the Tunnel Table Manager (TTM).

When LDP can no longer resolve a FEC and de-activates it, it de-programs the NHLFE in the data path. It will however delay deleting the LDP tunnel entry in the TTM until the tunnel-down-damp-time timer expires. This means users of the LDP tunnel, such as SDPs (all services) and BGP (L3 VPN), will not be notified immediately. Traffic is still blackholed because the forwarding engine NHLFE has been de-programmed.

If the FEC gets resolved before the tunnel-down-damp-time timer expires, then LDP programs the forwarding engine with the new NHLFE and performs a tunnel modify event in TTM updating the dampened entry in TTM with the new NHLFE information. If the FEC does not get resolved and the tunnel-down-damp-time timer expires, LDP posts a tunnel down event to TTM which deletes the LDP tunnel.

When there is an upper layer (user of LDP) which depends of LDP control plane for failover detection then label withdrawal delay and tunnel-down-damp-time options must be set to 0.

An example is pseudowire redundancy where the primary PW does not have its own fast failover detection mechanism and the node depends on LDP tunnel down event to activate the standby PW.

The **no** form of this command resumes the default value of this command.

Default

no tunnel-down-damp-time (which equals a value of 3 seconds)

Parameters

seconds

Specifies the time interval (in s), that LDP waits before posting a tunnel down event to the Tunnel Table Manager.

Platforms

7705 SAR Gen 2

30.123 tunnel-endpoint

tunnel-endpoint

Syntax

tunnel-endpoint [**tunnel-spf**] [**tunnel-leak** *ip-address*]

no tunnel-endpoint

Context

[Tree] (debug>router>isis tunnel-endpoint)

Full Context

debug router isis tunnel-endpoint

Description

This command enables debugging for an ISIS tunnel endpoint.

The **no** form of the command disables the debugging.

Parameters

tunnel-spf

Debugs tunnel SPF information.

ip-address

When specified, only packets with the specified address are debugged.

Platforms

7705 SAR Gen 2

tunnel-endpoint

Syntax

tunnel-endpoint [**tunnel-spf** *ip-address*] [**tunnel-leak** *ip-address*]

Context

[Tree] (debug>router>ospf3 tunnel-endpoint)

[Tree] (debug>router>ospf tunnel-endpoint)

Full Context

debug router ospf3 tunnel-endpoint

debug router ospf tunnel-endpoint

Description

This command enables debugging for OSPF tunnel endpoints.

Parameters

tunnel-spf

Specifies the tunnel SPF IP address.

tunnel-leak

Specifies the tunnel leak IP address.

ip-address

Specifies the IP address.

Platforms

7705 SAR Gen 2

30.124 tunnel-far-end

tunnel-far-end

Syntax

tunnel-far-end *ip-address* | *ipv6-address*

no tunnel-far-end [*ip-address* | *ipv6-address*]

Context

[Tree] (config>service>sdp tunnel-far-end)

Full Context

configure service sdp tunnel-far-end

Description

This command enables the user to specify an SDP tunnel destination address that is different from the configuration in the SDP far-end option. The SDP must be shutdown first to add or change the configuration of the **tunnel-far-end** option.

When this option is enabled, service packets are encapsulated using an LDP LSP with a FEC prefix matching the value entered in ip-address. By default, service packets are encapsulated using an LDP LSP with a FEC prefix matching the address entered in the SDP far-end option.

The T-LDP session to the remote PE is still targeted to the address configured under the **far-end option**. This means that targeted hello messages are sent to the far-end address, which is also the LSR-ID of the remote node. TCP based LDP messages, such as initialization and label mapping messages, are sent to the address specified in the transport-address field of the "hello" message received from the remote PE. This address can be the same as the remote PE LSR-ID, or a different address. This feature works, however, if the signaling option in the SDP is set to off instead of tldp, in which case, the service labels are statically configured.

This feature operates on an SDP of type LDP only. It can be used with VLL, VPLS, and VPRN services when an explicit binding to an SDP with the **tunnel-far-end** is specified. It also operates with a spoke interface on an IES or VPRN service. Finally, this feature operates with a BGP AD based VPLS service when the **use-provisioned-sdp** option is enabled in the pseudowire template.

This feature is not supported in an SDP of type MPLS when an RSVP LSP name is configured under the SDP. It also does not work with a mixed-lsp SDP.

The **no** form of this command disables the use of the **tunnel-far-end** option and returns to using the address specified in the far-end.

Default

no tunnel-far-end

Parameters

ip-address* | *ipv6-address

Specifies the system address of the far-end router for the SDP in dotted decimal notation.

Platforms

7705 SAR Gen 2

30.125 tunnel-group

tunnel-group

Syntax

tunnel-group *tunnel-group-id* [**create**]

tunnel-group *tunnel-group-id* **isa-scale-mode** *isa-scale-mode* [**create**]

no tunnel-group *tunnel-group-id*

Context

[Tree] (config>isa tunnel-group)

Full Context

configure isa tunnel-group

Description

This command allows a tunnel group to be created or edited. A tunnel group is a set of one or more MS-ISAs that support the origination and termination of IPsec and IP/GRE tunnels. All of the MS-ISAs in a tunnel group must have **isa-tunnel** as their configured mda-type.

The **no** form of this command deletes the specified tunnel group from the configuration

Parameters

tunnel-group-id

Identifies the tunnel group.

Values 1 to 16

isa-scale-mode

Defines the maximum number of tunnels (all types combined) which can be established on each ISA of the tunnel group and for the whole tunnel-group. When it is not explicitly specified, **isa-scale-mode** has a default value, but that value is different on different platforms.

Values tunnel-limit-2k, tunnel-limit-32k, tunnel-limit-64k, tunnel-limit-8, tunnel-limit-32

create

Mandatory keyword used when creating tunnel group in the ISA context. The create keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

tunnel-group

Syntax

tunnel-group *tunnel-group-id* [**create**]

no tunnel-group *tunnel-group-id*

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ipsec tunnel-group)

Full Context

configure redundancy multi-chassis peer mc-ipsec tunnel-group

Description

This command enables multi-chassis redundancy for specified tunnel-group; or enters an already configured tunnel-group context. The configured tunnel-group could failover independently.

The **no** form of this command removes the tunnel group ID from the configuration.

Parameters

tunnel-group-id

Specifies the tunnel-group identifier.

Values 1 to 16

Platforms

7705 SAR Gen 2

tunnel-group

Syntax

tunnel-group *tunnel-group-id* [**create**]

no tunnel-group *tunnel-group-id*

Context

[Tree] (config>redundancy>multi-chassis>peer>mc-ipsec tunnel-group)

Full Context

configure redundancy multi-chassis peer mc-ipsec tunnel-group

Description

This command enables multi-chassis redundancy for specified tunnel-group; or enters an already configured tunnel-group context. The configured tunnel-group could failover independently.

The **no** form of this command removes the tunnel group ID from the configuration.

Parameters

tunnel-group-id

Specifies the tunnel-group identifier.

Values 1 to 16

Platforms

7705 SAR Gen 2

tunnel-group

Syntax

tunnel-group *tunnel-group-id* **sync-tag** *tag-name* [**create**]

no tunnel-group *tunnel-group-id*

Context

[Tree] (config>redundancy>multi-chassis>peer>sync tunnel-group)

Full Context

configure redundancy multi-chassis peer sync tunnel-group

Description

This command enables multi-chassis synchronization of IPsec states of specified tunnel-groups with a peer. The **sync-tag** parameter is used to match corresponding tunnel-group on both peers. IPsec states will be synchronized between tunnel-groups with same sync-tag.

Parameters

tunnel-group-id

Specifies the ID of the tunnel group.

tag-name

Specifies the name of the sync-tag.

Platforms

7705 SAR Gen 2

tunnel-group

Syntax

tunnel-group *tunnel-group-id*

no tunnel-group

Context

[Tree] (config>redundancy>multi-chassis>ipsec-domain tunnel-group)

Full Context

configure redundancy multi-chassis ipsec-domain tunnel-group

Description

This command specifies the tunnel group ID for the IPsec domain.

The **no** form of this command removes the tunnel group ID from the configuration.

Default

no tunnel-group

Parameters

tunnel-group-id

Specifies the tunnel group ID, up to 64 characters.

Platforms

7705 SAR Gen 2

30.126 tunnel-interface

tunnel-interface

Syntax

[**no**] **tunnel-interface** {**rsvp-p2mp** *lsp-name* | **ldp-p2mp** *p2mp-id* **sender** *sender-address* [**root-node**]}

Context

[Tree] (config>router tunnel-interface)

Full Context

configure router tunnel-interface

Description

This command creates a tunnel interface associated with an RSVP P2MP LSP. IPv4 multicast packets are forwarded over the P2MP LSP at the ingress LER based on a static join configuration of the multicast group against the tunnel interface associated with the originating P2MP LSP. At the egress LER, packets of a multicast group are received from the P2MP LSP via a static assignment of the specific <S,G> to the tunnel interface associated with a terminating LSP.

At ingress LER, the tunnel interface identifier consists of a string of characters representing the LSP name for the RSVP P2MP LSP. The user can create one or more tunnel interfaces and associate each to a different RSVP P2MP LSP.

At egress LER, the tunnel interface identifier consists of a couple of string of characters representing the LSP name for the RSVP P2MP LSP followed by the system address of the ingress LER. The LSP name must correspond to a P2MP LSP name configured by the user at the ingress LER. The LSP name string must not contain "::" (two :s) nor contain a ":" (single ":") at the end of the LSP name. However, a ":" (single ":") can appear anywhere in the string except at the end of the name.

Parameters

rsvp-p2mp *lsp-name*

Specifies the LSP. The LSP name can be up to 32 characters long and must be unique.

ldp-p2mp *p2mp-id*

Identifier used for signaling MLDP P2MP LSP.

Values 1 to 4294967296 (on leaf node)
1 to 8192 (on root node)

sender *sender-address*

Specifies the sender IP address: a.b.c.d.

Platforms

7705 SAR Gen 2

tunnel-interface

Syntax

tunnel-interface [**rsvp-p2mp** *lsp-name*] [**sender** *ip-address*] [**detail**]

tunnel-interface [**ldp-p2mp** *p2mp-id*] [**sender** *ip-address*] [**detail**]

no tunnel-interface [**rsvp-p2mp** *lsp-name*] [**sender** *ip-address*]

no tunnel-interface [**ldp-p2mp** *p2mp-id*] [**sender** *ip-address*]

Context

[Tree] (debug>router>pim tunnel-interface)

Full Context

debug router pim tunnel-interface

Description

This command enables debugging for PIM tunnel interfaces.

The **no** form of this command disables debugging for PIM tunnel interfaces.

Parameters

lsp-name

Specifies the LSP for RSVP P2MP.

ip-address

Specifies the IP address of the sender.

p2mp-id

Specifies the P2MP ID for LDP P2MP.

detail

Displays detailed information for PIM tunnel interfaces.

Platforms

7705 SAR Gen 2

30.127 tunnel-member-pool

tunnel-member-pool

Syntax

tunnel-member-pool *name* [**create**]

no tunnel-member-pool *name*

Context

[Tree] (config>isa tunnel-member-pool)

Full Context

configure isa tunnel-member-pool

Description

Commands in this context configure associated ESA VM and MDAs.

The **no** form of this command removes the pool name from the configuration.

Parameters

name

Specifies the tunnel member pool name of the command, up to 32 characters.

create

Keyword used to create the command instance.

Platforms

7705 SAR Gen 2

30.128 tunnel-mtu

tunnel-mtu

Syntax

tunnel-mtu *bytes*

no tunnel-mtu

Context

[Tree] (config>router>isis>segment-routing tunnel-mtu)

Full Context

configure router isis segment-routing tunnel-mtu

Description

This command configures the MTU of all SR tunnels within each IGP instance.

The MTU of a SR tunnel populated into TTM is determined like in the case of an IGP tunnel; for example, LDP LSP, based on the outgoing interface MTU minus the label stack size. Remote LFA can add, at most, one more label to the tunnel for a total of two labels. There is no default value for this command. If the user does not configure an SR tunnel MTU, the MTU is determined by IGP as explained below.

The MTU of the SR tunnel in bytes is then determined as follows:

*SR_Tunnel_MTU = MIN {Cfg_SR_MTU, IGP_Tunnel_MTU- (1+frr-overhead)*4}*

Where:

Cfg_SR_MTU is the MTU configured by the user for all SR tunnels within a given IGP instance using the above CLI. If no value was configured by the user, the SR tunnel MTU will be determined by the IGP interface calculation explained next.

IGP_Tunnel_MTU is the minimum of the IS-IS or OSPF interface MTU among all the ECMP paths or among the primary and LFA backup paths of this SR tunnel.

frr-overhead is set to 1 if **segment-routing** and **remote-lfa** options are enabled in the IGMP instance. Otherwise, it is set to 0.

The SR tunnel MTU is dynamically updated anytime any of the above parameters used in its calculation changes. This includes when the set of the tunnel next-hops changes or the user changes the configured SR MTU or interface MTU value.

Default

no tunnel-mtu

Parameters

bytes

Specifies the size of the Maximum Transmission Unit (MTU) in bytes.

Values 512 to 9786

Platforms

7705 SAR Gen 2

tunnel-mtu

Syntax

tunnel-mtu *bytes*

no tunnel-mtu

Context

[Tree] (config>router>ospf>segm-rtng tunnel-mtu)

Full Context

configure router ospf segment-routing tunnel-mtu

Description

This command configures the MTU of all SR tunnels within each IGP instance.

The MTU of a SR tunnel populated into the TTM is determined as the same as an IGP tunnel; for example, for an LDP LSP, based on the outgoing interface MTU minus the label stack size. Remote LFA can add, at most, one more label to the tunnel for a total of two labels. There is no default value for this command. If the user does not configure an SR tunnel MTU, the MTU will be determined by IGP as follows:

The MTU of the SR tunnel in bytes is then determined as follows:

SR_Tunnel_MTU = MIN {Cfg_SR_MTU, IGP_Tunnel_MTU- (1+frr—overhead)X4}

Where:

- *Cfg_SR_MTU* is the MTU configured by the user for all SR tunnels within an IGP instance using the tunnel-mtu command. If no value is configured by the user, the SR tunnel MTU is determined by the IGP interface calculation explained in the next bullet point.
- *IGP_Tunnel_MTU* is the minimum of the IS-IS or OSPF interface MTU among all the ECMP paths or among the primary and LFA backup paths of this SR tunnel.
- *frr-overhead* is set to 1 if the **segment-routing** and **remote-lfa** options are enabled in the IGMP instance. Otherwise, it is set to 0.

The SR tunnel MTU is dynamically updated whenever any of the above parameters used in its calculation changes. This includes if the set of the tunnel next-hops changes or the user changes the configured SR MTU or interface MTU value.

Default

no tunnel-mtu

Parameters

bytes

Specifies the size of the MTU in bytes.

Values 512 to 9786

Platforms

7705 SAR Gen 2

30.129 tunnel-nearest-bridge

tunnel-nearest-bridge

Syntax

[**no**] **tunnel-nearest-bridge**

Context

[Tree] (cfg>port>eth>lldp>dstmac tunnel-nearest-bridge)

Full Context

configure port ethernet lldp dest-mac tunnel-nearest-bridge

Description

This command allows LLDP packets received on the port with the destination address of the nearest bridge to be tunneled without being intercepted on the local port. The dest-mac nearest-bridge must be disabled for tunneling to occur. This is applicable to NULL SAP Epipe and VPLS services only.

Default

no tunnel-nearest-bridge

Platforms

7705 SAR Gen 2

tunnel-nearest-bridge

Syntax

[**no**] **tunnel-nearest-bridge**

Context

[Tree] (config>lag>lldp-member-template>dstmac tunnel-nearest-bridge)

Full Context

configure lag lldp-member-template dest-mac tunnel-nearest-bridge

Description

This command allows LLDP packets received on the port with the destination address of the nearest bridge to be tunneled without being intercepted on the local port. The **dest-mac nearest-bridge** must be disabled for tunneling to occur. This is applicable to NULL SAP Epipe and VPLS services only.

The **no** form of this command disables the packets from being tunneled.

Default

no tunnel-nearest-bridge

Platforms

7705 SAR Gen 2

30.130 tunnel-next-hop

tunnel-next-hop

Syntax

tunnel-next-hop

Context

[Tree] (config>router>static-route-entry>indirect tunnel-next-hop)

Full Context

configure router static-route-entry indirect tunnel-next-hop

Description

Commands in this context configure the static route's nexthop to be resolved to an indirect tunnel next-hop.

Platforms

7705 SAR Gen 2

tunnel-next-hop

Syntax

tunnel-next-hop

Context

[Tree] (config>router>isis>igp-shortcut tunnel-next-hop)

Full Context

configure router isis igp-shortcut tunnel-next-hop

Description

Commands in this context configure the resolution of IGP IPv4 prefix families, IGP IPv6 prefix families, SR-ISIS IPv4 tunnel families, SR-ISIS IPv6 tunnel families, and SR-OSPF IPv4 tunnel families using IGP shortcuts.

The **resolution** node is introduced to provide flexibility in the selection of the tunnel types for each of the IP prefix and SR tunnel families.

The IPv4 **family** option causes the IS-IS or OSPF SPF to include the IPv4 IGP shortcuts in the IP reach calculation of IPv4 nodes and prefixes. RSVP-TE or SR-TE LSPs terminating on a node identified by its router ID can be used to reach IPv4 prefixes owned by this node or for which this node is the IPv4 next hop.

The IPv6 **family** option causes the IS-IS or OSPFv3 SPF to include the IPv4 IGP shortcuts in the IP reach calculation of IPv6 nodes and prefixes. RSVP-TE or SR-TE LSPs terminating on a node identified by its router ID can be used to reach IPv6 prefixes owned by this node or for which this node is the IPv6 next-hop. The resolution of IPv6 prefixes is supported in OSPFv3 and in both IS-IS MT=0 and MT=2.

The IS-IS and OSPFv3 IPv6 routes resolved to IPv4 IGP shortcuts are used to:

- forward packets of IS-IS or OSPFv3 prefixes matching these routes
- forward CPM-originated IPv6 packets
- resolve the BGP next hop of BGP IPv6 prefixes
- resolve the indirect next hop of static IPv6 routes

In the data path, a packet for an IPv6 prefix has a label stack that consists of the IPv6 Explicit-Null label value of 2 at the bottom of the label stack followed by the label stack of the IPv4 RSVP-TE LSP.

There is no default behavior for IPv4 prefixes to automatically resolve to RSVP-TE or SR-TE LSPs used as IGP shortcuts by only enabling the **igp-shortcut** context. Instead, the user must enable the **ipv4 family** or **ipv6 family** and set the resolution to the value of **rsvp-te** to select the RSVP-TE tunnel type, or to the value of **sr-te** to select the SR-TE tunnel type.

Setting the **resolution** to the **any** value means that IGP selects the tunnels used as IGP shortcuts according to the TTM preference for the tunnel type. The RSVP-TE LSP type is of higher priority than the SR-TE LSP type.

An IP prefix of family=ipv4 or family=ipv6 always resolves to a single type of tunnel **rsvp-te** or **sr-te**. **Rsvp-te** type is preferred if both types are allowed by the prefix family resolution and both types exist in the set of tunnel next-hops of the prefix. The feature does not support mixing tunnel types per prefix.

If **resolution** for the IPv4 or IPv6 family is set to **disabled**, the corresponding prefixes are resolved to IP next-hops in the multicast routing table.

The **srv4 family** enables the resolution of SR-OSPF IPv4 tunnels and SR-ISIS IPv4 tunnels in MT=0 over RSVP-TE IPv4 IGP shortcuts. A maximum of 32 ECMP tunnel next-hops can be programmed for an SR-OSPF or an SR-ISIS IPv4 tunnel.

The **srv6 family** enables the resolution of SR-ISIS IPv6 tunnels in MT=0 over RSVP-TE IPv4 IGP shortcuts. A maximum of 32 ECMP tunnel next-hops can be programmed for an SR-ISIS IPv6 tunnel.

One or more RSVP-TE LSPs can be selected if **resolution**=**match-family-ip** and the corresponding IPv4 or IPv6 prefix resolves to RSVP-TE LSPs.

Note:
An SR tunnel cannot resolve to SR-TE IGP shortcuts.

If **resolution** for the SRv4 or SRv6 tunnel family is set to **disabled**, the corresponding tunnels are resolved to IP next-hops in the multicast routing table.

To enable (disable) IGP shortcuts in the IGP instance, the user must perform a **shutdown** or **no shutdown** in the **igp-shortcut** context.

Platforms

7705 SAR Gen 2

tunnel-next-hop

Syntax

tunnel-next-hop

Context

[Tree] (config>router>ospf3>igp-shortcut tunnel-next-hop)

[Tree] (config>router>ospf>igp-shortcut tunnel-next-hop)

Full Context

configure router ospf3 igp-shortcut tunnel-next-hop

configure router ospf igp-shortcut tunnel-next-hop

Description

Commands in this context configure the resolution of IGP IPv4 prefix families, IGP IPv6 prefix families, SR-ISIS IPv4 tunnel families, SR-ISIS IPv6 tunnel families, and SR-OSPF IPv4 tunnel families using IGP shortcuts.

The **resolution** node is introduced to provide flexibility in the selection of the tunnel types for each of the IP prefix and SR tunnel families.

The IPv4 **family** option causes the IS-IS or OSPF SPF to include the IPv4 IGP shortcuts in the IP reach calculation of IPv4 nodes and prefixes. RSVP-TE or SR-TE LSPs terminating on a node identified by its router ID can be used to reach IPv4 prefixes owned by this node or for which this node is the IPv4 next hop.

The IPv6 **family** option causes the IS-IS or OSPFv3 SPF to include the IPv4 IGP shortcuts in the IP reach calculation of IPv6 nodes and prefixes. RSVP-TE or SR-TE LSPs terminating on a node identified by its router ID can be used to reach IPv6 prefixes owned by this node or for which this node is the IPv6 next hop. The resolution of IPv6 prefixes is supported in OSPFv3 and in both IS-IS MT=0 and MT=2.

The IS-IS and OSPFv3 IPv6 routes resolved to IPv4 IGP shortcuts are used to:

- forward packets of IS-IS or OSPFv3 prefixes matching these routes
- forward CPM-originated IPv6 packets
- resolve the BGP next hop of BGP IPv6 prefixes
- resolve the indirect next hop of static IPv6 routes

In the data path, a packet for an IPv6 prefix has a label stack that consists of the IPv6 Explicit-Null label value of 2 at the bottom of the label stack followed by the label stack of the IPv4 RSVP-TE LSP.

There is no default behavior for IPv4 prefixes to automatically resolve to RSVP-TE or SR-TE LSPs used as IGP shortcuts by only enabling the **igp-shortcut** context. Instead, the user must enable the **ipv4 family** or **ipv6 family** and set the resolution to the value of **rsvp-te** to select the RSVP-TE tunnel type, or to the value of **sr-te** to select the SR-TE tunnel type.

Setting the **resolution** to the **any** value means that IGP selects the tunnels used as IGP shortcuts according to the TTM preference for the tunnel type. The RSVP-TE LSP type is of higher priority than the SR-TE LSP type.

An IP prefix of family=ipv4 or family=ipv6 always resolves to a single type of tunnel **rsvp-te** or **sr-te**. **Rsvp-te** type is preferred if both types are allowed by the prefix family resolution and both types exist in the set of tunnel next-hops of the prefix. The feature does not support mixing tunnel types per prefix.

If **resolution** for the IPv4 or IPv6 family is set to **disabled**, the corresponding prefixes are resolved to IP next-hops in the multicast routing table.

The **srv4 family** enables the resolution of SR-OSPF IPv4 tunnels and SR-ISIS IPv4 tunnels in MT=0 over RSVP-TE IPv4 IGP shortcuts. A maximum of 32 ECMP tunnel next-hops can be programmed for an SR-OSPF or an SR-ISIS IPv4 tunnel.

The **srv6 family** enables the resolution of SR-ISIS IPv6 tunnels in MT=0 over RSVP-TE IPv4 IGP shortcuts. A maximum of 32 ECMP tunnel next-hops can be programmed for an SR-ISIS IPv6 tunnel.

One or more RSVP-TE LSPs can be selected if **resolution**=**match-family-ip** and the corresponding IPv4 or IPv6 prefix resolves to RSVP-TE LSPs.

Note:
An SR tunnel cannot resolve to SR-TE IGP shortcuts.

If **resolution** for the SRv4 or SRv6 tunnel family is set to **disabled**, the corresponding tunnels are resolved to IP next-hops in the multicast routing table.

To enable or disable IGP shortcuts in the IGP instance, the user must perform a **shutdown** or **no shutdown** in the **igp-shortcut** context.

Platforms

7705 SAR Gen 2

30.131 tunnel-qinq

tunnel-qinq

Syntax

[**no**] **tunnel-qinq**

Context

[Tree] (config>port>ethernet>dot1x tunnel-qinq)

Full Context

configure port ethernet dot1x tunnel-qinq

Description

This command configures the tunneling of double tagged (QinQ) dot1x packets. When enabled, the router extracts the packets to the CPM.

The **no** form of this command disables the tunelling of the QinQ dot1x packets on the port.

Default

tunnel-qinq

Platforms

7705 SAR Gen 2

30.132 tunnel-table

tunnel-table

Syntax

tunnel-table [*ip-address*] [{**ldp** | **rsvp** [**tunnel-id** *tunnel-id*] | **sdp** [**sdp-id** *sdp-id*}]]

Context

[Tree] (debug>router>ip tunnel-table)

Full Context

debug router ip tunnel-table

Description

This command enables debugging for tunnel tables.

Platforms

7705 SAR Gen 2

30.133 tunnel-table-pref

tunnel-table-pref

Syntax

tunnel-table-pref *preference*

no tunnel-table-pref

Context

[Tree] (config>router>ldp tunnel-table-pref)

Full Context

configure router ldp tunnel-table-pref

Description

This command configures the tunnel table preference for LDP tunnel type away from its default value.

The tunnel table preference applies to the next-hop resolution of BGP routes of the following families: EVPN, IPv4, IPv6, VPN-IPv4, VPN-IPv6, label-IPv4, and label-IPv6 in the tunnel table.

This feature does not apply to a VPRN, VPLS, or VLL service with explicit binding to an SDP that enabled the **mixed-lsp-mode** option. The tunnel preference in such an SDP is fixed and is controlled by the service manager. The configuration of the tunnel table preference parameter does not modify the behavior of such an SDP and the services that bind to it.

It is recommended to not set two or more tunnel types to the same preference value. In such a situation, the tunnel table prefers the tunnel type which was first introduced in SR OS implementation historically.

The **no** form of this command reverts to the default value.

Default

tunnel-table-pref 9

Parameters

preference

Specifies the preference value.

Values 1 to 255

Default 9

Platforms

7705 SAR Gen 2

tunnel-table-pref

Syntax

tunnel-table-pref

Context

[Tree] (config>router>mpls tunnel-table-pref)

Full Context

configure router mpls tunnel-table-pref

Description

Commands in this context configure the tunnel table preference for RSVP-TE LSP and SR-TE LSP tunnel types.

Platforms

7705 SAR Gen 2

tunnel-table-pref

Syntax

tunnel-table-pref *preference*

no tunnel-table-pref

Context

[Tree] (config>router>isis>segment-routing tunnel-table-pref)

Full Context

configure router isis segment-routing tunnel-table-pref

Description

This command configures the TTM preference of SR tunnels created by the IGP instance. This is used in the case of BGP shortcuts, VPRN auto-bind, or BGP transport tunnel when the new tunnel binding commands are configured to the **any** value which parses the TTM for tunnels in the protocol preference order. The user can choose to either go with the global TTM preference or list explicitly the tunnel types they want to use. When they list the tunnel types explicitly, the TTM preference will still be used to select one type over the other. In both cases, a fallback to the next preferred tunnel type is performed if the selected one fails. Also, a reversion to a more preferred tunnel type is performed as soon as one is available.

The segment routing module adds to TTM a SR tunnel entry for each resolved remote node SID prefix and programs the data path with the corresponding LTN with the push operation pointing to the primary and LFA backup NHLFEs.

The default preference for SR tunnels in the TTM is set lower than LDP tunnels but higher than BGP tunnels to allow controlled migration of customers without disrupting their current deployment when they enable segment routing. The following is the setting of the default preference of the various tunnel types. This includes the preference of SR tunnels based on shortest path (referred to as **SR-ISIS** and **SR-OSPF**).

The global default TTM preference for the tunnel types is as follows:

- ROUTE_PREF_RSVP 7
- ROUTE_PREF_SR_TE 8
- ROUTE_PREF_LDP 9
- ROUTE_PREF_OSPF_TTM 10
- ROUTE_PREF_ISIS_TTM 11
- ROUTE_PREF_BGP_TTM 12
- ROUTE_PREF_GRE 255

The default value for SR-ISIS or SR-OSPF is the same regardless if one or more IS-IS or OSPF instances programmed a tunnel for the same prefix. The selection of a SR tunnel in this case will be based on lowest IGP instance-id.

It is recommended to not set two or more tunnel types to the same preference value. In such a situation, the tunnel table prefers the tunnel type which was first introduced in SR OS implementation historically.

Default

tunnel-table-pref 11

Parameters

preference

Specifies the integer value to represent the preference of IS-IS or OSPF SR tunnels in TTM.

Values 1 to 255

Platforms

7705 SAR Gen 2

tunnel-table-pref

Syntax

tunnel-table-pref *preference*

no tunnel-table-pref

Context

[Tree] (config>router>ospf>segm-rtng tunnel-table-pref)

Full Context

configure router ospf segment-routing tunnel-table-pref

Description

This command configures the TTM preference of shortest path SR tunnels created by the IGP instance. This is used for BGP shortcuts, VPRN auto-bind, or BGP transport tunnel when the tunnel binding commands are configured to the **any** value, which parses the TTM for tunnels in the protocol preference order. The user can choose to either accept the global TTM preference or explicitly list the tunnel types they want to use. If the user lists the tunnel types explicitly, the TTM preference is still used to select one type over the other. In both cases, a fallback to the next preferred tunnel type is performed if the selected type fails. A reversion to a more preferred tunnel type is performed as soon as one is available.

The segment routing module adds to the TTM an SR tunnel entry for each resolved remote node SID prefix and programs the data path having the corresponding LTN with the push operation pointing to the primary and LFA backup NHLFEs.

The default preference for shortest path SR tunnels in the TTM is set lower than LDP tunnels but higher than BGP tunnels to allow controlled migration of customers without disrupting their current deployment when they enable segment routing. The following is the value of the default preference for the various tunnel types. This includes the preference of SR tunnels based on shortest path (referred to as SR-ISIS and SR-OSPF).

Note:
The preference of an SR-TE LSP is not configurable and is the second most preferred tunnel type after RSVP-TE. The preference is the same whether if the SR-TE LSP was resolved in IS-IS or OSPF.

The global default TTM preference for the tunnel types is as follows:

- ROUTE_PREF_RSVP 7
- ROUTE_PREF_SR_TE 8
- ROUTE_PREF_LDP 9
- ROUTE_PREF_OSPF_TTM 10
- ROUTE_PREF_ISIS_TTM 11
- ROUTE_PREF_BGP_TTM 12
- ROUTE_PREF_GRE 255

The default value for SR-ISIS or SR-OSPF is the same regardless if one or more IS-IS or OSPF instances programmed a tunnel for the same prefix. The selection of a SR tunnel in this case will be based on the lowest IGP instance ID. Similarly, IPv6 SR-ISIS and SR-OSPF3 tunnels are programmed into TTMv6 with the same default preference value as IPv4 SR-ISIS and IPv4 SR-OSPF respectively.

It is recommended to not set two or more tunnel types to the same preference value. In such a situation, the tunnel table prefers the tunnel type which was first introduced in SR OS implementation historically.

Default

tunnel-table-pref 10

Parameters

preference

Specifies the integer value to represent the preference of IS-IS, OSPF, or OSPF3 SR tunnels in the TTM.

Values 1 to 255

Platforms

7705 SAR Gen 2

30.134 tunnel-template

tunnel-template

Syntax

tunnel-template *tunnel-template-id*

no tunnel-template

Context

[Tree] (config>ipsec>client-db>client tunnel-template)

Full Context

configure ipsec client-db client tunnel-template

Description

This command specifies the tunnel template to be used for tunnel setup.

The **no** form of this command reverts to the default.

Default

no tunnel-template

Parameters

tunnel-template-id

Specifies the identifier of the tunnel template.

Values 1 to 2048

Platforms

7705 SAR Gen 2

tunnel-template

Syntax

tunnel-template *ipsec-template-identifier* [**create**]

no tunnel-template *ipsec-template-identifier*

Context

[Tree] (config>ipsec tunnel-template)

Full Context

configure ipsec tunnel-template

Description

This command creates a tunnel template. Up to 2000 templates are allowed.

Parameters

ipsec-template-identifier

Specifies the template identifier.

Values 1 to 2048

create

Mandatory keyword used when creating a tunnel-template in the IPsec context. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

30.135 tunneling

tunneling

Syntax

[**no**] **tunneling**

Context

[Tree] (config>port>ethernet>dot1x tunneling)

Full Context

configure port ethernet dot1x tunneling

Description

This command enables the tunneling of untagged 802.1x frames received on a port and is supported only when **dot1x port-control** is set to **force-auth**. 802.1x tunneling is applicable to both Epipe and VPLS services using either a null SAP or a default SAP on a dot1q port. When configured, untagged 802.1x frames will be switched into the service with the corresponding supported SAP.

The **no** form of this command disables tunneling of untagged 802.1x frames.

Default

no tunneling

Platforms

7705 SAR Gen 2

tunneling

Syntax

[**no**] **tunneling**

Context

[Tree] (config>router>ldp>targ-session>peer-template tunneling)

[Tree] (config>router>ldp>targ-session>peer tunneling)

Full Context

configure router ldp targeted-session peer-template tunneling

configure router ldp targeted-session peer tunneling

Description

This command enables LDP over tunnels.

The **no** form of this command disables tunneling.

Default

no tunneling

Platforms

7705 SAR Gen 2

tunneling

Syntax

[**no**] **tunneling**

Context

[Tree] (config>router>ldp>targ-session>auto-rx>ipv4 tunneling)

[Tree] (config>router>ldp>targ-session>auto-tx>ipv4 tunneling)

Full Context

configure router ldp targeted-session auto-rx ipv4 tunneling

configure router ldp targeted-session auto-tx ipv4 tunneling

Description

This command enables the local system to use the targeted LDP session to send FEC/label bindings that it has advertised to other LDP peers. For LDP rLFA, the source node requires the PQ node's label binding information in order to reach the destination. Therefore, this command must be enabled for the **auto-rx** context. However, because **auto-rx** has lower precedence, **tunneling** must be enabled under the **auto-tx** command, in case **auto-rx** is in a **no shutdown** state on the same system.

The **no** form of this command disables the local system from sending FEC/label bindings.

Default

no tunneling

Platforms

7705 SAR Gen 2

30.136 twamp

twamp

Syntax

twamp

Context

[Tree] (config>test-oam twamp)

Full Context

configure test-oam twamp

Description

This command enables TWAMP functionality.

Platforms

7705 SAR Gen 2

30.137 twamp-light

twamp-light

Syntax

twamp-light [**test-id** *test-id*] [**create**]

no twamp-light

Context

[Tree] (config>oam-pm>session>ip twamp-light)

Full Context

configure oam-pm session ip twamp-light

Description

This command assigns an identifier to the TWAMP Light test and creates the individual test.

The **no** form of this command removes the TWAMP Light test function from the OAM-PM session.

Parameters

test-id

Specifies the value of the 4-byte local test identifier not sent in the TWAMP Light packets.

Values 0 to 2147483647 | **auto**

auto - automatically assigns a *test-id*

create

Creates the test.

Platforms

7705 SAR Gen 2

twamp-light

Syntax

twamp-light

Context

[Tree] (config>router twamp-light)

[Tree] (config>test-oam>twamp twamp-light)

[Tree] (config>service>vprn twamp-light)

Full Context

configure router twamp-light

configure test-oam twamp twamp-light

configure service vprn twamp-light

Description

Commands in this context configure TWAMP Light parameters.

Platforms

7705 SAR Gen 2

30.138 tx-credit-max

tx-credit-max

Syntax

tx-credit-max *count*

no tx-credit-max

Context

[Tree] (config>system>lldp tx-credit-max)

Full Context

configure system lldp tx-credit-max

Description

This command configures the maximum consecutive LLDPDUs transmitted.

The **no** form of this command reverts to the default value.

Default

no tx-credit-max

Parameters

count

Specifies the maximum consecutive LLDPDUs transmitted.

Values 1 to 100

Default 5

Platforms

7705 SAR Gen 2

30.139 tx-dus

tx-dus

Syntax

[**no**] **tx-dus**

Context

[Tree] (config>port>ethernet>ssm tx-dus)

Full Context

configure port ethernet ssm tx-dus

Description

This command forces the QL value transmitted from the SSM channel of the SONET/SDH port or the Synchronous Ethernet port to be set to QL-DUS/QL-DNU. This capability is provided to block the use of the interface from the 7705 SAR Gen 2 for timing purposes.

This command is supported on TDM satellite.

Default

no tx-dus

Platforms

7705 SAR Gen 2

30.140 tx-hold-multiplier

tx-hold-multiplier

Syntax

tx-hold-multiplier *multiplier*

no tx-hold-multiplier

Context

[Tree] (config>system>lldp tx-hold-multiplier)

Full Context

configure system lldp tx-hold-multiplier

Description

This command configures the multiplier of the tx-interval.

The **no** form of this command reverts to the default value.

Default

no tx-hold-multiplier

Parameters

multiplier

Specifies the multiplier of the tx-interval.

Values 2 to 10

Default 4

Platforms

7705 SAR Gen 2

30.141 tx-interval

tx-interval

Syntax

tx-interval *interval*

no tx-interval

Context

[Tree] (config>system>lldp tx-interval)

Full Context

configure system lldp tx-interval

Description

This command configures the LLDP transmit interval time.

The **no** form of this command reverts to the default value.

Default

no tx-interval

Parameters

interval

Specifies the LLDP transmit interval time.

Values 5 to 32768

Default 30

Platforms

7705 SAR Gen 2

30.142 tx-mgmt-address

tx-mgmt-address

Syntax

tx-mgmt-address [**system**] [**system-ipv6**] [**oob**] [**oob-ipv6**]

no tx-mgmt-address

Context

[Tree] (config>port>ethernet>lldp>dstmac tx-mgmt-address)

Full Context

configure port ethernet lldp dest-mac tx-mgmt-address

Description

This command specifies which management address to transmit. The operator can choose to send the system IPv4 address, the system IPv6 address, the out-of-band IPv4 address, the out-of-band IPv6 address, or any combination of these. The system address is sent only once. The address must be configured for the specific version of the protocol in order to send the management address.

The **no** form of the command resets value to the default.

Default

no tx-mgmt-address

Parameters

system

Specifies to use the system IP address. The system address will only be transmitted once it has been configured if this parameter is specified.

system-ipv6

Specifies to use the system IPv6 address. The system address will only be transmitted once it has been configured if this parameter is specified.

oob

Specifies to use the out-of-band IPv4 address for active CPM.

oob-ipv6

Specifies to use the out-of-band IPv6 address for active CPM.

Platforms

7705 SAR Gen 2

tx-mgmt-address

Syntax

tx-mgmt-address [**system**] [**system-ipv6**] [**oob**] [**oob-ipv6**]

no tx-mgmt-address

Context

[Tree] (config>lag>lldp-member-template>dstmac tx-mgmt-address)

Full Context

configure lag lldp-member-template dest-mac tx-mgmt-address

Description

This command configures the management address to transmit. The operator can choose to send the system IPv4 address, system IPv6 address, out-of-band IPv4 address, out-of-band IPv6 address, or any combination of these. The system address is sent only once. The address must be configured for the specific version of the protocol to send the management address.

The **no** form of this command reverts to the default value.

Default

no tx-mgmt-address

Parameters

system

Keyword to use the system IP address. The system address is only transmitted after it has been configured if this keyword is specified.

system-ipv6

Keyword to use the system IPv6 address. The system address must be configured before it can be transmitted using this keyword.

oob

Keyword to use the out-of-band IPv4 address for active CPM.

oob-ipv6

Keyword to use the out-of-band IPv6 address for active CPM.

Platforms

7705 SAR Gen 2

30.143 tx-tlvs

tx-tlvs

Syntax

tx-tlvs [**port-desc**] [**sys-name**] [**sys-desc**] [**sys-cap**]

no tx-tlvs

Context

[Tree] (config>port>ethernet>lldp>dstmac tx-tlvs)

Full Context

configure port ethernet lldp dest-mac tx-tlvs

Description

This command specifies which LLDP TLVs to transmit. The TX TLVs, defined as a bitmap, includes the basic set of LLDP TLVs whose transmission is allowed on the local LLDP agent by the network management. Each bit in the bitmap corresponds to a TLV type associated with a specific optional TLV. Organizationally-specific TLVs are excluded from this bitmap.

There is no bit reserved for the management address TLV type since transmission of management address TLVs are controlled by another object.

The **no** form of this command resets the value to the default.

Default

no tx-tlvs

Parameters

port-desc

Indicates that the LLDP agent should transmit port description TLVs.

sys-name

Indicates that the LLDP agent should transmit system name TLVs.

sys-desc

Indicates that the LLDP agent should transmit system description TLVs.

sys-cap

Indicates that the LLDP agent should transmit system capabilities TLVs.

Platforms

7705 SAR Gen 2

tx-tlvs

Syntax

tx-tlvs [**port-desc**] [**sys-name**] [**sys-desc**] [**sys-cap**]

no tx-tlvs

Context

[Tree] (config>lag>lldp-member-template>dstmac tx-tlvs)

Full Context

configure lag lldp-member-template dest-mac tx-tlvs

Description

This command configures which LLDP TLVs to transmit. The TX TLVs, defined as a bitmap, include the basic set of LLDP TLVs whose transmission is allowed on the local LLDP agent by the network management. Each bit in the bitmap corresponds to a TLV type associated with a specific optional TLV. Organizationally specific TLVs are excluded from this bitmap.

No bit is reserved for the management address TLV type because transmission of these TLVs is controlled by another object.

The **no** form of this command reverts to the default value.

Default

no tx-tlvs

Parameters

port-desc

Keyword to specify that the LLDP agent transmits port description TLVs.

sys-name

Keyword to specify that the LLDP agent transmits system name TLVs.

sys-desc

Keyword to specify that the LLDP agent transmits system description TLVs.

sys-cap

Keyword to specify that the LLDP agent transmits system capabilities TLVs.

Platforms

7705 SAR Gen 2

30.144 type

type

Syntax

type *reflector-type*

Context

[Tree] (config>router>twamp-light>reflector type)

[Tree] (config>service>vprn>twamp-light>refl type)

Full Context

configure router twamp-light reflector type

configure service vprn twamp-light reflector type

Description

This command configures the processing behavior of the TWAMP Light reflector. When the value is **twamp-light**, the reflector does not check the received PDU as a traditional base TWAMP Light packet without TLV processing. When the value is **stamp**, the reflector attempts to find and process supported STAMP TLVs that follow the base STAMP packet.

In mixed environments where different types of session senders may be targeting a common TWAMP Light reflector, the value should be set to stamp. When the reflector is operating in stamp mode, the primary parsing is based on STAMP, checking and processing known TLVs, and also determining when TLVs are not present and the arriving PDU is a TWAMP Light PDU. A session sender launching a TWAMP Light-based packet must use all zeros and a padding pattern zero when the pad size is non zero.

Default

type twamp-light

Parameters

reflector-type

Specifies the type of processing behavior for the reflector.

Values stamp, twamp-light

Platforms

7705 SAR Gen 2

type

Syntax

[**no**] **type** {**internal** | **external**}

Context

[Tree] (config>service>vprn>bgp>group>neighbor type)

[Tree] (config>service>vprn>bgp>group type)

Full Context

configure service vprn bgp group neighbor type

configure service vprn bgp group type

Description

This command designates the BGP peer as type internal or external.

The type of **internal** indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.

By default, the OS derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, then the peer is considered **external**.

The **no** form of this command used at the group level reverts to the default value.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no type

Parameters

internal

Configures the peer as internal.

external

Configures the peer as external.

no type

Type of neighbor is derived on the local AS specified.

Platforms

7705 SAR Gen 2

type

Syntax

type [**hub** | **spoke** | **subscriber-split-horizon**]

no type

Context

[Tree] (config>service>vprn type)

Full Context

configure service vprn type

Description

This command designates the type of VPRN instance being configured for hub and spoke topologies. Use the **no** form to reset to the default of a fully meshed VPRN.

Default

no type

Parameters

hub

Specifies a hub VPRN which allows all traffic from the hub SAPs to be routed to the destination directly, while all traffic from spoke VPRNs or network interfaces can only be routed to a hub SAP.

spoke

Specifies a spoke VPRN which allows traffic from associated SAPs or spoke terminations to only be forwarded through routes learned from separate VPRN, which should be configured as a type Hub VPRN.

subscriber-split-horizon

Controls the flow of traffic for wholesale subscriber applications.

Platforms

7705 SAR Gen 2

type

Syntax

[**no**] **type**

Context

[Tree] (config>saa>test type)

Full Context

configure saa test type

Description

This command creates the context to provide the test type for the named test. Only a single test type can be configured.

A test can only be modified while the test is in shut down mode.

Once a test type has been configured, the command can be modified by re-entering the command. However, the test type must be the same as the previously entered test type.

To change the test type, the old command must be removed using the **config>saa>test>no type** command.

The **no** form of this command removes the test type parameters from the configuration.

Platforms

7705 SAR Gen 2

type

Syntax

type *filter-type*

no type

Context

[Tree] (config>qos>sap-ingress>mac-criteria type)

Full Context

configure qos sap-ingress mac-criteria type

Description

This command sets the mac-criteria type.

Default

type normal

Parameters

filter-type

Specifies which type of entries this MAC filter can contain.

Values **normal** — Regular match criteria are allowed; ISID match not allowed.

vid — Configures the VID filter type used to match on ethernet_II frame types. This allows matching VLAN tags for explicit filtering.

Platforms

7705 SAR Gen 2

type

Syntax

type {**cpm-np**}

no type

Context

[Tree] (config>router>bfd>bfd-template type)

Full Context

configure router bfd bfd-template type

Description

This command selects the CPM network processor as the local termination point for the BFD session. This is enabled by default.

The **no** form of this command reverts to the default behavior.

Default

no type

Platforms

7705 SAR Gen 2

type

Syntax

type *file-url* [**no-redirect**] [**client-tls-profile** *profile*] [**proxy** *proxy-url*]

Context

[Tree] (file type)

Full Context

file type

Description

This command displays the contents of a text file.

Parameters

file-url

Specifies the file contents to display.

Values		
	local-url	[*cflash-id/*][*file-path*] up to 200 characters, including cflash-id directory length up to 99 each
	remote-url	[{ftp:// | tftp:// | http:// | https://}*login*:*pswd@remote-locn*/][*file-path*]
		up to 247 characters
		directory length up to 99 characters each
	remote-locn	[hostname | ipv4-address | [ipv6-address]]
	ipv4-address	*a.b.c.d*
	ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]
		x:x:x:x:x:x:d.d.d.d[-*interface*]
		x - [0 to FFFF]H
		d - [0 to 255]D
		interface - up to 32 characters, for link local addresses 255
	cflash-id	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

profile

Specifies the TLS client profile configured under **config**>**system**>**security**>**tls**> **client-tls-profile** to use.

proxy-url

Specifies the URL of an HTTP proxy. For example, http://proxy.mydomain.com:8000. This URL must be an HTTP URL and not an HTTPS URL.

no-redirect

Specifies to automatically refuse any HTTP redirects without prompting the user.

Platforms

7705 SAR Gen 2

type

Syntax

type *schedule-type*

Context

[Tree] (config>system>cron>sched type)

Full Context

configure system cron schedule type

Description

This command specifies how the system should interpret the commands contained within the schedule node.

Default

type periodic

Parameters

schedule-type

Specifies the type of schedule for the system to interpret the commands contained within the schedule node.

Values periodic — Specifies a schedule which runs at a given interval. The interval must be specified for this feature to run successfully.

calendar — Specifies a schedule which runs based on a calendar. The month, weekday, day-of-month, and minute parameters must be specified for this feature to run successfully.

oneshot — Specifies a schedule which runs one time only. As soon as the first event specified in these parameters takes place and the associated event occurs, the schedule enters a shutdown state. The month, weekday, day-of-month, and minute parameters must be specified for this feature to run successfully.

Default periodic

Platforms

7705 SAR Gen 2

type

Syntax

type *indicator-type*

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>progress-indicator type)

Full Context

configure system management-interface cli md-cli environment progress-indicator type

Description

This command specifies the type of progress indicator used in the MD-CLI.

Default

type dots

Parameters

indicator-type

Specifies the progress indicator type.

Values **dots**: displays the progress indicator as dynamically changing dots

Platforms

7705 SAR Gen 2

type

Syntax

type all

type [**gnmi-capabilities**] [**gnmi-get**] [**gnmi-set**] [**gnmi-subscribe**] [**gnoi-cert-mgmt-rpcs**]

no type

Context

[Tree] (debug>system>grpc type)

Full Context

debug system grpc type

Description

This command enables debugging for all RPCs or a particular RPC.

The **no** form of this command deactivates debugging for all RPCs.

Parameters

all

Specifies that debugging is enabled for all RPCs.

gnmi-capabilities

Specifies that debugging is enabled for gNMI capability RPC.

gnmi-get

Specifies that debugging is enabled for gNMI get RPC.

gnmi-set

Specifies that debugging is enabled for gNMI set RPC.

gnmi-subscribe

Specifies that debugging is enabled for gNMI subscribe RPC.

gnoi-cert-mgmt-rpcs

Specifies that debugging is enabled for gNOI certificate management RPCs.

Platforms

7705 SAR Gen 2

type

Syntax

[**no**] **type** {**internal** | **external**}

Context

[Tree] (config>router>bgp>group type)

[Tree] (config>router>bgp>group>neighbor type)

Full Context

configure router bgp group type

configure router bgp group neighbor type

Description

This command designates the BGP peer as type internal or external.

The type of **internal** indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.

By default, the router derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, then the peer is considered **external**.

The **no** form of this command used at the group level reverts to the default value.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no type

Parameters

internal

Configures the peer as internal.

external

Configures the peer as external.

Platforms

7705 SAR Gen 2

type

Syntax

type {**1** | **2**}

no type

Context

[Tree] (config>router>policy-options>policy-statement>entry>from type)

Full Context

configure router policy-options policy-statement entry from type

Description

This command configures an OSPF type metric as a match criterion in the route policy statement entry.

If no type is specified, any OSPF type is considered a match.

The **no** form of this command removes the OSPF type match criterion.

Default

no type

Parameters

1

Matches OSPF routes with type 1 LSAs.

2

Matches OSPF routes with type 2 LSAs.

Platforms

7705 SAR Gen 2

type

Syntax

type {*type* | *param-name*}

no type

Context

[Tree] (config>router>policy-options>policy-statement>default-action type)

[Tree] (config>router>policy-options>policy-statement>entry>action type)

Full Context

configure router policy-options policy-statement default-action type

configure router policy-options policy-statement entry action type

Description

This command sets the subtype for the Type 5 LSA (external LSA).

The **no** form of this command disables assigning a type in the route policy entry.

Default

type 2

Parameters

type

Specifies the type metric.

Values Subtype 1 — The external metric in the external LSA is comparable with the internal metric, and thus one can sum up all the metrics along the path (both internal and external) to get the total cost to the destination.

Subtype 2 — The metric in the external LSA is much more important than the internal metric, so the internal metrics should only be considered when comparing two external routes that have the same external metric.

param-name

The type parameter variable name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@”.

Platforms

7705 SAR Gen 2

30.145 type-multi-line

type-multi-line

Syntax

[**no**] **type-multi-line**

Context

[Tree] (config>saa>test type-multi-line)

Full Context

configure saa test type-multi-line

Description

This command creates the context to configure the OAM probe type and its parameters in a flexible multi-line format.

The **no** form of this command removes the context.

Platforms

7705 SAR Gen 2

31 u Commands

31.1 udp

udp

Syntax

udp [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no udp

Context

[Tree] (config>service>nat>nat-policy>timeouts udp)

Full Context

configure service nat nat-policy timeouts udp

Description

This command configures the UDP mapping timeout.

Default

udp min 5

Parameters

hours

Specifies the timeout hours field.

Values 1 to 24

minutes

Specifies the timeout minutes field.

Values 1 to 59

seconds

Specifies the timeout seconds field.

Values 1 to 59

Platforms

7705 SAR Gen 2

udp

Syntax

[**no**] **udp**

Context

[Tree] (config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter udp)

Full Context

configure router bgp next-hop-resolution labeled-routes transport-tunnel family resolution-filter udp

Description

This command selects UDP tunnel in TTM for next-hop resolution.

Platforms

7705 SAR Gen 2

31.2 udp-dns

udp-dns

Syntax

udp-dns [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no udp-dns

Context

[Tree] (config>service>nat>nat-policy>timeouts udp-dns)

Full Context

configure service nat nat-policy timeouts udp-dns

Description

This command configures the timeout applied to a UDP session with destination port 53.

Default

udp-dns sec 15

Parameters

hours

Specifies the timeout hours field.

Values 1 to 24

minutes

Specifies the timeout minutes field.

Values 1 to 59

seconds

Specifies the timeout seconds field.

Values 1 to 59

Platforms

7705 SAR Gen 2

31.3 udp-inbound-refresh

udp-inbound-refresh

Syntax

[**no**] **udp-inbound-refresh**

Context

[Tree] (config>service>nat>nat-policy udp-inbound-refresh)

Full Context

configure service nat nat-policy udp-inbound-refresh

Description

This command enables UDP session timeout extended on inbound traffic.

The **no** form of the command disables UDP session timeout extended on inbound traffic.

Default

no udp-inbound-refresh

Platforms

7705 SAR Gen 2

31.4 udp-initial

udp-initial

Syntax

udp-initial [**min** *minutes*] [**sec** *seconds*]

no udp-initial

Context

[Tree] (config>service>nat>nat-policy>timeouts udp-initial)

Full Context

configure service nat nat-policy timeouts udp-initial

Description

This command configures the UDP mapping timeout applied to new sessions.

Default

udp-initial sec 15

Parameters

minutes

Specifies the timeout minutes field.

Values 1 to 59

seconds

Specifies the timeout seconds field.

Values 1 to 59

Platforms

7705 SAR Gen 2

31.5 unavailability-event

unavailability-event

Syntax

unavailability-event {**forward** | **backward** | **aggregate**} **threshold** *raise-threshold* [**clear** *clear-threshold*]

no unavailability-event {**forward** | **backward** | **aggregate**}

Context

[Tree] (config>oam-pm>session>ip>twamp-light>loss-events unavailability-event)

Full Context

configure oam-pm session ip twamp-light loss-events unavailability-event

Description

This command sets the threshold to be applied to the overall count of the unavailability indicators, not transitions, per configured direction. This value is compared to the 32 bit unavailability counter specific to the direction which tracks the number of individual delta-ts that have been recorded as unavailable. The aggregate is a function of summing forward and backward. This value is only used as a threshold mechanism and is not part of the stored statistics. If the optional **clear** *clear-threshold* parameter is not specified, the traffic crossing alarm is stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and regardless of any previous window. Each unique event can only be raised once within measurement interval. If the optional **clear** *clear-threshold* parameter is specified, the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to following measurement intervals. If a threshold crossing event is raised another is raised until a measurement interval completes and the clear threshold has not been exceeded. A clear event is raised under that condition.

The **no** form of this command removes the event threshold for frame loss ratio. The direction must be included with the **no** command.

Default

no unavailability-event forward

no unavailability-event backward

no unavailability-event aggregate

Parameters

forward

Specifies the threshold is applied to the forward direction count.

backward

Specifies the threshold is applied to the backward direction count.

aggregate

Specifies the threshold is applied to the aggregate count (sum of forward and backward).

raise-threshold

Specifies a numerical value compared to the unavailability counter that is the rising threshold that determines when the event is to be generated, when value reached.

Values 1 to 864000

clear-threshold

Specifies an optional value used for stateful behavior that allows the operator to configure a percentage of loss value lower than the rising percentage to indicate when the clear event should be generated.

Values 0 to 863999

A value of zero means that the unavailability counter must be 0.

Platforms

7705 SAR Gen 2

31.6 uncoloured-octets-offered-count

uncoloured-octets-offered-count

Syntax

[**no**] **uncoloured-octets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters uncoloured-octets-offered-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters uncoloured-octets-offered-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters uncoloured-octets-offered-count

configure log accounting-policy custom-record policer e-counters uncoloured-octets-offered-count

Description

This command includes the uncoloured octets offered count.

The **no** form of this command excludes the uncoloured octets offered count.

Default

no uncoloured-octets-offered-count

Platforms

7705 SAR Gen 2

uncoloured-octets-offered-count

Syntax

[**no**] **uncoloured-packets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters uncoloured-octets-offered-count)

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters uncoloured-octets-offered-count)

[Tree] (config>log>acct-policy>cr>queue>i-counters uncoloured-octets-offered-count)

[Tree] (config>log>acct-policy>cr>policer>i-counters uncoloured-octets-offered-count)

Full Context

configure log accounting-policy custom-record ref-policer i-counters uncoloured-octets-offered-count

configure log accounting-policy custom-record ref-queue i-counters uncoloured-octets-offered-count

configure log accounting-policy custom-record queue i-counters uncoloured-octets-offered-count

configure log accounting-policy custom-record policer i-counters uncoloured-octets-offered-count

Description

This command includes the uncoloured octets offered in the count.

The **no** form of this command excludes the uncoloured octets offered in the count.

Default

no uncoloured-octets-offered-count

Platforms

7705 SAR Gen 2

31.7 uncoloured-packets-offered-count

uncoloured-packets-offered-count

Syntax

[**no**] **uncoloured-packets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>e-counters uncoloured-packets-offered-count)

[Tree] (config>log>acct-policy>cr>policer>e-counters uncoloured-packets-offered-count)

Full Context

configure log accounting-policy custom-record ref-policer e-counters uncoloured-packets-offered-count

configure log accounting-policy custom-record policer e-counters uncoloured-packets-offered-count

Description

This command includes the uncoloured packets offered count.

The **no** form of this command excludes the uncoloured packets offered count.

Default

no uncoloured-packets-offered-count

Platforms

7705 SAR Gen 2

uncoloured-packets-offered-count

Syntax

[**no**] **uncoloured-packets-offered-count**

Context

[Tree] (config>log>acct-policy>cr>ref-policer>i-counters uncoloured-packets-offered-count)

[Tree] (config>log>acct-policy>cr>queue>i-counters uncoloured-packets-offered-count)

[Tree] (config>log>acct-policy>cr>ref-queue>i-counters uncoloured-packets-offered-count)

[Tree] (config>log>acct-policy>cr>policer>i-counters uncoloured-packets-offered-count)

Full Context

configure log accounting-policy custom-record ref-policer i-counters uncoloured-packets-offered-count

configure log accounting-policy custom-record queue i-counters uncoloured-packets-offered-count

configure log accounting-policy custom-record ref-queue i-counters uncoloured-packets-offered-count

configure log accounting-policy custom-record policer i-counters uncoloured-packets-offered-count

Description

This command includes the uncolored packets offered count.

The **no** form of this command excludes the uncoloured packets offered count.

Default

no uncoloured-packets-offered-count

Platforms

7705 SAR Gen 2

31.8 uncommitted-changes-indicator

uncommitted-changes-indicator

Syntax

[**no**] **uncommitted-changes-indicator**

Context

[Tree] (config>system>management-interface>cli>md-cli>environment>prompt uncommitted-changes-indicator)

Full Context

configure system management-interface cli md-cli environment prompt uncommitted-changes-indicator

Description

This command displays the change indicator.

The **no** form of this command suppresses the change indicator.

Default

uncommitted-changes-indicator

Platforms

7705 SAR Gen 2

31.9 undet-availability-event

undet-availability-event

Syntax

undet-availability-event {**forward** | **backward** | **aggregate**} **threshold** *raise-threshold* [**clear** *clear-threshold*]

no undet-availability-event {**forward** | **backward** | **aggregate**}

Context

[Tree] (config>oam-pm>session>ip>twamp-light>loss-events undet-availability-event)

Full Context

configure oam-pm session ip twamp-light loss-events undet-availability-event

Description

This command sets the threshold to be applied to the overall count of the undetermined availability indicators, not transitions, per configured direction. This value is compared to the 32 bit unavailability counter specific to the direction which tracks the number of individual delta-ts that have been recorded as undetermined available. The aggregate is a function of summing forward and backward. This value is only used as a threshold mechanism and is not part of the stored statistics. If the optional **clear** *clear-threshold* parameter is not specified, the traffic crossing alarm is stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and regardless of any previous window. Each unique event can only be raised once within measurement interval. If the optional **clear** *clear-threshold* parameter is specified, the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to following measurement intervals. If a threshold crossing event is raised another is raised until a measurement interval completes and the clear threshold has not been exceeded. A clear event is raised under that condition.

The **no** form of this command removes the event threshold for frame loss ratio. The direction must be included with the **no** command.

Default

no undet-availability-event forward

no undet-availability-event backward

no undet-availability-event aggregate

Parameters

forward

Specifies the threshold is applied to the forward direction count.

backward

Specifies the threshold is applied to the backward direction count.

aggregate

Specifies the threshold is applied to the aggregate count (sum of forward and backward).

raise-threshold

Specifies the rising threshold that determines when the event is to be generated, when value reached.

Values 1 to 864000

clear-threshold

Specifies an optional value used for stateful behavior that allows the operator to configure a percentage of loss value lower than the rising percentage to indicate when the clear event should be generated.

Values 0 to 863999

A value of zero means that the undetermined availability counter must be 0.

Platforms

7705 SAR Gen 2

31.10 undet-unavailability-event

undet-unavailability-event

Syntax

undet-unavailability-event {**forward** | **backward** | **aggregate**} **threshold** *raise-threshold* [**clear** *clear-threshold*]

no undet-unavailability-event {**forward** | **backward** | **aggregate**}

Context

[Tree] (config>oam-pm>session>ip>twamp-light>loss-events undet-unavailability-event)

Full Context

configure oam-pm session ip twamp-light loss-events undet-unavailability-event

Description

This command sets the threshold to be applied to the overall count of the undetermined unavailability indicators, not transitions, per configured direction. This value is compared to the 32 bit unavailability counter specific to the direction which tracks the number of individual delta-ts that have been recorded as undetermined unavailable. The aggregate is a function of summing forward and backward. This value is only used as a threshold mechanism and is not part of the stored statistics. If the **clear** *clear-threshold* parameter is not specified the traffic crossing alarm is stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and without regard to any previous window. Each unique event can only be raised once within measurement interval. If the optional clear threshold is specified the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to following measurement intervals. If a threshold crossing event is raised another is not raised until a measurement interval completes and the clear threshold has not been exceeded. A clear event is raised under that condition.

The **no** form of this command removes the event threshold for frame loss ratio. The direction must be included with the **no** command.

Default

no undet-unavailable-event forward

no undet-unavailable-event backward

no undet-unavailable-event aggregate

Parameters

forward

Specifies the threshold is applied to the forward direction count.

backward

Specifies the threshold is applied to the backward direction count.

aggregate

Specifies the threshold is applied to the aggregate count (sum of forward and backward).

raise-threshold

Specifies the rising threshold that determines when the event is to be generated, when value reached.

Values 1 to 864000

clear-threshold

Specifies an optional value used for stateful behavior that allows the operator to configure a percentage of loss value lower than the rising percentage to indicate when the clear event should be generated.

Values 0 to 863999

A value of zero means that the undetermined availability counter must be 0.

Platforms

7705 SAR Gen 2

31.11 undo

undo

Syntax

undo [*count*]

Context

[Tree] (candidate undo)

Full Context

candidate undo

Description

This command removes the most recent change(s) done to the candidate. The changes can be reapplied using the **redo** command. All undo or redo history is lost when the operator exits the **edit-cfg** mode. Undo can not be used to recover a candidate that has been discarded with **candidate discard**.

An **undo** command is blocked if another user has made changes in the same CLI branches that would be impacted during the undo.

Parameters

count

Specifies the number of previous changes to remove.

Values	1 to 50
Default	1

Platforms
7705 SAR Gen 2

31.12 uni

uni

Syntax
uni

Context
[Tree] (config>system>security>keychain>direction uni)

Full Context
configure system security keychain direction uni

Description
This command configures keys for send or receive stream directions.

Platforms
7705 SAR Gen 2

31.13 unicast-address

unicast-address

Syntax
[**no**] **unicast-address** *ip-address*

Context
[Tree] (config>service>vprn>rip>group>neighbor unicast-address)

Full Context
configure service vprn rip group neighbor unicast-address

Description

This command configures the unicast IPv4 address, RIP updates messages will be sent to if the RIP **send** command is set to **send unicast**.

Multiple unicast-address entries can be configured, in which case unicast messages will be sent to each configured unicast IPv4 address.

The **no** form of this command deletes the specified IPv4 unicast address from the configuration.

Parameters

ip-address

Specifies the unicast IPv4 address in a.b.c.d format.

Platforms

7705 SAR Gen 2

unicast-address

Syntax

[**no**] **unicast-address** *ipv6-address*

Context

[Tree] (config>service>vprn>ripng>group>neighbor unicast-address)

Full Context

configure service vprn ripng group neighbor unicast-address

Description

This command configures the unicast IPv6 address, RIPng updates messages will be sent to if the RIPng **send** command is set to **send unicast**.

Multiple unicast-address entries can be configured, in which case unicast messages will be sent to each configured unicast IPv6 address.

The **no** form of this command deletes the specified IPv6 unicast address from the configuration.

Parameters

ipv6-address

Specifies the unicast IPv6 address.

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x [0 to FFFF]H
	d [0 to 255]D

Platforms

7705 SAR Gen 2

unicast-address

Syntax

[**no**] **unicast-address** *ipv6-address*

Context

[Tree] (config>router>ripng>group>neighbor unicast-address)

[Tree] (config>router>rip>group>neighbor unicast-address)

Full Context

configure router ripng group neighbor unicast-address

configure router rip group neighbor unicast-address

Description

This command configures the unicast IPv6 address that RIP and RIPng update messages will be sent to if the **send** command is set to **send unicast**.

Multiple unicast-address entries can be configured, in which case unicast messages will be sent to each configured unicast IPv6 address.

The **no** form of the command deletes the specified IPv6 unicast address from the configuration.

Parameters

ipv6-address

Specifies the IPv6 unicast address to which unicast RIP or RIPng updates should be sent.

Platforms

7705 SAR Gen 2

31.14 unicast-import-disable

unicast-import-disable

Syntax

[**no**] **unicast-import-disable** [**ipv4**]

[**no**] **unicast-import-disable** [**ipv6**]

[**no**] **unicast-import-disable** [**both**]

Context

[Tree] (config>service>vprn>isis unicast-import-disable)

Full Context

configure service vprn isis unicast-import-disable

Description

This command allows one IGP to import its routes into RPF RTM while another IGP imports routes only into the unicast RTM. Import policies can redistribute routes from an IGP protocol into the RPF RTM (the multicast routing table). By default, the IGP routes will not be imported into RPF RTM as such an import policy must be explicitly configured.

Default

no unicast-import-disable

Parameters

ipv4

Allows importation of IPv4 routes only.

ipv6

Allows importation of IPv6 routes only.

both

Allows importation of both IPv4 and IPv6 routes.

Platforms

7705 SAR Gen 2

unicast-import-disable

Syntax

[**no**] **unicast-import-disable**

Context

[Tree] (config>service>vprn>ospf unicast-import-disable)

Full Context

configure service vprn ospf unicast-import-disable

Description

This command allows one IGP to import its routes into RPF RTM while another IGP imports routes only into the unicast RTM.

Import policies can redistribute routes from an IGP protocol into the RPF RTM (the multicast routing table). By default, the IGP routes will not be imported into RPF RTM as such an import policy must be explicitly configured

Default

no unicast-import-disable

Platforms

7705 SAR Gen 2

unicast-import-disable

Syntax

[**no**] **unicast-import-disable** [**ipv4**]

[**no**] **unicast-import-disable** [**ipv6**]

[**no**] **unicast-import-disable** [**both**]

Context

[Tree] (config>router>isis unicast-import-disable)

Full Context

configure router isis unicast-import-disable

Description

This command allows one IGP to import its routes into RPF RTM while another IGP imports routes only into the unicast RTM.

Import policies can redistribute routes from an IGP protocol into the RPF RTM (the multicast routing table). By default, the IGP routes are not imported into RPF RTM, thus, an import policy must be explicitly configured.

Default

no unicast-import-disable both

Parameters

ipv4

Allows importation of IPv4 routes only.

ipv6

Allows importation of IPv6 routes only.

both

Allows importation of both IPv4 and IPv6 routes.

Platforms

7705 SAR Gen 2

unicast-import-disable

Syntax

[**no**] **unicast-import-disable**

Context

[Tree] (config>router>ospf unicast-import-disable)

[Tree] (config>router>ospf3 unicast-import-disable)

Full Context

configure router ospf unicast-import-disable

configure router ospf3 unicast-import-disable

Description

This command allows one IGP to import its routes into RPF RTM while another IGP imports routes only into the unicast RTM. Import policies can redistribute routes from an IGP protocol into the RPF RTM (the multicast routing table). By default, the IGP routes are not imported into RPF RTM as such an import policy must be explicitly configured.

Default

no unicast-import-disable

Platforms

7705 SAR Gen 2

31.15 unicast-rt-test

unicast-rt-test

Syntax

[**no**] **unicast-rt-test**

Context

[Tree] (config>filter>redirect-policy>dest unicast-rt-test)

Full Context

configure filter redirect-policy destination unicast-rt-test

Description

This command configures a unicast route test for this destination. A destination is eligible for redirect if a valid unicast route to that destination exists in the routing instance specified by **config>filter>redirect-policy>router**. The unicast route test is mutually exclusive with other redirect-policy test types.

The test cannot be configured if **no router** is configured for this redirect policy.

The **no** form of the command disables the test.

Default

no unicast-rt-test

Platforms

7705 SAR Gen 2

31.16 unknown-arp-request-flood-evpn

unknown-arp-request-flood-evpn

Syntax

[**no**] **unknown-arp-request-flood-evpn**

Context

[Tree] (config>service>vpls>proxy-arp unknown-arp-request-flood-evpn)

Full Context

configure service vpls proxy-arp unknown-arp-request-flood-evpn

Description

This command controls whether unknown ARP-requests are flooded into the EVPN network. By default, the system floods ARP-requests, including EVPN (with source squelching), if there is no active proxy-arp entry for the requested IP.

The **no** form of the command will only flood to local SAPs/SDP-bindings and not to EVPN destinations.

Default

unknown-arp-request-flood-evpn

Platforms

7705 SAR Gen 2

31.17 unknown-mac-route

unknown-mac-route

Syntax

[**no**] **unknown-mac-route**

Context

[Tree] (config>service>vpls>bgp-evpn unknown-mac-route)

Full Context

configure service vpls bgp-evpn unknown-mac-route

Description

This command enables the advertisement of the unknown-mac-route in BGP. This will be coded in an EVPN MAC route where the MAC address is zero and the MAC address length 48. By using this unknown-mac-route advertisement, the user may decide to optionally turn off the advertisement of MAC addresses learned from SAPs and SDP-bindings, hence reducing the control plane overhead and the size of the FDB tables in the data center. All the receiving NVEs supporting this concept will send any unknown-unicast packet to the owner of the unknown-mac-route, as opposed to flooding the unknown-unicast traffic to all other nodes part of the same VPLS. Although the 7705 SAR Gen 2 can be configured to generate and advertise the unknown-mac-route, the router will never honor the unknown-mac-route and will flood to the vpls flood list when an unknown-unicast packet arrives to an ingress SAP/SDP-binding.

Use of the unknown-mac-route is only supported for BGP-EVPN VXLAN.

Default

no unknown-mac-route

Platforms

7705 SAR Gen 2

31.18 unknown-message-rate

unknown-message-rate

Syntax

unknown-message-rate *integer*

no unknown-message-rate

Context

[Tree] (config>router>pcep>pcc unknown-message-rate)

Full Context

configure router pcep pcc unknown-message-rate

Description

This command configures the maximum rate of unknown messages which can be received on a PCEP session.

When the rate of received unrecognized or unknown messages reaches the configured limit, the PCEP speaker closes the session to the peer.

The **no** form of the command returns the unknown message rate to the default value.

Default

unknown-message-rate 10

Parameters

integer

the rate of unknown messages, in messages per minute

Values 1 to 255

Platforms

7705 SAR Gen 2

31.19 unknown-ns-flood-evpn

unknown-ns-flood-evpn

Syntax

[**no**] **unknown-ns-flood-evpn**

Context

[Tree] (config>service>vpls>proxy-nd unknown-ns-flood-evpn)

Full Context

configure service vpls proxy-nd unknown-ns-flood-evpn

Description

This command controls whether unknown Neighbor Solicitation messages are flooded into the EVPN network. By default, the system floods NS (with source squelching) to SAPs/SDP-bindings including EVPN, if there is no active proxy-nd entry for the requested IPv6.

The **no** form of the command will only flood to local SAPs/SDP-bindings but not to EVPN destinations.

Default

unknown-ns-flood-evpn

Platforms

7705 SAR Gen 2

31.20 unknown-policer

unknown-policer

Syntax

unknown-policer *policer-id* [**fp-redirect-group**]

no unknown-policer

Context

[Tree] (config>qos>sap-ingress>fc unknown-policer)

Full Context

configure qos sap-ingress fc unknown-policer

Description

Within a **sap-ingress** QoS policy forwarding class context, the **unknown-policer** command is used to map packets that match the forwarding class and are considered unknown in nature to the specified policer-id. The specified policer-id must already exist within the **sap-ingress** QoS policy. While the system is determining the forwarding class of a packet, it is also looking up its forwarding destination based on the ingress service type and the service instance forwarding records. If the service type is VPLS and the destination MAC address is unicast, but the MAC has not been learned and populated within the VPLS services FDB, the packet is classified into the unknown forwarding type.

Unknown forwarding type packets are mapped to either an ingress multipoint queue (using the **unknown** *queue-id* or **unknown** *queue-id* **group** *ingress-queue-group* commands) or an ingress policer (**unknown-policer** *policer-id*). The **unknown** and **unknown-policer** commands within the forwarding class context are mutually exclusive. By default, the unknown forwarding type is mapped to the SAP ingress default multipoint queue. If the **unknown-policer** *policer-id* command is executed, any previous policer mapping or queue mapping for the unknown forwarding type within the forwarding class is overridden if the policer mapping is successful.

A policer defined within the **sap-ingress** policy is not actually created on an ingress SAP or a subscriber using an **sla-profile** where the policy is applied until at least one forwarding type (unicast, broadcast, unknown, or multicast) from one of the forwarding classes is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or multiservice site, or ingress policing is not supported on the port associated with the SAP or subscriber or multiservice site, the initial forwarding class forwarding type mapping will fail.

The **unknown-policer** command is ignored for instances of the policer applied to SAPs or subscribers' multiservice site where unknown packets are not supported.

When the unknown forwarding type within a forwarding class is mapped to a policer, the unknown packets classified to the subclasses within the forwarding class are also mapped to the policer.

The **no** form of this command is used to restore the mapping of the unknown forwarding type within the forwarding class to the default multipoint queue. If all forwarding class forwarding types had been removed from the default multipoint queue, the queue will not exist on the SAPs or subscriber or multiservice site associated with the QoS policy and the **no broadcast-policer** command will cause the system to attempt to create the default multipoint queue on each object. If the system cannot create the queue on each instance, the **no unknown-policer** command will fail and the unknown forwarding type within the forwarding class will continue its mapping to the existing *policer-id*. If the **no unknown-policer** command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscriber will be lost.

Parameters

policer-id

When the forwarding class **unknown-policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the **sap-ingress** QoS policy.

Values 1 to 63

fp-redirect-group

Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

Platforms

7705 SAR Gen 2

31.21 unknown-queue

unknown-queue

Syntax

unknown-queue *queue-id* [**group** *queue-group-name*]

no unknown-queue

Context

[Tree] (config>qos>sap-ingress>fc unknown-queue)

Full Context

configure qos sap-ingress fc unknown-queue

Description

This command overrides the default unknown unicast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. When the forwarding class mapping is executed, all unknown traffic on a SAP using this policy is forwarded using the *queue-id*.

The unknown forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.

The **no** form of this command sets the unknown forwarding type *queue-id* back to the default of tracking the multicast forwarding type queue mapping.

Parameters

queue-id

Specifies an existing multipoint queue defined in the **config>qos>sap-ingress** context.

Values Any valid multipoint *queue-id* in the policy including 2 through 32.

Default 11

group *queue-group-name*

This optional parameter is used to redirect the forwarding type within the forwarding class to the specified queue-id within the queue-group-name. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the specified queue group. The *queue-group-name* are configured in the **config>qos>queue-group-templates** egress and ingress contexts.

Platforms

7705 SAR Gen 2

31.22 unnumbered

unnumbered

Syntax

unnumbered {*ip-int-name* | *ip-address*}

no unnumbered

Context

[Tree] (config>service>ies>if unnumbered)

Full Context

configure service ies interface unnumbered

Description

This command configures the interface as an unnumbered interface. Unnumbered IP interfaces are supported on a SONET/SDH access port with the PPP, ATM, Frame Relay, cisco-HDLC encapsulation. It is not supported on access ports that do not carry IP traffic, but are used for native TDM circuit emulation.

Parameters

ip-int-name

Specifies the name of an IP interface. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

ip-address

Specifies an IP address.

Platforms

7705 SAR Gen 2

unnumbered

Syntax

unnumbered {*ip-int-name* | *ip-address*}

no unnumbered

Context

[Tree] (config>service>ies>if unnumbered)

Full Context

configure service ies interface unnumbered

Description

This command configures the interface as an unnumbered interface. Unnumbered IP interfaces are supported on a SONET/SDH access port with the PPP, ATM, Frame Relay, cisco-HDLC encapsulation. It is not supported on access ports that do not carry IP traffic, but are used for native TDM circuit emulation.

Parameters

ip-int-name

Specifies the name of an IP interface. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

ip-address

Specifies an IP address.

Platforms

7705 SAR Gen 2

unnumbered

Syntax

unnumbered [*ip-int-name* | *ip-address*]

no unnumbered

Context

[Tree] (config>service>vprn>if unnumbered)

Full Context

configure service vprn interface unnumbered

Description

This command configures the interface as an unnumbered interface. An unnumbered IP interface is supported on a SONET/SDH access port with the PPP, ATM, Frame Relay, cisco-HDLC encapsulation. It is not supported on access ports that do not carry IP traffic, but are used for native TDM circuit emulation.

Parameters

ip-int-name

Specifies the name of an IP interface. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

ip-address

Specifies an IP address.

Platforms

7705 SAR Gen 2

unnumbered

Syntax

unnumbered [*ip-int-name* | *ip-address*]

no unnumbered

Context

[Tree] (config>service>vprn>if unnumbered)

Full Context

configure service vprn interface unnumbered

Description

This command configures the interface as an unnumbered interface. An unnumbered IP interface is supported on a SONET/SDH access port with the PPP, ATM, Frame Relay, cisco-HDLC encapsulation. It is not supported on access ports that do not carry IP traffic, but are used for native TDM circuit emulation.

Parameters

ip-int-name

Specifies the name of an IP interface. If the string contains special characters (#, ?, space), the entire string must be enclosed between double quotes.

ip-address

Specifies an IP address.

Platforms

7705 SAR Gen 2

unnumbered

Syntax

unnumbered [{*ip-int-name* | *ip-address*}]

no unnumbered

Context

[Tree] (config>router>if unnumbered)

Full Context

configure router interface unnumbered

Description

This command sets an IP interface as an unnumbered interface and specifies the IP address to be used for the interface.

To conserve IP addresses, unnumbered interfaces can be configured. The address used when generating packets on this interface is the *ip-addr* parameter configured.

An error message will be generated if an **unnumbered** interface is configured, and an IP address already exists on this interface.

The **no** form of this command removes the IP address from the interface, effectively removing the unnumbered property. The interface must be **shutdown** before **no unnumbered** is issued to delete the IP address from the interface, or an error message will be generated.

Default

no unnumbered

Parameters

ip-int-name* | *ip-address

Optional. Specifies the IP address or IP interface name to associate with the unnumbered IP interface in dotted decimal notation. The configured IP address must exist on this node. It is recommended to use the system IP address as it is not associated with a specific interface and is therefore always reachable. The system IP address is the default if no *ip-addr* or *ip-int-name* is configured.

Platforms

7705 SAR Gen 2

31.23 unreachables

unreachables

Syntax

unreachables [*number seconds*]

no unreachables[*number seconds*]

Context

[Tree] (config>service>vprn>if>ipv6>icmp6 unreachables)

[Tree] (config>service>vprn>if>icmp unreachables)

[Tree] (config>service>vprn>nw-if>icmp unreachables)

[Tree] (config>service>ies>if>icmp unreachables)

Full Context

configure service vprn interface ipv6 icmp6 unreachables

configure service vprn interface icmp unreachables

configure service vprn network-interface icmp unreachables

configure service ies interface icmp unreachables

Description

This command configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

The **unreachables** command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.

By default, generation of ICMP destination unreachables messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of ICMP destination unreachable messages on the router interface and reverts to the default values.

Default

unreachables 100 10

Parameters

number

Specifies the maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 to 2000

seconds

Specifies the time frame, in seconds, used to limit the *number* of ICMP unreachable messages that can be issued.

Values 1 to 60

Platforms

7705 SAR Gen 2

unreachables

Syntax

unreachables [*number seconds*]

no unreachables

Context

[Tree] (config>service>ies>if>ipv6>icmp6 unreachables)

Full Context

configure service ies interface ipv6 icmp6 unreachables

Description

This command specifies that ICMPv6 host and network unreachable messages are generated by this interface.

When disabled, ICMPv6 host and network unreachable messages are not sent.

The **no** form of this command reverts to the default.

Default

unreachables 100 10

Parameters

number

Specifies the number of destination unreachable ICMPv6 messages are issued in the time frame specified by the *seconds* parameter.

Values 10 to 2000

seconds

Specifies the time frame, in seconds, that is used to limit the number of destination unreachable ICMPv6 messages to be issued.

Values 1 to 60

Platforms

7705 SAR Gen 2

unreachables

Syntax

unreachables [*number seconds*]

no unreachables

Context

[Tree] (config>router>if>icmp unreachables)

Full Context

configure router interface icmp unreachables

Description

This command enables and configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

The **unreachables** command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of destination unreachable messages that can be issued on the interface for a given time interval.

By default, generation of ICMP destination unreachables messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of ICMP destination unreachables on the router interface.

Default

unreachables 100 10 — Maximum of 100 unreachable messages in 10 seconds.

Parameters

number

The maximum number of ICMP unreachable messages to send, expressed as a decimal integer. The *seconds* parameter must also be specified.

Values 10 to 2000

seconds

The time frame, in seconds, used to limit the *number* of ICMP unreachable messages that can be issued, expressed as a decimal integer.

Values 1 to 60

Platforms

7705 SAR Gen 2

unreachables

Syntax

unreachables [*number seconds*]

no unreachables

Context

[Tree] (config>router>if>ipv6>icmp6 unreachables)

Full Context

configure router interface ipv6 icmp6 unreachables

Description

This command configures the rate for ICMPv6 unreachable messages. When enabled, ICMPv6 host and network unreachable messages are generated by this interface.

The **no** form of this command disables the generation of ICMPv6 host and network unreachable messages by this interface.

Default

unreachables 100 10 (when IPv6 is enabled on the interface)

Parameters

number

Determines the number destination unreachable ICMPv6 messages to issue in the time frame specified in *seconds* parameter.

Values 10 to 2000

seconds

Sets the time frame, in seconds, to limit the number of destination unreachable ICMPv6 messages issued per time frame.

Values 1 to 60

Platforms

7705 SAR Gen 2

31.24 untrusted

untrusted

Syntax

untrusted [**default-forwarding** *default-forwarding*]

no untrusted

Context

[Tree] (config>router>if untrusted)

Full Context

configure router interface untrusted

Description

This command configures the state of untrusted for a network IP interface.

The untrusted state identifies the participating interfaces in the label security feature for prefixes of a VPN family at an inter-AS boundary. The router supports a maximum of 15 network interfaces that can participate in this feature.

The user normally applies this command to an inter-AS interface. PIP keeps track of the untrusted status of each interface. In the data path, such an interface causes the default forwarding to be set to the *default-forwarding* value.

For backward compatibility reasons, the interface **default-forwarding** is set to the **forward** value; this means that labeled packets are checked in the normal way against the table of programmed ILMs to decide if they should be dropped or forwarded in a GRT, a VRF, or a L2 service context.

If the user sets the *default-forwarding* value to **drop**, all labeled packets received on that interface are automatically dropped.

This command sets the default behavior for an untrusted interface in the data path and for all ILMs. When enabling the label security for VPN IPv4 or VPN IPv6 prefixes, BGP programs the data path to provide an exception to the normal way of forwarding handling away from the default for those VPRN ILMs.

The **no** form of this command returns the interface into the default state of trusted.

Default

no untrusted

Parameters

default-forwarding

Specifies the default forwarding behavior of labeled packets received on this interface.

Values forward, drop

Default forward

Platforms

7705 SAR Gen 2

31.25 unzip

unzip

Syntax

unzip *source-file-url* [*dest-file-url*] **list**

unzip *source-file-url dest-file-url* [**create-destination**] [**force**]

Context

[Tree] (file unzip)

Full Context

file unzip

Description

This command expands the contents of a ZIP file to the local file system. Any file that is zipped using the store, deflate, or zip64 compression methods can be unzipped. The source ZIP file location can be a locally installed solid-state storage device or a remote FTP or TFTP server. Files can only be unzipped to the active CPM.

Parameters

source-file-url, dest-file-url

Specifies the source or destination file URL.

Values

local-url	[*cflash-id*/]*file-path*
	200 chars max, including cflash-id
	directory length 99 chars max each

remote-url	{ftp | tftp}://[*login*:*pswd*@]
	remote-locn / [*file-path*]
	247 chars max, *file-path* 199 chars max
remote-locn	{*hostname* | *ipv4-address* | "["i*pv6-address*"]" }[:*port*]|
ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x[-interface]
	x:x:x:x:x:x:d.d.d.d[-interface]
	x: [0 to FFFF]H
	d: [0 to 255]D
	interface - 32 characters max, for link local addresses
port	[0 to 65535]
cflash-id	cf1: | cf1-A: | cf2: | cf2-A: | cf3: | cf3-A:

create-destination

Specifies that a non-existent directory structure that is explicitly entered as the destination file URL is created as part of the unzip operation. This parameter is required to create new directories.

list

Lists the content of the ZIP file without performing the unzip operation.

force

Overwrites without prompting, any file or directory contained within the ZIP file that already exists in the destination URL. This keyword does not automatically create new directories explicitly specified by *dest-file-url*. To create these directories, use the **create-destination** flag.

Platforms

7705 SAR Gen 2

31.26 up

up

Syntax

up ip *seconds*

no up ip

up ipv6 *seconds*

no up ipv6

Context

[Tree] (config>service>ies>if>hold-time up)

[Tree] (config>service>vprn>nw-if>hold-time up)

[Tree] (config>service>vprn>if>hold-time up)

[Tree] (config>service>vpls>if>hold-time up)

Full Context

configure service ies interface hold-time up

configure service vprn network-interface hold-time up

configure service vprn interface hold-time up

configure service vpls interface hold-time up

Description

This command causes a delay in the deactivation of the associated IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface down.

The **no** form of this command removes the command from the active configuration and removes the delay in deactivating the associated IP interface. If the configuration is removed during a delay period, the currently running delay will continue until it expires.

Default

no up ip

Parameters

seconds

The time delay, in seconds, to make the interface operational.

Values 1 to 1200

Platforms

7705 SAR Gen 2

31.27 upa-lifetime

upa-lifetime

Syntax

upa-lifetime *upa-lifetime*

no upa-lifetime

Context

[Tree] (config>router>isis>upa upa-lifetime)

Full Context

configure router isis prefix-unreachable upa-lifetime

Description

This command configures the amount of time a UPA is advertised.

The **no** form of this command reverts to the default.

Default

180

Parameters

upa-lifetime

Specifies the amount of time, in seconds, the UPA is advertised.

Values 30 to 1800

Platforms

7705 SAR Gen 2

31.28 upa-metric

upa-metric

Syntax

upa-metric *upa-metric*

no upa-metric

Context

[Tree] (config>router>isis>upa upa-metric)

Full Context

configure router isis prefix-unreachable upa-metric

Description

This command configures a specific metric to an advertised UPA.

The **no** form of this command reverts to the default.

Default

4261412865

Parameters

upa-metric

Specifies the metric to an advertised UPA.

Values 4261412865 to 4294967294

Platforms

7705 SAR Gen 2

31.29 update

update

Syntax

update [**neighbor** *ip-address* | **group** *name*]

no update

Context

[Tree] (debug>router>bgp update)

Full Context

debug router bgp update

Description

This command decodes and logs all sent and received update messages in the debug log.

The **no** form of this command disables debugging.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

Values ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x:x [-interface] (eight 16-bit pieces)
- x:x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H

- d: [0 to 255]D
- interface: up to 32 characters for link local addresses

group *name*

Debugs only events affecting the specified peer group name, up to 64 characters, and associated neighbors.

Platforms

7705 SAR Gen 2

Output

The following output is an example of debug router BGP update information.

Output Example

```
debug router bgp update
```

```
17 2022/05/04 17:39:07.566 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 76
    Flag: 0x90 Type: 14 Len: 32 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.4
        [VPLS/VPWS] preflen 21, veid: 4, vbo: 5, vbs: 1, label-base: 524252, RD
192.0.2.4:801, csv: 0x00000000, type 1, len 1,
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:801
        l2-vpn/vrf-imp:Encap=5: Flags=-TRC: MTU=1514: PREF=0
```

```
158 2022/05/10 08:05:21.767 UTC MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2
"Peer 1: 2001:db8::2: UPDATE
Peer 1: 2001:db8::2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-AD Len: 25 RD: 192.0.2.5:201 ESI: ESI-0, tag: 5 Label: 838804
8 (Raw Label: 0x7ffdd0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:201
        l2-attribute:MTU: 1514 C: 1 F: 1 P: 0 B: 0
        bgp-tunnel-encap:MPLS
"

367 2022/05/10 08:04:47.560 UTC MINOR: DEBUG #2001 Base Peer 1: 2001:db8::5
"Peer 1: 2001:db8::5: UPDATE
```

```
Peer 1: 2001:db8::5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 77
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.2:500, tag: 0, orig_addr len:
32, orig_addr: 192.0.2.2
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:500
        l2-attribute:MTU: 1514 C: 1 F: 1 P: 0 B: 0
        bgp-tunnel-encap:MPLS
    Flag: 0xc0 Type: 22 Len: 9 PMSI:
        Tunnel-type Ingress Replication (6)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label 8388512
        Tunnel-Endpoint 192.0.2.2
```

```
2840 2024/09/04 18:22:17.332 UTC MINOR: DEBUG #2001 Base Peer 1: 2001:db8::1
"Peer 1: 2001:db8::1: UPDATE
Peer 1: 2001:db8::1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 93
    Flag: 0x90 Type: 14 Len: 48 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 16 Global NextHop 2001:db8::2
        Type: EVPN-AD Len: 25 RD: 192.0.2.2:500 ESI: 01:66:00:00:00:00:00:00:00:
00, tag: 23 Label: 8388416 (Raw Label: 0x7fff40) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:500
        l2-attribute:MTU: 1514 V: Double-VID M: Default F: 1 C: 1 P: 1 B: 0
        bgp-tunnel-encap:MPLS
"
```

31.30 update-fault-tolerance

update-fault-tolerance

Syntax

[**no**] **update-fault-tolerance**

Context

[Tree] (config>service>vprn>bgp>group>neighbor>error-handling update-fault-tolerance)

[Tree] (config>service>vprn>bgp>group>error-handling update-fault-tolerance)

[Tree] (config>service>vprn>bgp>error-handling update-fault-tolerance)

Full Context

configure service vprn bgp group neighbor error-handling update-fault-tolerance

configure service vprn bgp group error-handling update-fault-tolerance

configure service vprn bgp error-handling update-fault-tolerance

Description

This command enables **treat-as-withdraw** and other similarly non-disruptive approaches for handling a wide range of UPDATE message errors, as long as there are no length errors that prevent all of the NLRI fields from being correctly identified and parsed.

Default

no update-fault-tolerance

Platforms

7705 SAR Gen 2

update-fault-tolerance

Syntax

[**no**] **update-fault-tolerance**

Context

[Tree] (config>router>bgp>group>neighbor>error-handling update-fault-tolerance)

[Tree] (config>router>bgp>group>error-handling update-fault-tolerance)

[Tree] (config>router>bgp>error-handling update-fault-tolerance)

Full Context

configure router bgp group neighbor error-handling update-fault-tolerance

configure router bgp group error-handling update-fault-tolerance

configure router bgp error-handling update-fault-tolerance

Description

This command enables **treat-as-withdraw** and other similarly non-disruptive approaches for handling a wide range of UPDATE message errors, as long as there are no length errors that prevent all of the NLRI fields from being correctly identified and parsed.

Default

no update-fault-tolerance

Platforms

7705 SAR Gen 2

31.31 update-interval

update-interval

Syntax

update-interval *minutes* [**jitter** *seconds*]

no update-interval

Context

[Tree] (config>ipsec>rad-acct-plcy update-interval)

Full Context

configure ipsec radius-accounting-policy update-interval

Description

This command enables the system to send RADIUS interim-update packets for IKEv2 remote-access tunnels. The RADIUS attributes in the interim-update packet are the same as acct-start. The value of the Acct-status-type in the interim-update message is 3.

Default

update-interval 10

Parameters

minutes

Specifies the interval in minutes.

Values 5 to 259200

seconds

Specifies the jitter as the number of seconds when the system sends each interim-update packet.

Values 0 to 3600

Platforms

7705 SAR Gen 2

31.32 update-key

update-key

Syntax

update-key card *cpm-slot* **serial-number** *cpm-serial-number* **confirmation-code** *code* **software-image** *file-url*

Context

[Tree] (admin>system>security>secure-boot update-key)

Full Context

admin system security secure-boot update-key

Description

This command updates secure boot keys.

Parameters

cpm-slot

Specifies the CPM slot.

Values A,B

cpm-serial-number

Specifies the CPM serial number, up to 256 characters.

code

Specifies the signed software confirmation code, up to 32 characters.

file-url

Specifies the URL for the software image.

Values [*local-url* | *remote-url*] (up to 180 characters)

where:

- *local-url* — [*cflash-id*/] [*file-path*]
  180 chars max, including *cflash-id*
  directory length 99 chars max each
- *remote-url* — [{ftp://| tftp://} *login*:*pswd*@*remote-locn*/][*file-path*]
  180 chars max
  directory length 99 chars max each
  where: *remote-locn* — [*hostname* | *ipv4-address* | *ipv6-address*]

  ipv4-address a.b.c.d

ipv6-address	x:x:x:x:x:x:x:x[-interface] x:x:x:x:x:x:d.d.d.d[-interface] x - [0..FFFF]H d - [0..255]D interface - 32 chars max, for link local addresses
cflash-id	cf1:| cf1-A:| cf1-B:| cf2:| cf2-A:| cf2-B:| cf3:| cf3-A:| cf3-B:

Platforms

7705 SAR Gen 2

31.33 update-timer

update-timer

Syntax

update-timer *seconds*

no update-timer

Context

[Tree] (config>router>rsvp>te-threshold-update update-timer)

Full Context

configure router rsvp te-threshold-update update-timer

Description

This command is to control timer-based IGP TE updates. Timer-based IGP updates can be enabled by specifying a non-zero time value. Default value of update-timer is 0.

The **no** form of this command should reset update-timer to the default value and disable timer-based IGP update.

Default

no update-timer

Parameters

seconds

Specifies the time in seconds.

Values 0 to 300

Platforms

7705 SAR Gen 2

31.34 updates

updates

Syntax

[**no**] **updates** [**neighbor** *ip-int-name* | *ip-address*]

Context

[Tree] (debug>router>rip updates)

Full Context

debug router rip updates

Description

This command enables debugging for RIP updates.

Parameters

ip-int-name* | *ip-address

Debugs the RIP updates sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

updates

Syntax

[**no**] **updates** [**neighbor** *ip-int-name* | *ipv6-address*]

Context

[Tree] (debug>router>ripng updates)

Full Context

debug router ripng updates

Description

This command enables debugging for RIP updates.

Parameters

ip-int-name*| *ipv6-address

Debugs the RIP updates sent on the neighbor IP address or interface.

Platforms

7705 SAR Gen 2

31.35 upstream-ip-filter

upstream-ip-filter

Syntax

upstream-ip-filter *filter-id*

no upstream-ip-filter

Context

[Tree] (config>router>nat>outside upstream-ip-filter)

[Tree] (config>service>vprn>nat>outside upstream-ip-filter)

Full Context

configure router nat outside upstream-ip-filter

configure service vprn nat outside upstream-ip-filter

Description

This command configures the ip-filter for upstream traffic. This filter is applied to the upstream traffic after the NAT function and before it enters the outside virtual router instance; it is useful for traffic that bypasses the ingress filters applied in the inside virtual router instance, such as DS-Lite traffic.

Default

no upstream-ip-filter

Parameters

filter-id

Specifies the identifier of an IP filter.

Platforms

7705 SAR Gen 2

31.36 url

url

Syntax

url *url*

no url

Context

[Tree] (config>system>security>pki>ca-prof>auto-crl-update>crl-urls>url-entry url)

Full Context

configure system security pki ca-profile auto-crl-update crl-urls url-entry url

Description

This command specifies the HTTP URL of the CRL file for the **url-entry**. The system supports both IPv4 and IPv6 HTTP connections.

Note:
The URL must point to a DER encoded CRL.

Default

no url

Parameters

url

Specifies the URL, which specifies the location, where an updated CRL can be downloaded from.

Platforms

7705 SAR Gen 2

url

Syntax

url *url-string* [**service-id** *service-id*]

url *url-string* [**service-name** *service-name*]

no url

Context

[Tree] (config>system>security>pki>ca-profile>cmpv2 url)

Full Context

configure system security pki ca-profile cmpv2 url

Description

This command specifies HTTP URL of the CMPv2 server. The URL must be unique across all configured ca-profiles.

The URL is resolved by the DNS server configured (if configured) in the corresponding router context.

If the *service-id* is 0 or omitted, then system tries to resolve the FQDN via DNS server configured in bof.cfg. After resolution, the system connects to the address in the management routing instance first, then the base routing instance.

Note:
If the service is VPRN, the system only allows HTTP ports 80 and 8080.

Parameters

url-string

Specifies the HTTP URL of the CMPv2 server, up to 180 characters.

service-id ***service-id***

Specifies the service instance that used to reach CMPv2 server.

This variant of this command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **url** *url-string* **service-name** *service-name* variant can be used in all configuration modes.

Values service-id: 1 to 2147483647 base-router: 0

service-name service-name

Identifies the service, up to 64 characters.

Platforms

7705 SAR Gen 2

31.37 url-entry

url-entry

Syntax

url-entry *entry-id* [**create**]

no url-entry *entry-id*

Context

[Tree] (config>system>security>pki>ca-prof>auto-crl-update>crl-urls url-entry)

Full Context

configure system security pki ca-profile auto-crl-update crl-urls url-entry

Description

This command creates a new **crl-url** entry with the **create** parameter, or enters an existing url-entry configuration context without **create** parameter.

The **no** form of this command removes the specified entry.

Parameters

entry-id

Specifies a URL configured on this system.

Values 1 to 8

create

Creates an auto URL entry.

Platforms

7705 SAR Gen 2

31.38 urpf-check

urpf-check

Syntax

[**no**] **urpf-check**

Context

[Tree] (config>service>ies>if urpf-check)

[Tree] (config>service>vprn>if urpf-check)

[Tree] (config>service>vprn>if>ipv6 urpf-check)

[Tree] (config>service>vprn>nw-if urpf-check)

[Tree] (config>service>ies>if>ipv6 urpf-check)

Full Context

configure service ies interface urpf-check

configure service vprn interface urpf-check

configure service vprn interface ipv6 urpf-check

configure service vprn network-interface urpf-check

configure service ies interface ipv6 urpf-check

Description

This command enables unicast RPF (uRPF) check on this interface.

The **no** form of this command disables unicast RPF (uRPF) Check on this interface.

Default

no urpf-check

Platforms

7705 SAR Gen 2

urpf-check

Syntax

urpf-check

no urpf-check

Context

[Tree] (config>service>vprn>network>ingress urpf-check)

Full Context

configure service vprn network ingress urpf-check

Description

This command enables the unicast RPF (uRPF) check of network ingress traffic to include traffic associated with the VPRN if the incoming network interface is configured with the **urpf-selected-vprns** command

If the command is not configured, then traffic associated with this VPRN that arrives on a network interface with **urpf-selected-vprns** configured bypasses the uRPF checking options specified for that network interface.

Default

no urpf-check

Platforms

7705 SAR Gen 2

urpf-check

Syntax

[**no**] **urpf-check**

Context

[Tree] (config>router>if urpf-check)

[Tree] (config>router>if>ipv6 urpf-check)

Full Context

configure router interface urpf-check

configure router interface ipv6 urpf-check

Description

This command enables unicast RPF (uRPF) Check on this interface.

The **no** form of this command disables unicast RPF (uRPF) Check on this interface.

Platforms

7705 SAR Gen 2

31.39 urpf-selected-vprns

urpf-selected-vprns

Syntax

[**no**] **urpf-selected-vprns**

Context

[Tree] (config>router>if urpf-selected-vprns)

Full Context

configure router interface urpf-selected-vprns

Description

This command enables uRPF checking of incoming traffic on the network interface for the following packets.

- Packets associated with the global routing table (base router) context.
- Packets associated with VPRNs that have enabled the uRPF check using the **config**>**service**>**vprn**>**network**> **ingress**>**urpf-check** command.

If the command is not configured, the default action is to perform uRPF checks for all ingress traffic on the network interface (associated with the base router and all VPRNs) based on the IPv4 and IPv6 **urpf-check** configuration options of the network interface.

Default

no urpf-selected-vprns

Platforms

7705 SAR Gen 2

31.40 use-arp

use-arp

Syntax

[**no**] **use-arp**

Context

[Tree] (config>service>ies>if>dhcp use-arp)

[Tree] (config>service>vprn>if>dhcp use-arp)

Full Context

configure service ies interface dhcp use-arp

configure service vprn interface dhcp use-arp

Description

This command enables the use of ARP to determine the destination hardware address.

The **no** form of this command disables the use of ARP to determine the destination hardware address.

Platforms

7705 SAR Gen 2

31.41 use-bgp-routes

use-bgp-routes

Syntax

[**no**] **use-bgp-routes**

Context

[Tree] (config>service>vprn>bgp>next-hop-res use-bgp-routes)

Full Context

configure service vprn bgp next-hop-resolution use-bgp-routes

Description

This command enables the use of BGP routes to resolve BGP next hops. When this command is enabled, any unlabeled IPv4 or IPv6 BGP route received from a VPRN BGP peer becomes resolvable by up to four other BGP routes in order to resolve the route to a VPRN IP interface.

This command also allows unlabeled IPv4 or IPv6 BGP routes leaked from the GRT with unresolved next hops (in the GRT) to be resolvable by BGP-VPN routes (of the VPRN).

The **no** form of this command reverts to the default behavior. By default, a VPRN BGP route is not resolvable by another VPRN BGP route or by a BGP-VPN route.

Default

no use-bgp-routes

Platforms

7705 SAR Gen 2

use-bgp-routes

Syntax

[**no**] **use-bgp-routes**

Context

[Tree] (config>router>bgp>next-hop-res use-bgp-routes)

Full Context

configure router bgp next-hop-resolution use-bgp-routes

Description

This command specifies whether to use BGP routes to recursively resolve the BGP next-hop of unlabeled IPv4 and unlabeled IPv6 routes. Up to four levels of recursion are supported.

The **no** form of this command reverts to the default behavior. By default, a BGP route is not resolvable by another BGP route.

Default

no use-bgp-routes

Platforms

7705 SAR Gen 2

use-bgp-routes

Syntax

use-bgp-routes

Context

[Tree] (config>router>bgp>next-hop-res>lbl-routes use-bgp-routes)

Full Context

configure router bgp next-hop-resolution labeled-routes use-bgp-routes

Description

Commands in this context configure labeled route options for next-hop resolution.

Platforms

7705 SAR Gen 2

31.42 use-default-template

use-default-template

Syntax

[**no**] **use-default-template**

Context

[Tree] (config>service>vprn>aaa>rmt-srv>tacplus use-default-template)

[Tree] (config>system>security>tacplus use-default-template)

Full Context

configure service vprn aaa remote-servers tacplus use-default-template

configure system security tacplus use-default-template

Description

This command specifies whether the **user-template tacplus_default** is actively applied to the TACACS+ user. When enabled, some parameters of the **user-template tacplus_default** are actively applied to all users that authenticate via TACACS+. See the **user-template tacplus_default** command for more details.

When disabled, the parameters of the template are not applied to TACACS+ users, and TACACS+ users cannot connect to an SR OS router since the user access parameters are not available. In this case, TACACS+ can only be used for accounting.

Default

use-default-template

Platforms

7705 SAR Gen 2

use-default-template

Syntax

[**no**] **use-default-template**

Context

[Tree] (config>system>security>radius use-default-template)

[Tree] (config>service>vprn>aaa>rmt-srv>radius use-default-template)

Full Context

configure system security radius use-default-template

configure service vprn aaa remote-servers radius use-default-template

Description

This command specifies whether the RADIUS default user template is actively applied to the RADIUS user if no VSAs are returned with the auth-accept from the RADIUS server. When enabled, the radius_default user-template is actively applied if no VSAs are returned with the auth-accept from the RADIUS server and radius authorization is enabled.

The **no** form of this command disables the use of the RADIUS default template.

Default

no use-default-template

Platforms

7705 SAR Gen 2

use-default-template

Syntax

[**no**] **use-default-template**

Context

[Tree] (config>system>security>ldap use-default-template)

Full Context

configure system security ldap use-default-template

Description

This command specifies whether the default template is to be actively applied to LDAP users.

Default

use-default-template

Platforms

7705 SAR Gen 2

31.43 use-gi-address

use-gi-address

Syntax

use-gi-address [**scope** *scope*]

Context

[Tree] (config>router>dhcp>server use-gi-address)

[Tree] (config>service>vprn>dhcp>server use-gi-address)

Full Context

configure router dhcp local-dhcp-server use-gi-address

configure service vprn dhcp local-dhcp-server use-gi-address

Description

This command enables the use of gi-address matching. If the gi-address flag is enabled, a pool can be used even if a subnets is not found. If the **local-user-db-name** is not used, the gi-address flag is used and addresses are handed out by GI only. If a user must be blocked from getting an address the server maps to a local user database and configures the user with no address.

A pool can include multiple subnets. Since the GI is shared by multiple subnets in a subscriber interface the pool may provide IP addresses from any of the subnets included when the GI is matched to any of its subnets. This allows a pool to be created that represents a sub-int.

The **no** form of the reverts to the default.

Parameters

scope

Specifies if addresses are handed out for a certain subnet where the gi-address belongs to only or for all subnets part of the pool.

Values **subnet** — Addresses are only handed out for the subnet where the gi-address is part.

pool — All subnets part of the pool which contain subnet where the gi-address is part of can hand out addresses.

Platforms

7705 SAR Gen 2

31.44 use-leaked-routes

use-leaked-routes

Syntax

use-leaked-routes

Context

[Tree] (config>router>bgp>next-hop-res use-leaked-routes)

[Tree] (config>service>vprn>bgp>next-hop-res use-leaked-routes)

Full Context

configure router bgp next-hop-resolution use-leaked-routes

configure service vprn bgp next-hop-resolution use-leaked-routes

Description

Commands in this context configure the use of leaked static routes to resolve BGP next hops.

Platforms

7705 SAR Gen 2

31.45 use-link-address

use-link-address

Syntax

use-link-address [**scope** *scope*]

no use-link-address

Context

[Tree] (config>service>vprn>dhcp6>server use-link-address)

[Tree] (config>router>dhcp6>server use-link-address)

Full Context

configure service vprn dhcp6 local-dhcp-server use-link-address

configure router dhcp6 local-dhcp-server use-link-address

Description

This command configures the local pool selection for IPv6 address or prefix assignment for the configured link-address under relay configuration. The selected pool will contain a prefix covering the link-address. The scope option defines the scope for the match. With scope **subnet**, the prefix or address selection is limited to the prefix in the pool that covers the link-address. With scope **pool**, all the prefixes in the selected pool are eligible for assignment.

The **no** form of the reverts to the default.

Default

scope subnet

Parameters

scope

Specifies the scope of the IP address selection.

Values **subnet** — Specifies that the prefix or address selection is limited to the prefix in the pool that covers the link address.

pool — Specifies that all prefixes in the selected pool are eligible for assignment.

Platforms

7705 SAR Gen 2

31.46 use-pool-from-client

use-pool-from-client

Syntax

use-pool-from-client delimiter *delimiter*

use-pool-from-client

no use-pool-from-client

Context

[Tree] (config>router>dhcp>server use-pool-from-client)

[Tree] (config>service>vprn>dhcp>server use-pool-from-client)

Full Context

configure router dhcp local-dhcp-server use-pool-from-client

configure service vprn dhcp local-dhcp-server use-pool-from-client

Description

This command enables the use of the pool indicated by DHCP client. When enabled, the IP address pool to be used by this server is the pool is indicated by the vendor-specific sub-option 13 of the DHCP option 82. When disabled or if there is no sub-option 13 in the DHCP message, the pool selection falls back to the **use-gi-address** configuration.

The **no** form of this command disables the use of the pool indicated by DHCP client.

Parameters

delimiter

A single ASCII character specifies the delimiter of separating primary and secondary pool names in Option82 VSO.

Platforms

7705 SAR Gen 2

31.47 use-virtual-mac

use-virtual-mac

Syntax

[**no**] **use-virtual-mac**

Context

[Tree] (config>service>vprn>router-advert>if use-virtual-mac)

Full Context

configure service vprn router-advertisement interface use-virtual-mac

Description

This command enables sending router advertisement messages using the VRRP virtual MAC address, provided that the virtual router is currently the master.

If the virtual router is not the master, no router advertisement messages are sent.

The **no** form of this command disables sending router advertisement messages.

Default

no use-virtual-mac

Platforms
7705 SAR Gen 2

use-virtual-mac

Syntax
[**no**] **use-virtual-mac**

Context
[Tree] (config>router>router-advert>if use-virtual-mac)

Full Context
configure router router-advertisement interface use-virtual-mac

Description
This command enables sending router advertisement messages using the VRRP virtual MAC address, provided that the virtual router is currently the master.

If the virtual router is not the master, no router advertisement messages are sent.

The **no** form of this command disables sending router advertisement messages.

Default
no use-virtual-mac

Platforms
7705 SAR Gen 2

31.48 user

user

Syntax
[**no**] **user** *user-name*

Context
[Tree] (config>system>security user)

Full Context
configure system security user

Description

This command creates a local user and a context to edit the user configuration.

If a new *user-name* is entered, the user is created. When an existing *user-name* is specified, the user parameters can be edited.

When creating a new user and then entering the **info** command, the system displays a password in the output. This is expected behavior in the hash2 scenario. However, when using that user name, there will be no password required. The user can login to the system and then <ENTER> at the password prompt, the user will be logged in.

Unless an administrator explicitly changes the password, it will be null. The hashed value displayed uses the username and null password field, so when the username is changed, the displayed hashed value will change.

The **no** form of this command deletes the user and all configuration data. Users cannot delete themselves.

Parameters

user-name

Specifies the name of the user up to 32 characters.

Platforms

7705 SAR Gen 2

31.49 user-db

user-db

Syntax

user-db *local-user-db-name* [**create**]

no user-db

Context

[Tree] (config>router>dhcp>server user-db)

Full Context

configure router dhcp local-dhcp-server user-db

Description

This command configures a local user database for authentication.

The **no** form of this command reverts to the default.

Parameters

local-user-db-name

Specifies the name of a user database, up to 32 characters.

create

Keyword used to create the user database. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

user-db

Syntax

user-db *local-user-db-name*

no user-db

Context

[Tree] (config>service>vprn>dhcp6>server user-db)

[Tree] (config>router>dhcp6>server user-db)

Full Context

configure service vprn dhcp6 local-dhcp-server user-db

configure router dhcp6 local-dhcp-server user-db

Description

This command assigns a local user database for DHCP lease data lookup.

The **no** form of this command removes the configuration.

Default

no user-db

Parameters

local-user-db-name

Specifies the local user database name, up to 32 characters.

Platforms

7705 SAR Gen 2

31.50 user-ident

user-ident

Syntax

user-ident *user-ident*

no user-ident

Context

[Tree] (config>service>vprn>dhcp>server user-ident)

[Tree] (config>router>dhcp>server user-ident)

Full Context

configure service vprn dhcp local-dhcp-server user-ident

configure router dhcp local-dhcp-server user-ident

Description

This command configures the user identification method for the DHCPv4 server.

The **no** form of the reverts to the default.

Default

user-ident mac-circuit-id

Parameters

user-ident

Specifies the user identification method

Values **client-id** — Specifies to use the DHCPv4 client identifier as the user identification method.

circuit-id — Specifies to use the circuit identifier of the DHCPv4 client as the user identification method.

mac — Specifies to use the MAC address of the DHCPv4 client as the user identification method.

mac-circuit-id — Specifies to use the MAC address and circuit identifier of the DHCPv4 client as the user identification method.

remote-id — Specifies to use the MAC address of the remote end as the user identification method.

Platforms

7705 SAR Gen 2

user-ident

Syntax

user-ident *user-ident*

no user-ident

Context

[Tree] (config>router>dhcp6>server user-ident)

[Tree] (config>service>vprn>dhcp6>server user-ident)

Full Context

configure router dhcp6 local-dhcp-server user-ident

configure service vprn dhcp6 local-dhcp-server user-ident

Description

This command configures the keys for identification of the DHCPv6 lease being held in the lease-database (for configured period after lease timeout). Subscriber requesting a lease via DHCPv6 that matches an existing lease based on this configured key is handed the matched prefix or address. This allows address and prefix "stickiness" for DHCPv6 assigned prefixes (IA_NA or PD).

The **no** form of the reverts to the default.

Default

user-ident duid

Parameters

user-ident

Specifies the user identification method.

Values **duid** — Specifies the IPv6 DHCP unique identifier from DHCPv6.

interface-id — Specifies the IPv6 interface-id option.

interface-id-link-local — Specifies the interface-id and link-local address.

Platforms

7705 SAR Gen 2

31.51 user-srlg-db

user-srlg-db

Syntax

user-srlg-db [**enable** | **disable**]

Context

[Tree] (config>router>mpls user-srlg-db)

Full Context

configure router mpls user-srlg-db

Description

This command enables the use of CSPF by the user SRLG database. When the MPLS module makes a request to CSPF for the computation of an SRLG secondary path, CSPF will query the local SRLG and compute a path after pruning links that are members of the SRLG IDs of the associated primary path. When MPLS makes a request to CSPF for an FRR bypass or detour path to associate with the primary path, CSPF queries the user SRLG database and computes a path after pruning links that are members of the SRLG IDs of the PLR outgoing interface.

If an interface was not entered into the user SRLG database, it is assumed that it does not have any SRLG membership. CSPF will not query the TE database for IGP advertised interface SRLG information.

The disable keyword disables the use of the user SRLG database. CSPF will then resume queries into the TE database for SRLG membership information. The user SRLG database is maintained.

Default

user-srlg-db disable

Platforms

7705 SAR Gen 2

31.52 user-template

user-template

Syntax

user-template {**tacplus_default** | **radius_default** | **ldap-default**}

Context

[Tree] (config>system>security user-template)

Full Context

configure system security user-template

Description

This command configures default security user template parameters.

Parameters

tacplus_default

Specifies the default TACACS+ user template. All parameters of the tacplus_default template except the "profile" are actively applied to all TACACS+ users if tacplus **use-default-template** is enabled. The **profile** parameter is used for AAA command authorization if TACACS+ authorization is disabled, or if the TACACS+ server does not return a priv-lvl for a user when **use-priv-lvl** is enabled under **tacplus authorization**. See the **tacplus authorization** command for more details.

radius_default

Specifies the default RADIUS user template. The radius_default template is actively applied to a RADIUS user if radius authorization is enabled, radius **use-default-template** is enabled, and no VSAs are returned with the auth-accept from the RADIUS server.

ldap_default

Specifies the default LDAP user template.

Platforms

7705 SAR Gen 2

31.53 usm-community

usm-community

Syntax

usm-community *community-string* [**hash** | **hash2** | **custom**] **group** *group-name* [**src-access-list** *list-name*]

no usm-community *community-string* [**hash** | **hash2** | **custom**]

Context

[Tree] (config>system>security>snmp usm-community)

Full Context

configure system security snmp usm-community

Description

This command is used to associate a community string with an SNMPv3 access group and its view. The access granted with a community string is restricted to the scope of the configured group.

Nokia's SR OS implementation of SNMP uses SNMPv3. In order to implement SNMPv1 and SNMPv2c configurations, several access groups are predefined. In order to implement SNMP with security features (Version 3), security models, security levels, and USM communities must be explicitly configured. Optionally, additional views which specify more specific OIDs (MIB objects in the subtree) can be configured.

The **no** form of this command removes a community string.

Parameters

community-string

Specifies the SNMPv1/SNMPv2c community string to determine the SNMPv3 access permissions to be used. Allowed values are any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (for example, #, $, spaces), the entire string must be enclosed within double quotes.

group

Specifies the group that governs the access rights of this community string. This group must be configured first in the **config**>**system**>**security**>**snmp**> **access group** context. Nokia does not recommend associating a **usm-community** with an SNMP access group that is configured with the **li** (lawful intercept) context.

list-name

Specifies the usm-community to reference a specific src-access-list that will be used to validate the source IP address of all received SNMP requests that use this usm-community. Multiple **community**, **usm-community**, or **vprn snmp community** instances can reference the same **src-access-list**.

Platforms

7705 SAR Gen 2

31.54 util-stats-interval

util-stats-interval

Syntax

util-stats-interval *seconds*

Context

[Tree] (config>port>ethernet util-stats-interval)

Full Context

configure port ethernet util-stats-interval

Description

This command configures the interval used to calculate the utilization statistics.

Port utilization statistics are only available for physical Ethernet ports on a host system. These statistics are not available for the following:

- Ethernet ports on an Ethernet satellite
- PXC ports
- vsm-cca-xp ports

Parameters

seconds

Specifies the size of the interval, in seconds.

Values 30 to 600

Default 300

Platforms

7705 SAR Gen 2

32 v Commands

32.1 v4-routed-override-filter

v4-routed-override-filter

Syntax

v4-routed-override-filter *ip-filter-id*

no v4-routed-override-filter

Context

[Tree] (config>service>ies>if>vpls>egress v4-routed-override-filter)

Full Context

configure service ies interface vpls egress v4-routed-override-filter

Description

This command configures an IPv4 filter ID that is applied to packets egressing the IES R-VPLS interface. The filter overrides existing egress IPv4 filter applied to VPLS service endpoints such as SAPs or SDPs, if configured.

The **no** form of this command removes the IPv4 routed override filter from the egress IES R-VPLS interface. When removed, egress IPv4 packets will use the IPv4 egress filter applied to the VPLS endpoint, if configured.

Default

no v4-routed-override-filter

Parameters

ip-filter-id

Specifies the IP filter ID. This parameter is required when executing the **v4-routed-override-filter** command. The specified filter ID must exist as an IPv4 filter within the system or the override command fails.

Platforms

7705 SAR Gen 2

v4-routed-override-filter

Syntax

v4-routed-override-filter *ip-filter-id*

no v4-routed-override-filter

Context

[Tree] (config>service>ies>if>vpls>ingress v4-routed-override-filter)

Full Context

configure service ies interface vpls ingress v4-routed-override-filter

Description

This command configures an IPv4 filter ID that is applied to routed unicast ingress packets entering the VPLS or I-VPLS service and destined to the R-VPLS interface MAC address. The filter overrides any existing ingress IPv4 filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional and when it is not defined or it is removed. The IPv4 routed packets use any existing ingress IPv4 filter on the VPLS virtual port.

The **no** form of this command removes the IPv4 routed override filter from the ingress IP interface. When removed, the IPv4 ingress routed packets within a VPLS service attached to the IP interface use the IPv4 ingress filter applied to the packets virtual port, when defined.

Parameters

ip-filter-id

Specifies the IP filter ID. This parameter is required when executing the **v4-routed-override-filter** command. The specified filter ID must exist as an IPv4 filter within the system or the override command fails.

Platforms

7705 SAR Gen 2

v4-routed-override-filter

Syntax

v4-routed-override-filter *ip-filter-id*

no v4-routed-override-filter

Context

[Tree] (config>service>vprn>if>vpls>egress v4-routed-override-filter)

Full Context

configure service vprn interface vpls egress v4-routed-override-filter

Description

This command configures an IPv4 filter ID that is applied to packets egressing the VPRN R-VPLS interface. The filter overrides the existing egress IPv4 filter applied to VPLS service endpoints such as SAPs or SDPs, if configured.

The **no** form of this command removes the IPv4 routed override filter from the egress VPRN R-VPLS interface. When removed, egress IPv4 packets will use the IPv4 egress filter applied to VPLS endpoint, if configured.

Parameters

ip-filter-id

Specifies the IP filter ID. This parameter is required when executing the **v4-routed-override-filter** command. The specified filter ID must exist as an IPv4 filter within the system or the override command fails.

Platforms

7705 SAR Gen 2

v4-routed-override-filter

Syntax

v4-routed-override-filter *ip-filter-id*

no v4-routed-override-filter

Context

[Tree] (config>service>vprn>if>vpls>ingress v4-routed-override-filter)

Full Context

configure service vprn interface vpls ingress v4-routed-override-filter

Description

This command configures an IPv4 filter ID that is applied to all ingress packets entering the VPLS service. The filter overrides any existing ingress IPv4 filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional and when it is not defined or it is removed, the IPv4 routed packet's will use the any existing ingress IPv4 filter on the VPLS virtual port.

The **no** form of this command removes the IPv4 routed override filter from the ingress IP interface. When removed, the IPv4 ingress routed packets within a VPLS service attached to the IP interface will use the IPv4 ingress filter applied to the packets virtual port, when defined.

Parameters

ip-filter-id

Specifies the IP filter ID. This parameter is required when executing the **v4-routed-override-filter** command. The specified filter ID must exist as an IPv4 filter within the system or the override command fails.

Platforms

7705 SAR Gen 2

32.2 v6-routed-override-filter

v6-routed-override-filter

Syntax

v6-routed-override-filter *ipv6-filter-id*

no v6-routed-override-filter

Context

[Tree] (config>service>ies>if>vpls>egress v6-routed-override-filter)

Full Context

configure service ies interface vpls egress v6-routed-override-filter

Description

This command configures an IPv6 filter ID that is applied to packets egressing the IES R-VPLS interface. The filter overrides existing egress IPv6 filter applied to VPLS service endpoints such as SAPs or SDPs, if configured.

The **no** form of this command removes the IPv4 routed override filter from the egress IES R-VPLS interface. When removed, egress IPv6 routed packets uses the IPv6 egress filter applied to VPLS endpoint, if configured

Parameters

ipv6-filter-id

Specifies the IPv6 filter ID. This parameter is required when executing the **v6-routed-override-filter** command. The specified filter ID must exist as an IPv6 filter within the system or the override command fails.

Platforms

7705 SAR Gen 2

v6-routed-override-filter

Syntax

v6-routed-override-filter *ipv6-filter-id*

no v6-routed-override-filter

Context

[Tree] (config>service>ies>if>vpls>ingress v6-routed-override-filter)

Full Context

configure service ies interface vpls ingress v6-routed-override-filter

Description

This command configures an IPv6 filter ID that is applied to routed unicast ingress packets entering the VPLS or I-VPLS service and destined to the R-VPLS interface MAC address. The filter overrides any existing ingress IPv6 filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional and when it is not defined or it is removed, the IPv6 routed packets use any existing ingress IPv6 filter on the VPLS virtual port.

The no v6-routed-override-filter command is used to remove the IPv6 routed override filter from the ingress IP interface. When removed, the IPv6 ingress routed packets within a VPLS service attached to the IP interface will use the IPv6 ingress filter applied to the packet's virtual port, when defined.

Default

no v6-routed-override-filter

Parameters

ipv6-filter-id

Specifies the IPv6 filter ID. This parameter is required when executing the **v6-routed-override-filter** command. The specified filter ID must exist as an IPv6 filter within the system or the override command fails.

Platforms

7705 SAR Gen 2

v6-routed-override-filter

Syntax

v6-routed-override-filter *ipv6-filter-id*

no v6-routed-override-filter

Context

[Tree] (config>service>vprn>if>vpls>egress v6-routed-override-filter)

Full Context

configure service vprn interface vpls egress v6-routed-override-filter

Description

This command configures an IPv6 filter ID that is applied to packets egressing the VPRN R-VPLS interface. The filter overrides existing egress IPv6 filter applied to VPLS service endpoints such as SAPs or SDPs, if configured.

The **no** form of the command removes the IPv4 routed override filter from the egress VPRN R-VPLS interface. When removed, egress IPv6 packets will use the IPv6 egress filter applied to the VPLS endpoint, if configured.

Parameters

ipv6-filter-id

Specifies the IPv6 filter ID. This parameter is required when executing the **v6-routed-override-filter** command. The specified filter ID must exist as an IPv6 filter within the system or the override command fails.

Platforms

7705 SAR Gen 2

v6-routed-override-filter

Syntax

v6-routed-override-filter *ipv6-filter-id*

no v6-routed-override-filter

Context

[Tree] (config>service>vprn>if>vpls>ingress v6-routed-override-filter)

Full Context

configure service vprn interface vpls ingress v6-routed-override-filter

Description

This command configures an IPv6 filter ID that is applied to all ingress packets entering the VPLS service. The filter overrides any existing ingress IPv6 filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional and when it is not defined or it is removed, the IPv6 routed packets use the any existing ingress IPv6 filter on the VPLS virtual port.

The **no** form of the command removes the IPv6 routed override filter from the ingress IP interface. When removed, the IPv6 ingress routed packets within a VPLS service attached to the IP interface uses the IPv6 ingress filter applied to the packet's virtual port, when defined.

Parameters

ipv6-filter-id

Specifies the IPv6 filter ID. This parameter is required when executing the **v6-routed-override-filter** command. The specified filter ID must exist as an IPv6 filter within the system or the override command fails.

Platforms

7705 SAR Gen 2

32.3 valid-lifetime

valid-lifetime

Syntax

valid-lifetime infinite

valid-lifetime [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no valid-lifetime

Context

[Tree] (config>service>vprn>dhcp6>local-dhcp-server>pool>prefix valid-lifetime)

[Tree] (config>router>dhcp6>server>pool>prefix valid-lifetime)

Full Context

configure service vprn dhcp6 local-dhcp-server pool prefix valid-lifetime

configure router dhcp6 local-dhcp-server pool prefix valid-lifetime

Description

This command configures the valid lifetime for the IPv6 prefix or address in the option.

The **no** form of this command reverts to the default.

Default

valid-lifetime days 1

Parameters

infinite

Sets the valid lifetime to infinite value.

valid-lifetime

Specifies the valid lifetime

Values	**days** *days*	0 to 49710
	hrs *hours*	0 to 23
	min *minutes*	0 to 59
	sec *seconds*	0 to 5

Platforms

7705 SAR Gen 2

valid-lifetime

Syntax

valid-lifetime {*seconds* | **infinite**}

no valid-lifetime

Context

[Tree] (config>service>vprn>router-advert>if>prefix valid-lifetime)

Full Context

configure service vprn router-advertisement interface prefix valid-lifetime

Description

This command specifies the length of time in seconds that the prefix is valid for the purpose of on-link determination. A value of all one bits (0xffffffff) represents infinity.

The address generated from an invalidated prefix should not appear as the destination or source address of a packet.

Default

valid-lifetime 2592000

Parameters

seconds

Specifies the remaining length of time in seconds that this prefix will continue to be valid.

Values 0 to 429496729

infinite

Specifies that the prefix will always be valid. A value of 4,294,967,295 represents infinity.

Platforms

7705 SAR Gen 2

valid-lifetime

Syntax

valid-lifetime {*seconds* | **infinite**}

no valid-lifetime

Context

[Tree] (config>router>router-advert>if>prefix valid-lifetime)

Full Context

configure router router-advertisement interface prefix valid-lifetime

Description

This command specifies the length of time in seconds that the prefix is valid for the purpose of on-link determination. A value of all one bits (0xffffffff) represents infinity.

The address generated from an invalidated prefix should not appear as the destination or source address of a packet.

Default

valid-lifetime 2592000

Parameters

seconds

Specifies the remaining length of time in seconds that this prefix will continue to be valid.

infinite

Specifies that the prefix will always be valid. A value of 4,294,967,295 represents infinity.

Platforms

7705 SAR Gen 2

32.4 validate

validate

Syntax

validate [*file-url*]

Context

[Tree] (admin>system>license validate)

Full Context

admin system license validate

Description

This command performs a validation on the license file pointed to by the command line argument. A validation ensures that the license is compatible with the current state of the target system but it does not change the existing license. Aspects that can cause a failure in the validation include:

- The license file was created for a different target system. The UUID encoded into the file must match that defined by the specific hardware platform.

- The license file does not include license information for the release of software currently running on the system.
- The current date/time reported to system is outside the validity period encoded in the license.
- The system is currently using a hardware upgrade license that is not included in the new file being validated.

Note:
If the CLM tool is being used for license management, it shall perform the validation and activation and there is no need to enter these commands manually.

Parameters

file-url

Specifies the file URL location to read the license file.

Values local-url, remote-url

Platforms

7705 SAR Gen 2

validate

Syntax

[**no**] **validate**

Context

[Tree] (configure>system>security>profile>netconf>base-op-authorization validate)

Full Context

configure system security profile netconf base-op-authorization validate

Description

This command enables the NETCONF <validate> RPC.

The **no** form of this command disables the RPC.

Default

no validate

Note:
The operation is enabled by default in the built-in system-generated administrative profile.

Platforms

7705 SAR Gen 2

validate

Syntax

validate software-image *file-url*

Context

[Tree] (admin>system>security>secure-boot validate)

Full Context

admin system security secure-boot validate

Description

This command validates the specified software image.

Parameters

file-url

Specifies the URL for the file.

Values [*local-url* | *remote-url*] (up to 180 characters)

where:

- *local-url* — [*cflash-id/*] [*file-path*]
  180 chars max, including *cflash-id*
  directory length 99 chars max each
- *remote-url* — [{ftp://| tftp://} *login*:*pswd*@*remote-locn/*][*file-path*]
  180 chars max
  directory length 99 chars max each
  where: *remote-locn* — [*hostname* | *ipv4-address* | *ipv6-address*]

ipv4-address	a.b.c.d
ipv6-address	x:x:x:x:x:x:x:x[-interface]
	x:x:x:x:x:x:d.d.d.d[-interface]
	x - [0..FFFF]H
	d - [0..255]D
	interface - 32 chars max, for link local addresses
cflash-id	cf1:| cf1-A:| cf1-B:| cf2:| cf2-A:| cf2-B:| cf3:| cf3-A:| cf3-B:

Platforms

7705 SAR Gen 2

32.5 validate-next-hop

validate-next-hop

Syntax

[**no**] **validate-next-hop**

Context

[Tree] (config>service>vprn>static-route-entry>next-hop validate-next-hop)

Full Context

configure service vprn static-route-entry next-hop validate-next-hop

Description

This optional command tracks the state of the next hop in the IPv4 ARP cache or IPv6 Neighbor Cache. When the next hop is not reachable and is removed from the ARP or Neighbor Cache, the next hop will no longer be considered valid and the associated static route state removed from the active route-table.

When the next hop is reachable again and present in the ARP/Neighbor Cache, the static route is considered valid and is subject to being placed into the active route-table.

Default

no validate-next-hop

Platforms

7705 SAR Gen 2

validate-next-hop

Syntax

[**no**] **validate-next-hop**

Context

[Tree] (config>router>static-route-entry>next-hop validate-next-hop)

Full Context

configure router static-route-entry next-hop validate-next-hop

Description

This optional command tracks the state of the next-hop in the IPv4 ARP cache or IPv6 Neighbor Cache. When the next-hop is not reachable and is removed from the ARP or Neighbor Cache, the next-hop will no longer be considered valid and the associated static-route state removed from the active route-table.

When the next-hop is reachable again and present in the ARP/Neighbor Cache, the static route is considered valid and is subject to being placed into the active route-table.

Default

no validate-next-hop

Platforms

7705 SAR Gen 2

32.6 vc-label

vc-label

Syntax

vc-label *egress-vc-label*

no vc-label [*egress-vc-label*]

Context

[Tree] (config>service>vprn>if>spoke-sdp>egress vc-label)

Full Context

configure service vprn interface spoke-sdp egress vc-label

Description

This command configures the egress VC label.

Parameters

vc-label

A VC egress value that indicates a specific connection.

Values 16 to 1048575

Platforms

7705 SAR Gen 2

vc-label

Syntax

vc-label *ingress-vc-label*

no vc-label [*ingress-vc-label*]

Context

[Tree] (config>service>vprn>if>spoke-sdp>ingress vc-label)

Full Context

configure service vprn interface spoke-sdp ingress vc-label

Description

This command configures the ingress VC label.

Parameters

vc-label

A VC ingress value that indicates a specific connection.

Values 2048 to 18431

Platforms

7705 SAR Gen 2

vc-label

Syntax

vc-label *egress-vc-label*

no vc-label [*egress-vc-label*]

Context

[Tree] (config>service>vpls>mesh-sdp>egress vc-label)

[Tree] (config>service>vpls>spoke-sdp>egress vc-label)

[Tree] (config>service>ies>if>spoke-sdp>egress vc-label)

Full Context

configure service vpls mesh-sdp egress vc-label

configure service vpls spoke-sdp egress vc-label

configure service ies interface spoke-sdp egress vc-label

Description

This command configures the static MPLS VC label used by this device to send packets to the far-end device in this service via this SDP.

Parameters

egress-vc-label

Specifies a VC egress value that indicates a specific connection.

Values 16 to 1048575

Platforms

7705 SAR Gen 2

vc-label

Syntax

vc-label *ingress-vc-label*

no vc-label [*ingress-vc-label*]

Context

[Tree] (config>service>vpls>mesh-sdp>ingress vc-label)

[Tree] (config>service>vpls>spoke-sdp>ingress vc-label)

[Tree] (config>service>ies>if>spoke-sdp>ingress vc-label)

Full Context

configure service vpls mesh-sdp ingress vc-label

configure service vpls spoke-sdp ingress vc-label

configure service ies interface spoke-sdp ingress vc-label

Description

This command configures the static MPLS VC label used by the far-end device to send packets to this device in this service via this SDP.

Parameters

ingress-vc-label

A VC ingress value that indicates a specific connection.

Values 2048 to 18431

Platforms

7705 SAR Gen 2

vc-label

Syntax

vc-label *egress-vc-label*

no vc-label [*egress-vc-label*]

Context

[Tree] (config>mirror>mirror-dest>spoke-sdp>egress vc-label)

[Tree] (config>mirror>mirror-dest>remote-src>spoke-sdp>egress vc-label)

Full Context

configure mirror mirror-dest spoke-sdp egress vc-label

configure mirror mirror-dest remote-source spoke-sdp egress vc-label

Description

This command configures the spoke SDP egress VC label.

The **no** form of this command removes the egress VC label value from the configuration.

Parameters

egress-vc-label

Specifies a VC egress value that indicates a specific connection.

Values 16 to 1048575

Platforms

7705 SAR Gen 2

vc-label

Syntax

vc-label *ingress-vc-label*

no vc-label [*ingress-vc-label*]

Context

[Tree] (config>mirror>mirror-dest>remote-src>spoke-sdp>ingress vc-label)

[Tree] (config>service>vprn>ipmirrorif>spoke-sdp>ingress vc-label)

Full Context

configure mirror mirror-dest remote-source spoke-sdp ingress vc-label

configure service vprn ip-mirror-interface spoke-sdp ingress vc-label

Description

This command configures the spoke SDP ingress VC label.

Parameters

vc-label

Specifies the VC ingress value that indicates a specific connection.

Values 32 to 18431

Platforms

7705 SAR Gen 2

32.7 vc-type

vc-type

Syntax

vc-type {**ether** | **vlan**}

Context

[Tree] (config>service>pw-template vc-type)

Full Context

configure service pw-template vc-type

Description

This command overrides the default VC type signaled for the binding to the far end SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the bindings VC type causes the binding to signal the new VC type to the far end when signaling is enabled. VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

- The VC type value for Ethernet is 0x0005.
- The VC type value for an Ethernet VLAN is 0x0004.

Parameters

ether

Defines the VC type as Ethernet. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke SDP binding. (hex 5)

vlan

Defines the VC type as VLAN. The top VLAN tag, if a VLAN tag is present, is stripped from traffic received on the pseudowire, and a vlan-tag is inserted when forwarding into the pseudowire. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings.

Note:
The system expects a symmetrical configuration with its peer, specifically it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. As some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. Consequently, protocol extractions will not necessarily function for asymmetrical configurations as they would with a symmetrical configurations resulting in an unexpected operation.

Platforms

7705 SAR Gen 2

32.8 ve-id

ve-id

Syntax

ve-id *value*

no ve-id

Context

[Tree] (config>service>epipe>bgp-vpws>ve-name ve-id)

[Tree] (config>service>epipe>bgp-vpws>remote-ve-name ve-id)

Full Context

configure service epipe bgp-vpws ve-name ve-id

configure service epipe bgp-vpws remote-ve-name ve-id

Description

This command configures a ve-id for either the local VPWS instance when configured under the ve-name, or for the remote VPWS instance when configured under the remote-ve-name.

A single ve-id can be configured per ve-name or remote-ve-name. The ve-id can be changed without shutting down the VPWS instance. When the ve-name ve-id changes, BGP withdraws the previously advertised route and sends a route-refresh to all the peers which would result in reception of all the remote routes again. The old PWs are removed and new ones are instantiated for the new ve-id value.

When the remote-ve-name ve-id changes, BGP withdraws the previously advertised route and send a new update matching the new ve-id. The old pseudowires are removed and new ones are instantiated for the new ve-id value.

NLRIs received whose advertised ve-id does not match the list of ve-ids configured under the remote ve-id will not have a spoke SDP binding auto-created but will remain in the BGP routing table but not in the Layer 2 route table. A change in the locally configured ve-ids may result in auto-sdp-bindings either being deleted or created, based on the new matching results.

Each ve-id configured within a service must be unique.

The **no** form of this command removes the configured ve-id. It can be used just when the BGP VPWS status is shutdown. The **no shutdown** command cannot be used if there is no ve-id configured.

Default

no ve-id

Parameters

value

A two bytes identifier that represents the local or remote VPWS instance and is advertised through the BGP NLRI.

Values 1 to 65535

Platforms

7705 SAR Gen 2

ve-id

Syntax

ve-id *ve-id-value*

no ve-id

Context

[Tree] (config>service>vpls>bgp-vpls>ve-name ve-id)

Full Context

configure service vpls bgp-vpls ve-name ve-id

Description

This command configures a ve-id. Just one ve-id can be configured per BGP VPLS instance. The VE-ID can be changed without shutting down the VPLS Instance. When the VE-ID changes, BGP is withdrawing its own previously advertised routes and sending a route-refresh to all the peers which would result in reception of all the remote routes again. The old pseudowires are removed and new ones are instantiated for the new VE-ID value.

The **no** form of this command removes the configured ve-id. It can be used just when the BGP VPLS status is shutdown. The **no shutdown** command cannot be used if there is no ve-id configured.

Default

no ve-id

Parameters

value

Specifies a two-byte identifier that represents the local instance in a VPLS and is advertised through the BGP NLRI. Must be lower or equal with the max-ve-id.

Values 1 to 65535

Platforms

7705 SAR Gen 2

32.9 ve-name

ve-name

Syntax

[**no**] **ve-name** *name*

Context

[Tree] (config>service>epipe>bgp-vpws ve-name)

Full Context

configure service epipe bgp-vpws ve-name

Description

This command configures the name of the local VPWS instance in this service.

The **no** form of this command removes the ve-name.

Parameters

name

Specifies a site name up to 32 characters in length.

Platforms

7705 SAR Gen 2

ve-name

Syntax

ve-name *name*

no ve-name

Context

[Tree] (config>service>vpls>bgp-vpls ve-name)

Full Context

configure service vpls bgp-vpls ve-name

Description

This command creates or edits a ve-name. Just one ve-name can be created per BGP VPLS instance.

The **no** form of this command removes the configured ve-name from the bgp vpls node. It can be used only when the BGP VPLS status is shutdown. The **no shutdown** command cannot be used if there is no ve-name configured.

Default

no ve-name

Parameters

name

Specifies the A character string to identify the VPLS Edge instance up to 32 characters in length

Platforms

7705 SAR Gen 2

32.10 vendor-id

vendor-id

Syntax

vendor-id *vendor-id*

no vendor-id

Context

[Tree] (config>system>ned>profile vendor-id)

Full Context

configure system network-element-discovery profile vendor-id

Description

This command configures the vendor ID to be advertised.

The **no** form of this command reverts to the default value.

Default

vendor-id "Nokia"

Parameters

vendor-id

Specifies the vendor ID to be advertised with the profile, up to 255 characters.

Platforms

7705 SAR Gen 2

32.11 vendor-specific-option

vendor-specific-option

Syntax

[**no**] **vendor-specific-option**

Context

[Tree] (config>service>vpls>sap>dhcp>option vendor-specific-option)

[Tree] (config>service>vprn>if>dhcp>option vendor-specific-option)

[Tree] (config>service>ies>if>dhcp>option vendor-specific-option)

Full Context

configure service vpls sap dhcp option vendor-specific-option

configure service vprn interface dhcp option vendor-specific-option

configure service ies interface dhcp option vendor-specific-option

Description

This command enables the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command reverts to the default.

Platforms

7705 SAR Gen 2

vendor-specific-option

Syntax

[**no**] **vendor-specific-option**

Context

[Tree] (config>router>if>dhcp>option vendor-specific-option)

Full Context

configure router interface dhcp option vendor-specific-option

Description

This command configures the Nokia vendor specific suboption of the DHCP relay packet.

Platforms

7705 SAR Gen 2

32.12 version

version

Syntax

version *version*

no version

Context

[Tree] (config>service>vpls>sap>igmp-snooping version)

[Tree] (config>service>vpls>spoke-sdp>mld-snooping version)

[Tree] (config>service>vpls>spoke-sdp>igmp-snooping version)

[Tree] (config>service>vpls>mesh-sdp>mld-snooping version)

[Tree] (config>service>vpls>mesh-sdp>igmp-snooping version)

[Tree] (config>service>vpls>sap>mld-snooping version)

Full Context

configure service vpls sap igmp-snooping version

configure service vpls spoke-sdp mld-snooping version

configure service vpls spoke-sdp igmp-snooping version

configure service vpls mesh-sdp mld-snooping version

configure service vpls mesh-sdp igmp-snooping version

configure service vpls sap mld-snooping version

Description

This command specifies the version of IGMP or MLD which is running on this SAP or SDP. This object can be used to configure a router capable of running either value. For IGMP or MLD to function correctly, all routers on a LAN must be configured to run the same version of IGMP or MLD on that LAN.

When the **send-query** command is configured, all type of queries generate ourselves are of the configured **version**. If a report of a version higher than the configured version is received, the report gets dropped and a new "wrong version" counter is incremented.

If the **send-query** command is not configured, the **version** command has no effect. The version used on that SAP or SDP is the version of the querier. This implies that, for example, when there is a v2 querier, a v3 group or group-source specific query when a host wants to leave a certain group will never be sent.

Parameters

version

Specifies the IGMP or MLD version

Values 1, 2, 3

Platforms

7705 SAR Gen 2

version

Syntax

version *version*

no version

Context

[Tree] (config>service>vprn>igmp>if version)

Full Context

configure service vprn igmp interface version

Description

This command specifies the IGMP version. If routers run different versions of IGMP, they will negotiate the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version. For IGMP to function correctly, all routers on a LAN should be configured to run the same version of IGMP on that LAN.

For IGMPv3, a multicast router that is also a group member performs both parts of IGMPv3, receiving and responding to its own IGMP message transmissions as well as those of its neighbors.

Default

version 3

Parameters

version

Specifies the IGMP version number.

Values 1, 2, 3

Platforms

7705 SAR Gen 2

version

Syntax

version *version*

no version

Context

[Tree] (config>service>vprn>mld>if version)

Full Context

configure service vprn mld interface version

Description

This command specifies the MLD version. If routers run different versions, they will negotiate the lowest common version of MLD that is supported by hosts on their subnet and operate in that version. For MLD to function correctly, all routers on a LAN should be configured to run the same version of MLD on that LAN.

Default

version 2

Parameters

version

Specifies the MLD version number.

Values 1, 2

Platforms

7705 SAR Gen 2

version

Syntax

version *version*

no version

Context

[Tree] (config>router>igmp>if version)

Full Context

configure router igmp interface version

Description

This command specifies the IGMP version. If routers run different versions of IGMP, they will negotiate the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version. For IGMP to function correctly, all routers on a LAN should be configured to run the same version of IGMP on that LAN.

For IGMPv3, a multicast router that is also a group member performs both parts of IGMPv3, receiving and responding to its own IGMP message transmissions as well as those of its neighbors.

Default

version 3

Parameters

version

Specifies the IGMP version number.

Values 1, 2, 3

Platforms

7705 SAR Gen 2

version

Syntax

version *version*

no version

Context

[Tree] (config>router>mld>interface version)

Full Context

configure router mld interface version

Description

This command specifies the MLD version. If routers run different versions of MLD, they will negotiate the lowest common version of MLD that is supported by hosts on their subnet and operate in that version. For MLD to function correctly, all routers on a LAN should be configured to run the same version of MLD on that LAN.

Default

version 2

Parameters

version

Specifies the MLD version number.

Values 1, 2

Platforms

7705 SAR Gen 2

version

Syntax

version *version*

no version

Context

[Tree] (config>service>pw-template>igmp-snooping version)

Full Context

configure service pw-template igmp-snooping version

Description

This command specifies the version of IGMP. This object can be used to configure a router capable of running either value. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN.

When the **send-query** command is configured, all type of queries generated are of the configured **version**. If a report of a version higher than the configured version is received, the report gets dropped and a new "wrong version" counter is incremented.

If the **send-query** command is not configured, the **version** command has no effect. The version used on that SAP or SDP is the version of the querier. This implies that, for example, when there is a v2 querier, a v3 group or group-source specific query when a host wants to leave a certain group is never sent.

Default

version 3

Parameters

version

Specifies the IGMP version.

Values 1, 2, 3

Platforms

7705 SAR Gen 2

version

Syntax

version *file-url* [**check**]

Context

[Tree] (file version)

Full Context

file version

Description

This command displays the version of an SR OS *.tim image file.

Parameters

file-url

Specifies the file name of the target file.

Values

local-url	[*cflash-id*/][*file-path*]
	up to 200 characters, including *cflash-id*
	directory length up to 99 characters each
remote-url	[{ftp://| tftp://}*login*:*pswd*@*remote-locn*/ [*file-path*]
	up to 247 characters
	directory length 199 characters each
remote-locn	[*hostname* | *ipv4-address* | *ipv6-address*]
ipv4-address	*a.b.c.d*
ipv6-address	*x:x:x:x:x:x:x:x*[-*interface*]

	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D
	interface - up to 32 characters, for link local addresses
cflash-id	cf1:, cf1-A:, cf1-B:

check

Validates the SR OS *.tim image file.

Platforms

7705 SAR Gen 2

Output

The following output is an example of SR OS version information.

Output Example

```
A:Redundancy>file cf3:\ # version  ftp://test:1234@192.0.2.79/usr/global/images/6.1/R4/cpm.tim
TiMOS-C-6.1.R4 for 7750
Thu Oct 30 14:21:09 PDT 2018 by builder in /relx.1/b1/Rx/panos/main
A:Redundancy>file cf3:\ # version check ftp://test:1234@192.0.2.79/usr/global/
images/6.1/R4/cpm.tim
TiMOS-C-6.1.R4 for 7750
Thu Oct 30 14:21:09 PDT 2018 by builder in /relx.1/b1/Rx/panos/main
Validation successful
A:Redundancy>file cf3:\ #
```

32.13 vi

vi

Syntax

vi *local-url*

Context

[Tree] (file vi)

Full Context

file vi

Description

Edit files with the text editor. For more information, refer to "Text Editor" in the *7705 SAR Gen 2 Basic System Configuration Guide*.

Parameters

local-url

Specifies the local source file or directory.

Values *[cflash-id/]file-path*

cflash-id: cf1:, cf2:, cf3:

Platforms

7705 SAR Gen 2

32.14 view

view

Syntax

view [*line*]

Context

[Tree] (candidate view)

Full Context

candidate view

Description

This command displays the candidate configuration along with line numbers that can be used for editing the candidate configuration.

Parameters

line

Displays the candidate configuration starting at the point indicated by the following options (the display is not limited to the current CLI context/branch).

Values line, offset, **first**, **edit-point**, **last**

line	absolute line number
offset	relative line number to current edit point. Prefixed with '+' or '-'

first	keyword - first line
edit-point	keyword - current edit point
last	keyword - last line that is not 'exit'

Platforms

7705 SAR Gen 2

view

Syntax

view [*checkpoint-id* | **rescue** | **latest-rb**]

Context

[Tree] (admin>rollback view)

Full Context

admin rollback view

Description

This command displays the checkpoint.

Parameters

latest-rb

Specifies the most recently created rollback checkpoint (corresponds to the file-url.rb rollback checkpoint file).

checkpoint-id

Indicates rollback checkpoint file to be viewed. Checkpoint-id of 1 corresponds to the file-url.rb.1 rollback checkpoint file. The higher the id, the older the checkpoint. Max is the highest rollback checkpoint supported or configured.

Values 1 to 9

rescue

Displays the rescue configuration.

Platforms

7705 SAR Gen 2

view

Syntax

view {**bootup-cfg** | **active-cfg** | **candidate-cfg** | **latest-rb**| *checkpoint-id* | **rescue**}

Context

[Tree] (admin view)

Full Context

admin view

Description

The context to configure administrative system viewing parameters. Only authorized users can execute the commands in the **admin** context.

Parameters

bootup-cfg

Specifies the bootup configuration.

active-cfg

Specifies current running configuration.

candidate-cfg

Specifies candidate configuration.

latest-rb

Specifies the latest configuration.

checkpoint-id

Specifies a specific checkpoint file configuration.

Values 1 to 9

rescue

Specifies a rescue checkpoint configuration.

Platforms

7705 SAR Gen 2

view

Syntax

view *view-name* **subtree** *oid-value*

no view *view-name* [**subtree** *oid-value*]

Context

[Tree] (config>system>security>snmp view)

Full Context

configure system security snmp view

Description

This command configures a view. Views control the accessibility of a MIB object within the configured MIB view and subtree. Object identifiers (OIDs) uniquely identify MIB objects in the subtree. OIDs are organized hierarchically with specific values assigned by different organizations.

Once the subtree (OID) is identified, a mask can be created to select the portions of the subtree to be included or excluded for access using this particular view. See the **mask** command.

The view(s) configured with this command can subsequently be used in read, write, and notify commands which are used to assign specific access group permissions to created views and assigned to particular access groups.

Multiple subtrees can be added or removed from a view name to tailor a view to the requirements of the user access group.

A subtree statement matches (covers) any OID that is a descendant of the specified OID value. For example, the subtree 1.3.6.1 matches 1.3.6.1.x (for any value of x), 1.3.6.1.x.y (for any values of x & y), and so on.

Subtrees that are not covered by **view** statements are not accessible in the view.

Per RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP),* each MIB view is defined by two sets of view subtrees, the included view subtrees, and the excluded view subtrees (see the **included** and **excluded** parameters of the **mask** command). Every such view subtree, both the included and the excluded ones, are defined in this table. To determine if a particular object instance is in a particular MIB view, compare the object instance's OID with each of the MIB view's active entries in this table. If none match, then the object instance is not in the MIB view. If one or more match, then the object instance is included in, or excluded from, the MIB view according to the value of vacmViewTreeFamilyType in the entry whose value of vacmViewTreeFamilySubtree has the most sub-identifiers.

The **no view** *view-name* command removes a view and all subtrees.

The **no view** *view-name* **subtree** *oid-value* removes a sub-tree from the view name.

Parameters

view-name

Specifies a view name, up to 32 characters.

oid-value

Specifies the object identifier (OID) value for the *view-name*. This value, for example, 1.3.6.1.6.3.11.2.1, combined with the mask and include and exclude statements, configures the access available in the view.

It is possible to have a view with different subtrees with their own masks and include and exclude statements. This allows for customizing visibility and write capabilities to specific user requirements**.**

Platforms

7705 SAR Gen 2

32.15 virtual-link

virtual-link

Syntax

[**no**] **virtual-link** *router-id* **transit-area** *area-id*

Context

[Tree] (config>service>vprn>ospf>area virtual-link)

[Tree] (config>service>vprn>ospf3>area virtual-link)

Full Context

configure service vprn ospf area virtual-link

configure service vprn ospf3 area virtual-link

Description

This command configures a virtual link to connect area border routers to the backbone via a virtual link.

The *router-id* specified in this command must be associated with the virtual neighbor. The transit area cannot be a stub area or a Not So Stubby Area (NSSA).

The **no** form of this command deletes the virtual link.

Default

No virtual link is defined.

Parameters

router-id

The router ID of the virtual neighbor in IP address dotted decimal notation.

transit-area ***area-id***

The area-id specified identifies the transit area that links the backbone area with the area that has no physical connection with the backbone.

Platforms

7705 SAR Gen 2

virtual-link

Syntax

[**no**] **virtual-link** *router-id* **transit-area** *area-id*

Context

[Tree] (config>router>ospf3>area virtual-link)

[Tree] (config>router>ospf>area virtual-link)

Full Context

configure router ospf3 area virtual-link

configure router ospf area virtual-link

Description

This command configures a virtual link to connect area border routers to the backbone via a virtual link.

The *router-id* specified in this command must be associated with the virtual neighbor. The transit area cannot be a stub area or a Not So Stubby Area (NSSA).

The **no** form of this command deletes the virtual link.

By default, no virtual link is defined.

Default

no virtual-link

Parameters

router-id

Specifies the router ID of the virtual neighbor in IP address dotted-decimal notation.

area-id

Specifies the area-id that identifies the transit area that links the backbone area with the area that has no physical connection with the backbone.

Platforms

7705 SAR Gen 2

32.16 virtual-neighbor

virtual-neighbor

Syntax

virtual-neighbor [*router-id*]

no virtual-neighbor

Context

[Tree] (debug>router>ospf virtual-neighbor)

[Tree] (debug>router>ospf3 virtual-neighbor)

Full Context

debug router ospf virtual-neighbor

debug router ospf3 virtual-neighbor

Description

This command enables debugging for an OSPF virtual neighbor.

Parameters

router-id

Specifies the router ID of the virtual neighbor.

Platforms

7705 SAR Gen 2

32.17 vlan

vlan

Syntax

vlan [*vlan-encap*]

Context

[Tree] (config>redundancy>mc>peer>mcr>l3ring>node>cv vlan)

Full Context

configure redundancy multi-chassis peer mc-ring l3-ring ring-node connectivity-verify vlan

Description

This command specifies the VLAN tag of the SAP used for ring-node connectivity verification of this ring node. It is only meaningful if the value of is not zero.

The **no** form of this command reverts to the default.

Parameters

vlan-encap

Specifies the node cc VLAN IP.

Platforms

7705 SAR Gen 2

32.18 vlan-aware-bundle

vlan-aware-bundle

Syntax

vlan-aware-bundle *name* [**eth-tag** *value*]

no vlan-aware-bundle

Context

[Tree] (config>service>vpls>bgp-evpn vlan-aware-bundle)

Full Context

configure service vpls bgp-evpn vlan-aware-bundle

Description

This command configures a name that is used to group a bundle of VPLS services (Broadcast Domains) that are part of the same VLAN-aware bundle instance. This name is optional and allows the user to execute **show** commands that are relevant to all the Broadcast Domains in a VLAN-aware bundle service group.

The optional Ethernet Tag ID can be encoded in the EVPN routes for control-plane interoperability mode with VLAN-aware bundle services. The configuration of a non-default value requires the previous configuration of a VLAN-aware bundle name on the service.

When the Ethernet Tag ID is set to a non-zero value, the EVPN routes advertised for the VPLS service are advertised with this value into the Ethernet Tag ID field of the routes.

On reception of EVPN routes with non-zero Ethernet Tag ID, BGP imports the routes based on the import route target as usual. However, the system checks the received Ethernet Tag ID field and processes only routes whose Ethernet Tag ID matches the local VLAN-aware bundle Ethernet Tag value.

The **no** form of this command removes the configuration.

Parameters

name

Specifies the VLAN-aware bundle name, up to 32 characters.

value

Specifies the Ethernet Tag ID.

Values 1 to 16777215

Platforms

7705 SAR Gen 2

32.19 vlan-range

vlan-range

Syntax

[**no**] **vlan-range** [*vlan-range*]

Context

[Tree] (config>service>vpls>stp>mst-instance vlan-range)

Full Context

configure service vpls stp mst-instance vlan-range

Description

This command specifies a range of VLANs associated with a certain mst-instance. This range applies to all SAPs of the M-VPLS.

Every VLAN range that is not assigned within any of the created **config>service>vpls>stp mst-instance** is automatically assigned to mst-instance 0. This instance is automatically maintained by the software and cannot be modified. Changing the VLAN range value can be performed only when the specified mst-instance is shutdown.

The **no** form of this command removes the **vlan-range** from the specified **config>service>vpls>stp mst-instance**.

Parameters

vlan-range

The first VLAN range specifies the left-bound (i.e., minimum value) of a range of VLANs that are associated with the M-VPLS SAP. This value must be smaller than (or equal to) the second VLAN range value. The second VLAN range specifies the right-bound (i.e., maximum value) of a range of VLANs that are associated with the M-VPLS SAP.

Values 1 to 4094 to 1 to 4094

Platforms

7705 SAR Gen 2

vlan-range

Syntax

vlan-range *from* [**to** *to*]

no vlan-range *from*

Context

[Tree] (config>connection-profile-vlan vlan-range)

Full Context

configure connection-profile-vlan vlan-range

Description

This command allows the user to configure different ranges in the connection-profile-vlan. The ranges have the following characteristics:

- Ranges can contain a single VID or start-and-end values. When the *to-vid* is not specified, the end vid value is the same as the start vid value.
- On the fly addition/removal of ranges is allowed.
- When removing an entry, the **no vlan-range** *vid* **to** *vid* must be configured by the user.
- Multiple ranges are allowed under the same connection-profile-vlan. No VLAN values should overlap within the same connection-profile-vlan.
- The index for connection-profile and connection-profile-vlan must be unique between the two. For example, if **connection-profile 100** is present, then **connection-profile-vlan 100** is disallowed.

Each connection-profile-vlan must be explicitly configured.

Parameters

from

Specifies the beginning of the **vlan-range** associated to the **connection-profile-vlan**.

Values 1 to 4094

to

Specifies the end of the **vlan-range** associated to the **connection-profile-vlan**. If not specified, the **vlan-range** is comprised of only the *from* VLAN ID.

Values 1 to 4094

Platforms

7705 SAR Gen 2

32.20 vlan-vc-etype

vlan-vc-etype

Syntax

vlan-vc-etype *ethernet-type*

no vlan-vc-etype [*ethernet-type*]

Context

[Tree] (config>service>sdp vlan-vc-etype)

Full Context

configure service sdp vlan-vc-etype

Description

This command configures the VLAN VC EtherType.

The **no** form of this command returns the value to the default.

Default

no vlan-vc-etype

Parameters

ethernet-type

Specifies a valid VLAN etype identifier.

Values 0x0600 to 0xffff

Platforms

7705 SAR Gen 2

32.21 vlan-vc-tag

vlan-vc-tag

Syntax

vlan-vc-tag *vlan-id*

no vlan-vc-tag [*vlan-id*]

Context

[Tree] (config>service>vpls>spoke-sdp vlan-vc-tag)

[Tree] (config>service>vpls>mesh-sdp vlan-vc-tag)

Full Context

configure service vpls spoke-sdp vlan-vc-tag

configure service vpls mesh-sdp vlan-vc-tag

Description

This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.

When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The **no** form of this command disables the command.

Default

no vlan-vc-tag

Parameters

vlan-id

Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

Values 0 to 4094

Platforms

7705 SAR Gen 2

vlan-vc-tag

Syntax

vlan-vc-tag *tag*

no vlan-vc-tag *tag*

Context

[Tree] (config>service>epipe>spoke-sdp vlan-vc-tag)

Full Context

configure service epipe spoke-sdp vlan-vc-tag

Description

This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.

When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The **no** form of this command disables the command.

Default

no vlan-vc-tag

Parameters

tag

Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

Values 0 to 4094

Platforms

7705 SAR Gen 2

vlan-vc-tag

Syntax

vlan-vc-tag *vlan-id*

no vlan-vc-tag

Context

[Tree] (config>service>pw-template vlan-vc-tag)

Full Context

configure service pw-template vlan-vc-tag

Description

This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.

When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The **no** form of this command disables the command.

Default

no vlan-vc-tag

Parameters

vlan-id

Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

Values 0 to 4094

Platforms

7705 SAR Gen 2

32.22 vpls

vpls

Syntax

vpls *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**m-vpls**] [**b-vpls** | **i-vpls**] [**etree**] [**name** *name*]

no vpls *service-id*

Context

[Tree] (config>service vpls)

Full Context

configure service vpls

Description

This command creates or edits a Virtual Private LAN Services (VPLS) instance. The **vpls** command is used to create or maintain a VPLS service. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites together acting like a zero-hop, Layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **create** keyword must be specified if the **create** command is enabled in the **environment** context. When creating a service, you must enter the **customer** keyword and specify a *customer-id* to associate the service with a customer. The *customer-id* must already exist, having been created using the **customer** command in the **service** context. The *customer-id* must already exist having been created using the **customer** command in the service context. Once a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and re-created with a new customer association.

Once a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

More than one VPLS service may be created for a single customer ID.

By default, no VPLS instances exist until they are explicitly created.

The **no** form of this command deletes the VPLS service instance with the specified *service-id*. The service cannot be deleted until all SAPs and SDPs defined within the service ID have been shut down and deleted, and the service has been shut down.

Parameters

service-id

Specifies unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values *service-id*: 1 to 2147483647

svc-name: 64 characters maximum

customer *customer-id*

Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn vpn-id

Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number

Values 1 to 2147483647

Default null (0)

create

Keyword used to create the service ID. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

m-vpls

Specifies a management VPLS

e-tree

Specifies a VPLS service as an E-Tree VPLS. E-Tree VPLS services have root and leaf attachment circuit (AC) and root leaf tag SAPs/SDP bindings for E-Tree interconnection. The access root AC SAP behaves as a SAP in non-E-tree VPLS services. The leaf AC SAP communicates only with root-connected services. Leaf and root AC SAPs behave externally the same as SAPs in non-E-Tree VPLS services.

The root AC SDP bind behaves as an SDP bind in non-E-tree VPLS services. The leaf AC SDP bind communicates only with root-connected services.

In the E-Tree VPLS, the root AC SAP/SDP bindings can communicate with other root and leaf AC SAP/SDP bind services locally and remotely. Root-originated traffic is marked internally with a root indication and the root is tagged externally on tag SAP/SDP binds. The leaf AC SAP/SDP bindings can communicate with other root SAP/SDP bindings locally and remotely. Leaf-originated traffic is marked internally with a leaf indication and tagged externally on leaf tag SAP/SDP bindings.

Any number of root or leaf AC SAPs can be used, up to the configured SAP limits in the E-Tree VPLS.

Network-side root leaf tag SAPs use additional SAP resources. These tag SAPs used two tags; one for root and one for leaf. Network-side tag SDPs use a hard coded tag of 1 for root and 2 for leaf. AC SDP bindings are designated as root or leaf SDP bindings but carry no tags marking traffic on the egress frames.

The E-Tree SAP type must be specified when the SAP is created. To change the SAP type, the SAP must be removed and recreated.

b-vpls | i-vpls

Creates a backbone-vpls or ISID-vpls

name name

Configures an optional service name identifier, up to 64 characters, to a given service. This service name can then be used in configuration references, display, and show commands throughout the system. A defined service name can help the service provider or administrator to identify and manage services within the SR OS platforms.

To create a service, you must assign a service ID; however, after it is created, either the service ID or the service name can be used to identify and reference a service.

If a name is not specified at creation time, then SR OS assigns a string version of the *service-id* as the name.

Platforms

7705 SAR Gen 2

vpls

Syntax

vpls *service-name*

no vpls

Context

[Tree] (config>service>ies>if vpls)

Full Context

configure service ies interface vpls

Description

The **vpls** command, within the IP interface context, is used to bind the IP interface to the specified service name (VPLS or I-VPLS).

The system does not attempt to resolve the service name provided until the IP interface is placed into the administratively up state (**no shutdown**). Once the IP interface is administratively up, the system will scan the available VPLS services that have the **allow-ip-int-bind** flag set for a VPLS service associated with the name. If the service name is bound to the service name when the IP interface is already in the administratively up state, the system will immediately attempt to resolve the given name.

If a VPLS service is found associated with the name and with the **allow-ip-int-bind** flag set, the IP interface is attached to the VPLS service allowing routing to and from the service virtual ports once the IP interface is operational.

A VPLS service associated with the specified name that does not have the **allow-ip-int-bind** flag set or a non-VPLS service associated with the name is ignored and will not be attached to the IP interface.

If the service name is applied to a VPLS service after the service name is bound to an IP interface and the VPLS service **allow-ip-int-bind** flag is set at the time the name is applied, the VPLS service is automatically resolved to the IP interface if the interface is administratively up or when the interface is placed in the administratively up state.

If the service name is applied to a VPLS service without the **allow-ip-int-bind** flag set, the system will not attempt to resolve the applied service name to an existing IP interface bound to the name. To rectify this condition, the flag must first be set and then the IP interface must enter or reenter the administratively up state.

While the specified service name may be assigned to only one service context in the system, it is possible to bind the same service name to more than one IP interface. If two or more IP interfaces are bound to the same service name, the first IP interface to enter the administratively up state (if currently administratively down) or to reenter the administratively up state (if currently administratively up) when a VPLS service is configured with the name and has the **allow-ip-int-bind** flag set is attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

Once an IP interface is attached to a VPLS service, the name associated with the service cannot be removed or changed until the IP interface name binding is removed. Also, the **allow-ip-int-bind** flag cannot be removed until the attached IP interface is unbound from the service name.

Unbinding the service name from the IP interface causes the IP interface to detach from the VPLS service context. The IP interface may then be bound to another service name or a SAP or SDP binding may be created for the interface using the **sap** or **spoke-sdp** commands on the interface.

IP Interface MTU and Fragmentation

A VPLS service is affected by two MTU values; port MTUs and the VPLS service MTU. The MTU on each physical port defines the largest Layer 2 packet (including all DLC headers and CRC) that may be transmitted out a port. The VPLS itself has a service level MTU that defines the largest packet supported by the service. This MTU does not include the local encapsulation overhead for each port (QinQ, Dot1Q, TopQ or SDP service delineation fields and headers) but does include the remainder of the packet. As virtual ports are created in the system, the virtual port cannot become operational unless the configured port MTU minus the virtual port service delineation overhead is greater than or equal to the configured VPLS service MTU. Thus, an operational virtual port is ensured to support the largest packet traversing the VPLS service. The service delineation overhead on each Layer 2 packet is removed before forwarding into a VPLS service. VPLS services do not support fragmentation and must discard any Layer 2 packet larger than the service MTU after the service delineation overhead is removed.

IP interfaces have a configurable up MTU that defines the largest packet that may egress the IP interface without being fragmented. This MTU encompasses the IP portion of the packet and does not include any of the egress DLC header or CRC. This MTU does not affect the size of the largest ingress packet on the IP interface. If the egress IP portion of the packet is larger than the IP interface MTU and the IP header do not fragment flag is not set, the packet is fragmented into smaller packets that will not exceed the configured MTU size. If the do not fragment bit is set, the packet is silently discarded at egress when it exceeds the IP MTU.

When the IP interface is bound to a VPLS service, the IP MTU must be at least 18 bytes less than the VPLS service MTU. This allows for the addition of the minimal Ethernet encapsulation overhead; 6 bytes for the DA, 6 bytes for the SA, 2 bytes for the Etype and 4 bytes for the trailing CRC. Any remaining egress virtual port overhead (Dot1P, Dot1Q, QinQ, TopQ or SDP) required above the minimum is known to be less than the egress ports MTU since the virtual port would not be operational otherwise.

If the IP interface IP MTU value is too large based on the VPLS service MTU, the IP interface will enter the operationally down state until either the IP MTU is adequately lowered or the VPLS service MTU is sufficiently increased.

The **no** form of this command on the IP interface is used to remove the service name binding from the IP interface. If the service name has been resolved to a VPLS service context and the IP interface has been attached to the VPLS service, the IP interface will also be detached from the VPLS service.

Parameters

service-name

The service-name parameter is required when using the IP interface vpls command and specifies the service name that the system will attempt to resolve to an **allow-ip-int-bind** enabled VPLS service associated with the name. The specified name is expressed as an ASCII string comprised of up to 32 characters. It does not need to already be associated with a service and the system does not check to ensure that multiple IP interfaces are not bound to the same name.

Platforms

7705 SAR Gen 2

32.23 vpls-group

vpls-group

Syntax

vpls-group *vpls-group-id* [**create**]

no vpls-group *vpls-group-id*

Context

[Tree] (config>service>vpls vpls-group)

Full Context

configure service vpls vpls-group

Description

This command defines a vpls-group index. Multiple vpls-group commands can be specified to allow the use of different VPLS and SAP templates for different ranges of service ids. A vpls-group can be deleted only in shutdown state. Multiple commands under different vpls-group ids can be issued and can be in progress at the same time.

Default

no vpls-group

Parameters

vpls-group-id

Specifies the ID associated with the VPLS group

Values 1 to 4094

Platforms

7705 SAR Gen 2

32.24 vpls-id

vpls-id

Syntax

vpls-id *vpls-id*

Context

[Tree] (config>service>vpls>bgp-ad vpls-id)

Full Context

configure service vpls bgp-ad vpls-id

Description

This command configures the VPLS ID component that is signaled in one of the extended community attributes (*ext-comm*).

Values and format (6 bytes, other 2 bytes of type-subtype is automatically generated)

Parameters

vpls-id

Specifies a globally unique VPLS ID for BGP auto-discovery in this VPLS service

Values vpls-id: <ip-addr:comm-val>| <as-number:ext-comm-val>
ip-addr: a.b.c.d
comm-val: [0 to 65535]
as-number: [1 to 65535]
ext-comm-val: [0 to 4294967295]

Platforms
7705 SAR Gen 2

32.25 vpls-sap-template

vpls-sap-template

Syntax
vpls-sap-template *name/id* **create**
[**no**] **vpls-sap-template** *name/id*

Context
[Tree] (config>service>template vpls-sap-template)

Full Context
configure service template vpls-sap-template

Description
This is the command used to create a SAP template to be used in a vpls-template. Only certain existing VPLS SAP attributes can be changed in the vpls-sap-template, not in the instantiated VPLS SAP

The following SAP attributes are set in the instantiated saps (no configuration allowed):

description: "Sap <sap-id> controlled by MVRP service <svc id>" – auto generated

shutdown: no shutdown

Parameters
name/id

Specifies the name in ASCII or the template ID

Values 1 to 2147483647

Platforms
7705 SAR Gen 2

32.26 vpls-template

vpls-template

Syntax

vpls-template *name/id* **create**

[**no**] **vpls-template** *name/id*

Context

[Tree] (config>service>template vpls-template)

Full Context

configure service template vpls-template

Description

This command is used to create a vpls-template to be used to auto-instantiate a range of VPLS services. Only certain existing VPLS attributes specified in the command reference section can be changed in the vpls-template, not in the instantiated VPLS. The following attributes are automatically set in the instantiated VPLSs (no template configuration necessary) and the operator cannot change these values.

vpn-id: none

description: "Service <svc id> auto-generated by control VPLS <svc-id>”

service-name: "Service <svc id>” (Auto-generated)

shutdown: no shutdown

Following existing attributes can be set by the user in the instantiated VPLSs:

[no] sap

All the other VPLS attributes are not supported.

Parameters

name/id

Specifies the name in ASCII or the template ID

Values name: ASCII string

Values ID: [1 to 2147483647]

Platforms

7705 SAR Gen 2

32.27 vpls-template-binding

vpls-template-binding

Syntax

vpls-template-binding *name/id*

no vpls-template-binding

Context

[Tree] (config>service>vpls>vpls-group vpls-template-binding)

Full Context

configure service vpls vpls-group vpls-template-binding

Description

This command configures the binding to a VPLS template to be used to instantiate pre-provisioned data VPLS using as input variables the service IDs generated by the vid-range command.

The **no** form of this command removes the binding and deletes the related VPLS instances. The command will fail if any of the affected VPLS instances have either a provisioned SAP or an active MVRP declaration/registration or if the related vpls-group id is in no shutdown state. Any changes to the vpls-template-binding require the vpls-group to be in shutdown state.

Default

no vpls-template-binding

Parameters

name/id

Specifies the name or the ID of the VPLS template

Values 1 to 1024

Platforms

7705 SAR Gen 2

32.28 vpn-apply-export

vpn-apply-export

Syntax

[**no**] **vpn-apply-export**

Context

[Tree] (config>router>bgp>group>neighbor vpn-apply-export)

[Tree] (config>router>bgp>group vpn-apply-export)

[Tree] (config>router>bgp vpn-apply-export)

Full Context

configure router bgp group neighbor vpn-apply-export

configure router bgp group vpn-apply-export

configure router bgp vpn-apply-export

Description

This command causes the base instance BGP export route policies to be applied to vpn-ipv4/6, mvpn-ipv4/6, l2-vpn, mdt-safi, mcast-vpn-ipv4, and evpn routes.

The **no** form of this command disables the application of the base instance BGP route policies to vpn-ipv4/6, mvpn-ipv4/6, l2-vpn, mdt-safi, mcast-vpn-ipv4, and evpn routes.

Default

no vpn-apply-export

Platforms

7705 SAR Gen 2

32.29 vpn-apply-import

vpn-apply-import

Syntax

[**no**] **vpn-apply-import**

Context

[Tree] (config>router>bgp>group vpn-apply-import)

[Tree] (config>router>bgp>group>neighbor vpn-apply-import)

[Tree] (config>router>bgp vpn-apply-import)

Full Context

configure router bgp group vpn-apply-import

configure router bgp group neighbor vpn-apply-import

configure router bgp vpn-apply-import

Description

This command causes the base instance BGP import route policies to be applied to vpn-ipv4/6, mvpn-ipv4/6, l2-vpn, mdt-safi, mcast-vpn-ipv4, and evpn routes.

The **no** form of this command disables the application of the base instance BGP import route policies to vpn-ipv4/6, mvpn-ipv4/6, l2-vpn, mdt-safi, mcast-vpn-ipv4, and evpn routes.

Default

no vpn-apply-import

Platforms

7705 SAR Gen 2

32.30 vpn-domain

vpn-domain

Syntax

vpn-domain [*type* {**0005** | **0105** | **0205** | **8005**}] **id** *id*

no vpn-domain

Context

[Tree] (config>service>vprn>ospf vpn-domain)

Full Context

configure service vprn ospf vpn-domain

Description

This command specifies type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID. This applies to VPRN instances of OSPF only. An attempt to modify the value of this object will result in an inconsistent value error when is not a VPRN instance. The parameters are mandatory and can be entered in either order. This command is not applicable in the **config>service>vprn>ospf3** context.

This command is not supported in OSPF3.

Default

no vpn-domain

Parameters

id

Specifies the OSPF VPN domain in the "xxxx.xxxx.xxxx" format. This is exchanged using BGP in the extended community attribute associated with a prefix. This object applies to VPRN instances of OSPF only.

type

Specifies the type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID.

Values 0005, 0105, 0205, 8005

Platforms

7705 SAR Gen 2

32.31 vpn-family-policy

vpn-family-policy

Syntax

vpn-family-policy *policy-name*

no vpn-family-policy

Context

[Tree] (config>router>bgp>next-hop-resolution vpn-family-policy)

Full Context

configure router bgp next-hop-resolution vpn-family-policy

Description

This command specifies the VPN family policy that is applied when filtering routes for consideration for next-hop resolution process for EVPN and IP-VPN families.

This policy is supported by the following families:

- VPN-IPv4 and VPN-IPv6
- EVPN (all routes types 1-6, although AD per-ES and AD per-EVI routes are always shown as resolved)
- MCAST-VPN-IPv4 and MCAST-VPN-IPv6

In a VPN family policy:

- only prefix-lists are used to match the next hop of a resolving route. No other policy qualifiers are supported.
- the route resolving the next hop is accepted or rejected

  In other words, if an imported route's next hop is resolved by route N (N is the preferred entry in tunnel-table for MPLS or the longest prefix match in the route-table for VXLAN), and route N is rejected by vpn-family-policy, then the route next hop is unresolved. This is irrespective of the existence of a route M that could potentially resolve the next hop in the tunnel-table or route-table.

The **no** form of this command removes the VPN family policy.

Default

no vpn-family-policy

Parameters

policy-name

Specifies the route policy name. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

Platforms

7705 SAR Gen 2

32.32 vpn-gre-source-ip

vpn-gre-source-ip

Syntax

vpn-gre-source-ip *ip-address*

no vpn-gre-source-ip

Context

[Tree] (config>service>system vpn-gre-source-ip)

Full Context

configure service system vpn-gre-source-ip

Description

This command configures a single system-wide alternate source IPv4 address of the GRE tunnels in all VPRN services using the **auto-bind-tunnel** or an explicit SDP binding (**config**>**service**>**vprn**>**spoke-sdp**) with a tunnel of encapsulation GRE.

A change to the value of the **vpn-gre-source-ip** parameter can be performed without disabling the service. Once the new value is configured, the system address is not used in services which bind to the GRE tunnel.

The primary IPv4 address of any local network IP interface, loopback or otherwise, may be used.

The address of the following interfaces are not supported, and the configuration is rejected:

- unnumbered network IP interface
- IES interface
- VPRN interface
- CSC VPRN interface

The **vpn-gre-source-ip** parameter value adheres to the following rules:

- This single source address counts towards the maximum of 15 distinct address values per system that are used by all GRE SDPs under the **config**>**service**>**sdp**>**local-end** context and all L2oGRE SDPs under the **config**>**service**>**system**>**gre-eth-bridged**>**tunnel-termination** context.
- The same source address can be used in both **vpn-gre-source-ip** and **config**>**service**>**sdp**>**local-end** contexts.
- The same source address cannot be used in both **vpn-gre-source-ip** and **config**>**service**>**system**>**gre-eth-bridged**>**tunnel-termination** contexts because an address configured for a L2oGRE SDP matches an internally created interface which is not available to other applications.
- The **vpn-gre-source-ip** address, when different from system, need not match the primary address of an interface which has the MPLS-over-GRE termination subnet configured, unless a GRE SDP or tunnel from the far-end router terminates on this address.

The **no** form of the command reverts to the default value.

Default

vpn-gre-source-ip ip-address (System interface primary IPv4 address)

Parameters

ip-address

Specifies the IPv4 address (a.b.c.d).

Platforms

7705 SAR Gen 2

32.33 vpn-ipv4

vpn-ipv4

Syntax

vpn-ipv4 send *send-limit* **receive** [**none**]

vpn-ipv4 send *send-limit*

no vpn-ipv4

Context

[Tree] (config>router>bgp>group>add-paths vpn-ipv4)

[Tree] (config>router>bgp>add-paths vpn-ipv4)

[Tree] (config>router>bgp>group>neighbor>add-paths vpn-ipv4)

Full Context

configure router bgp group add-paths vpn-ipv4

configure router bgp add-paths vpn-ipv4

configure router bgp group neighbor add-paths vpn-ipv4

Description

This command configures the add-paths capability for VPN-IPv4 routes. By default, add-paths is not enabled for VPN-IPv4 routes.

The maximum number of paths per VPN-IPv4 NLRI to send is the configured send-limit, which is a mandatory parameter. The capability to receive multiple paths per prefix from a peer is configurable using the **receive** keyword, which is optional. If the **receive** keyword is not included in the command the receive capability is enabled by default.

The **no** form of this command disables add-paths support for VPN-IPv4 routes, causing sessions established using add-paths for VPN-IPv4 to go down and come back up without the add-paths capability.

Default

no vpn-ipv4

Parameters

send-limit

Specifies the maximum number of paths per VPN-IPv4 NLRI that are allowed to be advertised to add-paths peers (the actual number of advertised routes may be less depending on the next-hop diversity requirement, other configuration options, route policies, or route advertisement rules). If the value is **multipaths**, then BGP advertises all of the used BGP multipaths for each VPN-IPv4 NLRI if the peer has signaled support for receiving multiple add paths. If the router has not installed any of the routes in its FIB then all BGP add-paths qualify for advertisement.

Values 1 to 16, none, multipaths

receive

Specifies that the router negotiates the add-paths receive capability for VPN-IPv4 routes with its peers.

none

Specifies that the router does not negotiate the add-paths receive capability for VPN-IPv4 routes with its peers.

Platforms

7705 SAR Gen 2

32.34 vpn-ipv6

vpn-ipv6

Syntax

vpn-ipv6 send *send-limit* **receive** [**none**]

vpn-ipv6 send *send-limit*

no vpn-ipv6

Context

[Tree] (config>router>bgp>group>add-paths vpn-ipv6)

[Tree] (config>router>bgp>add-paths vpn-ipv6)

[Tree] (config>router>bgp>group>neighbor>add-paths vpn-ipv6)

Full Context

configure router bgp group add-paths vpn-ipv6

configure router bgp add-paths vpn-ipv6

configure router bgp group neighbor add-paths vpn-ipv6

Description

This command configures the add-paths capability for VPN-IPv6 routes. By default, add-paths is not enabled for VPN-IPv6 routes.

The maximum number of paths per VPN-IPv6 NLRI to send is the configured send-limit, which is a mandatory parameter. The capability to receive multiple paths per prefix from a peer is configurable using the **receive** keyword, which is optional. If the **receive** keyword is not included in the command the receive capability is enabled by default.

The **no** form of this command disables add-paths support for VPN-IPv6 routes, causing sessions established using add-paths for VPN-IPv6 to go down and come back up without the add-paths capability.

Default

no vpn-ipv6

Parameters

send-limit

Specifies the maximum number of paths per VPN-IPv6 NLRI that are allowed to be advertised to add-paths peers (the actual number of advertised routes may be less depending on the next-hop diversity requirement, other configuration options, route policies, or route advertisement rules). If the value is **multipaths**, then BGP advertises all

of the used BGP multipaths for each VPN-IPv6 NLRI if the peer has signaled support for receiving multiple add paths. If the router has not installed any of the routes in its FIB then all BGP add-paths qualify for advertisement.

Values 1 to 16, none, multipaths

receive

Specifies that the router negotiates the add-paths receive capability for VPN-IPv6 routes with its peers.

none

Specifies that the router does not negotiate the add-paths receive capability for VPN-IPv6 routes with its peers.

Platforms

7705 SAR Gen 2

32.35 vpn-tag

vpn-tag

Syntax

vpn-tag *vpn-tag*

no vpn-tag

Context

[Tree] (config>service>vprn>ospf vpn-tag)

Full Context

configure service vprn ospf vpn-tag

Description

This command specifies the route tag for an OSPF VPN on a PE router. This field is set in the tag field of the OSPF external LSAs generated by the PE. This is mainly used to prevent routing loops. This applies to VPRN instances of OSPF only. An attempt to modify the value of this object will result in an inconsistent value error when is not a VPRN instance.

This command is not supported in OSPF3.

Default

vpn-tag 0

Platforms

7705 SAR Gen 2

32.36 vprn

vprn

Syntax

vprn *service-id* [**name** *name*] [**customer** *customer-id*] [**create**]

no vprn *service-id*

Context

[Tree] (config>service vprn)

Full Context

configure service vprn

Description

This command creates or edits a Virtual Private Routed Network (VPRN) service instance.

If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

VPRN services allow the creation of customer-facing IP interfaces in the same routing instance used for service network core routing connectivity. VPRN services require that the IP addressing scheme used by the subscriber must be unique between it and other addressing schemes used by the provider and potentially the entire Internet.

IP interfaces defined within the context of an VPRN service ID must have a SAP created as the access point to the subscriber network.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the customer command in the service context. When a service is created with a customer association, it is not possible to edit the customer association. The service must be deleted and re-created with a new customer association.

When a service is created, the use of the **customer** *customer-id* is optional to navigate into the service configuration context. If attempting to edit a service with the incorrect *customer-id* results in an error.

Multiple VPRN services are created to separate customer-owned IP interfaces. More than one VPRN service can be created for a single customer ID. More than one IP interface can be created within a single VPRN service ID. All IP interfaces created within an VPRN service ID belongs to the same customer.

The **no** form of this command deletes the VPRN service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces and all routing protocol configurations defined within the service ID have been shut down and deleted.

Parameters

service-id

Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service

of any type. The *service-id* must be the same number used for every 7705 SAR Gen 2 router on which this service is defined.

Values
service-id: 1 to 2147483648
svc-name: 64 characters maximum

customer-id

Specifies an existing customer identification number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

name name

This parameter configures an optional VPRN name, up to 64 characters, which adds a name identifier to a given vprn to then use that vprn name in configuration references as well as display and use vprn names in show commands throughout the system. This helps the service provider/administrator to identify and manage vprn within the SR OS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a given service once it is initially created.

If a name is not specified at creation time, then SR OS assigns a string version of the service-id as the name.

Service names may not begin with an integer (0 to 9).

create

Keyword used to create a service ID. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

Platforms

7705 SAR Gen 2

vprn

Syntax

vprn *service-id*

no vprn

Context

[Tree] (config>system>security>vprn-aaa-server vprn)

Full Context

configure system security vprn-aaa-server vprn

Description

This command configures TACACS+ or RADIUS servers in a VPRN to be used for AAA by that VPRN and by sessions in VPRNs without a AAA server configured.

The **no** form of this command disables the use of servers in a VPRN.

Default

no vprn

Parameters

service-id

Specifies the VPRN server for AAA to use for sessions in VPRNs without a AAA server.

Values *service-id*: 1 to 2147483648

svc-name: 64 characters maximum

Platforms

7705 SAR Gen 2

32.37 vprn-aaa-server

vprn-aaa-server

Syntax

vprn-aaa-server

Context

[Tree] (config>system>security vprn-aaa-server)

Full Context

configure system security vprn-aaa-server

Description

Commands in this context configure the use of AAA servers in a VPRN.

Platforms

7705 SAR Gen 2

32.38 vprn-auto-bind

vprn-auto-bind

Syntax

vprn-auto-bind [**include** | **exclude**]

Context

[Tree] (config>router>mpls>lsp vprn-auto-bind)

[Tree] (config>router>mpls>lsp-template vprn-auto-bind)

Full Context

configure router mpls lsp vprn-auto-bind

configure router mpls lsp-template vprn-auto-bind

Description

This command determines whether the associated names LSP can be used or not as part of the auto-bind feature for VPRN services. By default, a names LSP is available for inclusion to be used for the auto-bind feature.

By configuring the command vprn-auto-bind exclude, the associated LSP will not be used by the auto-bind feature within VPRN services.

The **no** form of this command resets the flag back to the default value.

Default

vprn-auto-bind include

Parameters

include

Allows an associated LSP to be used by auto-bin for vprn services

exclude

Disables the use of the associated LSP to be used with the auto-bind feature for VPRN services.

Platforms

7705 SAR Gen 2

32.39 vprn-local

vprn-local

Syntax

vprn-local [{**none** | **all** | **vc-only**}]

Context

[Tree] (config>router>ttl-propagate vprn-local)

Full Context

configure router ttl-propagate vprn-local

Description

This command configures the TTL propagation for locally generated packets which are forwarded over a MPLS LSPs in all VPRN service contexts.

For vpn-ipv4 and vpn-ipv6 packets forwarded in the context of all VPRN services in the system, including 6VPE packets, the all value of the command enables TTL propagation from the IP header into all labels in the stack:

The user can enable the TTL propagation behavior separately for locally generated packets by CPM (vprn-local) and for user and control packets in transit at the node (vprn-transit).

The vc-only value reverts to the default behavior by which the IP TTL is propagated into the VC label but not to the transport labels in the stack. The user can explicitly set the default behavior by configuring the vc-only value. This command does not have a no version.

The value none allows the user to disable the propagation of the IP TTL to all labels in the stack, including the VC label. This is needed for a transparent operation of UDP traceroute in VPRN inter-AS option B such that the ingress and egress ASBR nodes are not traced.

The user can override the global configuration within each VPRN instance using the following commands:

- config service vprn ttl-propagate local [inherit | none | vc-only | all]
- config service vprn ttl-propagate transit [inherit | none | vc-only | all]

The default behavior for a given VPRN instance is to inherit the global configuration for the same command. The user can explicitly set the default behavior by configuring the inherit value.

When a packet is received in a VPRN context but is looked up in the Global Routing Table (GRT), for example, leaking to GRT is enabled, the behavior of the TTL propagation is governed by the RSVP or LDP shortcut configuration when the matching routing is a LSP shortcut route. It is governed by the BGP label route configuration when the matching route is a RFC 8277 label route or a 6PE route.

When a packet is received on one VPRN instance and is redirected using Policy Based Routing (PBR) to be forwarded in another VPRN instance, the TTL propagation is governed by the configuration of the outgoing VPRN instance.

Default

vprn-local vc-only

Parameters

none

Specifies that the TTL of the IP packet is not propagated into the VC label or labels in the transport label stack

all

Specifies that the TTL of the IP packet is propagated into the VC label and all labels in the transport label stack.

vc-only

Specifies that the TTL of the IP packet is propagated into the VC label and not into the labels in the transport label stack.

Platforms

7705 SAR Gen 2

32.40 vprn-network-exceptions

vprn-network-exceptions

Syntax

vprn-network-exceptions *number seconds*

no vprn-network-exceptions

Context

[Tree] (config>system>security vprn-network-exceptions)

Full Context

configure system security vprn-network-exceptions

Description

This command configures the rate to limit the processing of packets with label TTL expiry received within an LSP shortcut, or within all VPRN instances in the system, and from all network IP interfaces. This includes labeled user and control plane packets, ping and traceroute packets within GRT and VPRN, and ICMP replies. Packets over the configured rate are dropped.

This feature does not rate limit MPLS and service OAM packets (vprn-ping, vprn-trace, lsp-ping, lsp-trace, vccv-ping, and vccv-trace).

The **no** form of this command disables the rate limiting of the reply to these packets.

Parameters

number

Specifies the number limit of MPLS exception messages.

Values 10 to 10,000

seconds

Specifies the rate limit of MPLS exception messages, in seconds.

Values 1 to 60

Platforms

7705 SAR Gen 2

32.41 vprn-transit

vprn-transit

Syntax

vprn-transit [{**none** | **all** | **vc-only**}]

Context

[Tree] (config>router>ttl-propagate vprn-transit)

Full Context

configure router ttl-propagate vprn-transit

Description

This command configures the TTL propagation for in transit packets which are forwarded over a MPLS LSPs in all VPRN service contexts. For vpn-ipv4 and vpn-ipv6 packets forwarded in the context of all VPRN services in the system, including 6VPE packets, the all value of the command enables TTL propagation from the IP header into all labels in the stack:

The user can enable the TTL propagation behavior separately for locally generated packets by CPM (vprn-local) and for user and control packets in transit at the node (vprn-transit).

The vc-only value reverts to the default behavior by which the IP TTL is propagated into the VC label but not to the transport labels in the stack. The user can explicitly set the default behavior by configuring the vc-only value. This command does not have a no version.

The value none allows the user to disable the propagation of the IP TTL to all labels in the stack, including the VC label. This is needed for a transparent operation of UDP trace-route in VPRN inter-AS option B such that the ingress and egress ASBR nodes are not traced.

The user can override the global configuration within each VPRN service instance using the following commands:

- config service vprn ttl-propagate local [inherit | none | vc-only | all]

- config service vprn ttl-propagate transit [inherit | none | vc-only | all]

The default behavior for a given VPRN instance is to inherit the global configuration for the same command. The user can explicitly set the default behavior by configuring the inherit value.

When a packet is received in a VPRN context but is looked up in the Global Routing Table (GRT), for example, leaking to GRT is enabled, the behavior of the TTL propagation is governed by the RSVP or LDP shortcut configuration when the matching routing is a LSP shortcut route. It is governed by the BGP label route configuration when the matching route is a RFC 8277 label route or a 6PE route.

When a packet is received on one VPRN instance and is redirected using Policy Based Routing (PBR) to be forwarded in another VPRN instance, the TTL propagation is governed by the configuration of the outgoing VPRN instance

Default

vprn-transit vc-only

Parameters

none

Specifies that the TTL of the IP packet is not propagated into the VC label or labels in the transport label stack

all

Specifies that the TTL of the IP packet is propagated into the VC label and all labels in the transport label stack.

vc-only

Specifies that the TTL of the IP packet is propagated into the VC label and not into the labels in the transport label stack.

Platforms

7705 SAR Gen 2

32.42 vrf-export

vrf-export

Syntax

vrf-export *plcy-or-long-expr* [*plcy-or-expr*]

no vrf-export

Context

[Tree] (config>service>vprn>bgp-ipvpn>mpls vrf-export)

[Tree] (config>service>vprn>bgp-evpn>mpls vrf-export)

Full Context

configure service vprn bgp-ipvpn mpls vrf-export

configure service vprn bgp-evpn mpls vrf-export

Description

This command configures route policies that control how routes are exported from the local VRF to other VRFs on the same or remote PE routers (using MP-BGP). Route policies are configured in the **configure router policy-options** context.

The **vrf-export** command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine final action to accept or reject the route.

Only one of the 15 objects referenced by the **vrf-export** command can be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters while the remaining 14 objects have a maximum length of 64 characters each.

When multiple **vrf-export** commands are issued, the last command entered overrides the previous command.

Aggregate routes are not advertised using MP-BGP protocols to the other MP-BGP peers.

The **no** form of this command removes all route policy names from the **vrf-export** list.

Default

no vrf-export

Parameters

plcy-or-long-expr

Specifies the route policy name (up to 64 characters) or a policy logical expression (up to 255 characters).

plcy-or-expr

Specifies the route policy name (up to 64 characters) or a policy logical expression (up to 255 characters). Up to 14 policies may be entered.

Platforms

7705 SAR Gen 2

vrf-export

Syntax

vrf-export

Context

[Tree] (config>service>vprn vrf-export)

Full Context

configure service vprn vrf-export

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

32.43 vrf-import

vrf-import

Syntax

vrf-import *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]

no vrf-import

Context

[Tree] (config>service>vprn>bgp-ipvpn>mpls vrf-import)

[Tree] (config>service>vprn>bgp-evpn>mpls vrf-import)

Full Context

configure service vprn bgp-ipvpn mpls vrf-import

configure service vprn bgp-evpn mpls vrf-import

Description

This command configures route policies that control how VPN-IP and EVPN-IFL routes exported by other VRFs, on the same or remote PEs, are imported into the local VRF. Route policies are configured in the **configure router policy-options** context.

The **vrf-import** command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine final action to accept or reject the route

Only one of the 15 objects referenced by the **vrf-import** command is allowed to be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters while the remaining 14 objects have a maximum length of 64 characters each.

When multiple **vrf-import** commands are issued, the last command entered overrides the previous command.

The **no** form of this command removes all route policy names from the import list

Note:
Unless the preference value is changed by the policy, BGP-VPN and EVPN-IFL routes imported with a **vrf-import** policy have the preference value of 170 when imported from remote PE routers, or retain the protocol preference value of the exported route when imported from other VRFs on the same router.

Default

no vrf-import

Parameters

plcy-or-long-expr

Specifies the route policy name (up to 64 characters) or a policy logical expression (up to 255 characters).

plcy-or-expr

Specifies the route policy name (up to 64 characters) or a policy logical expression (up to 255 characters).

Platforms

7705 SAR Gen 2

vrf-import

Syntax

vrf-import

Context

[Tree] (config>service>vprn vrf-import)

Full Context

configure service vprn vrf-import

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

32.44 vrf-target

vrf-target

Syntax

vrf-target {*ext-community* | **export** *ext-community* | **import** *ext-community* | **export** *ext-community* **import** *ext-community*}

no vrf-target

Context

[Tree] (config>service>vprn>bgp-ipvpn>mpls vrf-target)

[Tree] (config>service>vprn>bgp-evpn>mpls vrf-target)

Full Context

configure service vprn bgp-ipvpn mpls vrf-target

configure service vprn bgp-evpn mpls vrf-target

Description

This command provides a simplified method to configure the route target added to advertised routes or compared against received routes from other VRFs on the same or remote PE routers (using MP-BGP).

BGP-VPN and EVPN-IFL routes imported with a VRF target policy use the BGP preference value of 170 when imported from remote PE routers, or retain the protocol preference value of the exported route when imported from other VRFs in the same router.

Specified VRF import or VRF export policies override the VRF target policy.

The **no** form of this command removes the VRF target policy.

Default

no vrf-target

Parameters

ext-comm

Specifies an extended BGP community in the *type:x:y* format. The value *x* can be an integer or IP address. The *type* can be the target or origin. *y* can be 16-bit integers.

Values

<ext-community>	: target:{<ip-addr:comm-val> | <2byte-asnumber:ext-comm-val> | <4byte-asnumber:comm-val>}	
	ip-addr:	a.b.c.d
	comm-val:	[0 to 65535]
	2byte-asnumber:	[0 to 65535]

ext-comm-val: [0 to 4294967295]

4byte-asnumber: [0 to 4294967295]

import *ext-community*

Specifies communities allowed to be received from remote PE neighbors.

export *ext-community*

Specifies communities allowed to be sent to remote PE neighbors.

Platforms

7705 SAR Gen 2

vrf-target

Syntax

vrf-target

Context

[Tree] (config>service>vprn vrf-target)

Full Context

configure service vprn vrf-target

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

32.45 vrrp

vrrp

Syntax

vrrp *virtual-router-id* [**owner**] [**passive**] [**monitor-oper-group** *group-name*]

no vrrp *virtual-router-id*

Context

[Tree] (config>service>ies>if>ipv6 vrrp)

[Tree] (config>service>ies>if vrrp)

Full Context

configure service ies interface ipv6 vrrp

configure service ies interface vrrp

Description

This command configures the router to create or edit a Virtual Router ID (VRID) on the service IP interface. A VRID is internally represented in conjunction with the IP interface name. This allows the VRID to be used on multiple IP interfaces while representing different virtual router instances.

Two VRRP nodes can be defined on an IP interface. The **vrrp** *virtual-router-id* command is used to define the configuration parameters for the VRID.

The **no** form of this command removes the specified VRID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the VRID. The VRID does not need to be shutdown to remove the virtual router instance.

Parameters

virtual-router-id

Specifies a virtual router ID or an ID that can be modified on the IP interface.

Values 1 to 255

owner

Keyword used to identify this virtual router instance as owning the virtual router IP addresses. If the **owner** keyword is not specified at the time of *vrid* creation, the **vrrp backup** commands must be specified to define the virtual router IP addresses. The **owner** keyword is not required when entering the *vrid* for editing purposes. When created as **owner**, a *vrid* on an IP interface cannot have the **owner** parameter removed. The *vrid* must be deleted, and then recreated without the **owner** keyword, to remove ownership.

passive

Keyword used to identify this virtual router instance as **passive**, and therefore, owning the virtual router IP addresses. A **passive** *vrid* does not send or receive VRRP advertisement messages, and is always in either the **master** state (if the interface is operationally up), or the **init** state (if the interface is operationally down). The **passive** keyword is not required when entering the *vrid* for editing purposes. When a *vrid* on an IP interface is created as **passive**, the parameter cannot be removed from the *vrid*. The *vrid* must be deleted, and then recreated without the **passive** keyword, to remove parameter.

group-name

Specifies the name of the **oper-group**, up to 32 characters to establish the associated VRRP instance as a following instance to the specified operational group. As a result of this association, the VRRP instance state follows that of the VRRP instance (the lead instance) associated with the specified operation group.

Platforms

7705 SAR Gen 2

vrrp

Syntax

vrrp *virtual-router-id* [**owner**] [**passive**] [**monitor-oper-group** *group-name*]

no vrrp *virtual-router-id*

Context

[Tree] (config>service>vprn>if vrrp)

Full Context

configure service vprn interface vrrp

Description

This command creates or edits a Virtual Router ID (VRID) on the service IP interface. A VRID is internally represented in conjunction with the IP interface name. This allows the VRID to be used on multiple IP interfaces while representing different virtual router instances.

Two VRRP nodes can be defined on an IP interface. One, both, or none may be defined as owner. The nodal context of **vrrp** *virtual-router-id* is used to define the configuration parameters for the VRID.

The **no** form of this command removes the specified VRID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the VRID. The VRID does not need to be shut down in order to remove the virtual router instance.

Parameters

virtual-router-id

Specifies a new virtual router ID or one that can be modified on the IP interface.

Values 1 to 255

owner

Identifies this virtual router instance as owning the virtual router IP addresses. If the **owner** keyword is not specified at the time of *vrid* creation, the **vrrp backup** commands must be specified to define the virtual router IP addresses. The **owner** keyword is not required when entering the *vrid* for editing purposes. Once created as **owner**, a *vrid* on an IP interface cannot have the **owner** parameter removed. The *vrid* must be deleted, and then recreated without the **owner** keyword, to remove ownership.

passive

Identifies this virtual router instance as **passive**, and therefore, owning the virtual router IP addresses. A **passive** *vrid* does not send or receive VRRP advertisement messages, and is always in either the **master** state (if the interface is operational-up), or the **init** state (if the interface is operational-down). The **passive** keyword is not required when entering the *vrid* for editing purposes. Once a *vrid* on an IP interface is created as **passive**, the parameter cannot be removed from the *vrid*. The *vrid* must be deleted, and then recreated without the **passive** keyword, to remove parameter.

group-name

Specifies the name of the **oper-group**, up to 32 characters to establish the associated VRRP instance as a following instance to the specified operational group. As a result of this association, the VRRP instance state follows that of the VRRP instance (the lead instance) associated with specified operation group.

Platforms

7705 SAR Gen 2

vrrp

Syntax

vrrp *virtual-router-id* [**owner**] [**passive**] [**monitor-oper-group** *group-name*]

no vrrp *virtual-router-id*

Context

[Tree] (config>service>vprn>if vrrp)

Full Context

configure service vprn interface vrrp

Description

This command configures the router to create or edit a Virtual Router ID (VRID) on the service IP interface. A VRID is internally represented in conjunction with the IP interface name. This allows the VRID to be used on multiple IP interfaces while representing different virtual router instances.

Two VRRP nodes can be defined on an IP interface. The **vrrp** *virtual-router-id* command is used to define the configuration parameters for the VRID.

The **no** form of this command removes the specified VRID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the VRID. The VRID does not need to be shutdown to remove the virtual router instance.

Parameters

virtual-router-id

Specifies a virtual router ID or an ID that can be modified on the IP interface.

Values 1 to 255

owner

Keyword used to identify this virtual router instance as owning the virtual router IP addresses. If the **owner** keyword is not specified at the time of *vrid* creation, the **vrrp backup** commands must be specified to define the virtual router IP addresses. The **owner** keyword is not required when entering the *vrid* for editing purposes. When created as **owner**, a *vrid* on an IP interface cannot have the **owner** parameter removed. The *vrid* must be deleted, and then recreated without the **owner** keyword, to remove ownership.

passive

Keyword used to identify this virtual router instance as **passive**, and therefore, owning the virtual router IP addresses. A **passive** *vrid* does not send or receive VRRP advertisement messages, and is always in either the **master** state (if the interface is operationally up), or the **init** state (if the interface is operationally down). The **passive** keyword is not required when entering the *vrid* for editing purposes. When a *vrid* on an IP interface is created as **passive**, the parameter cannot be removed from the *vrid*. The *vrid* must be deleted, and then recreated without the **passive** keyword, to remove parameter.

group-name

Specifies the name of the **oper-group**, up to 32 characters to establish the associated VRRP instance as a following instance to the specified operational group. As a result of this association, the VRRP instance state follows that of the VRRP instance (the lead instance) associated with the specified operation group.

Platforms

7705 SAR Gen 2

vrrp

Syntax

vrrp *virtual-router-id* [**owner**] [**passive**] [**monitor-oper-group** *group-name*]

no vrrp *virtual-router-id*

Context

[Tree] (config>router>if vrrp)

[Tree] (config>router>if>ipv6 vrrp)

Full Context

configure router interface vrrp

configure router interface ipv6 vrrp

Description

This command creates the context to configure a VRRP virtual router instance. A virtual router is defined by its virtual router identifier (VRID) and a set of IP addresses.

The optional **owner** keyword indicates that the **owner** controls the IP address of the virtual router and is responsible for forwarding packets sent to this IP address. The **owner** assumes the role of the master virtual router.

All other virtual router instances participating in this message domain must have the same *vrid* configured and cannot be configured as **owner**. Once created, the **owner** keyword is optional when entering the *vrid* for configuration purposes.

A *vrid* is internally associated with the IP interface. This allows the *vrid* to be used on multiple IP interfaces while representing different virtual router instances.

For IPv4, up to four VRRP VRID nodes can be configured on a router interface. Each virtual router instance can manage up to 16 backup IP addresses. For IPv6, only one VRID can be configured on a router interface.

The optional **passive** keyword indicates that a *vrid* can be configured as **passive**, in which case, the VRRP advertisement messages are suppressed on transmission and reception, and all routers configured with the same *vrid* become master. Passive *VRIDs* can exceed the limit of four VRRP VRID nodes on a router interface.

The **no** form of the command removes the specified *vrid* from the IP interface. This terminates VRRP participation and deletes all references to the *vrid* in conjunction with the IP interface. The *vrid* does not need to be shut down to remove the virtual router instance.

Default

no vrrp — No VRRP virtual router instance is associated with the IP interface.

Parameters

virtual-router-id

The virtual router ID for the IP interface expressed as a decimal integer.

Values 1 to 255

owner

Keyword used to identify this virtual router instance as owning the virtual router IP addresses. If the **owner** keyword is not specified at the time of *vrid* creation, the **vrrp backup** commands must be specified to define the virtual router IP addresses. The **owner** keyword is not required when entering the *vrid* for editing purposes. When created as **owner**, a *vrid* on an IP interface cannot have the **owner** parameter removed. The *vrid* must be deleted, and then recreated without the **owner** keyword, to remove ownership.

passive

Keyword used to identify this virtual router instance as **passive**, therefore owning the virtual router IP addresses. A **passive** *vrid* does not send or receive VRRP advertisement messages and is always in either the **master** state (if the interface is operationally up), or the **init** state (if the interface is operationally down). The **passive** keyword is not required when entering the *vrid* for editing purposes. When a *vrid* on an IP interface is created as **passive**, the parameter cannot be removed from the *vrid*. The *vrid* must be deleted, and then recreated without the **passive** keyword, to remove the parameter.

group-name

Specifies the name of the **oper-group**, up to 32 characters to establish the associated VRRP instance as a following instance to the specified operational group. As a result of this association, the VRRP instance state follows that of the VRRP instance (the lead instance) associated with the specified operation group.

Platforms

7705 SAR Gen 2

32.46 vsi-export

vsi-export

Syntax

vsi-export *policy-name* [*policy-name*]

no vsi-export

Context

[Tree] (config>service>vpls>bgp vsi-export)

Full Context

configure service vpls bgp vsi-export

Description

This command specifies the name of the VSI export policies to be used for BGP EVPN, BGP auto discovery, BGP VPLS, BGP VPWS, and BGP multi-homing if these features are configured in this VPLS service.

If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

The **no** form of this command removes the policy from the configuration.

Default

no vsi-export

Parameters

policy-name

Specifies up to five policy names, up to 32 characters.

Platforms

7705 SAR Gen 2

vsi-export

Syntax

vsi-export *policy-name* [*policy-name*]

no vsi-export

Context

[Tree] (config>service>epipe>bgp vsi-export)

Full Context

configure service epipe bgp vsi-export

Description

This command specifies the name of the VSI export policies to be used for BGP EVPN, BGP VPWS and BGP multi-homing if these features are configured in this Epipe service.

If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

The **no** form of this command removes the policy from the configuration.

Default

no vsi-export

Parameters

policy-name

Specifies up to five policy names, up to 32 characters.

Platforms

7705 SAR Gen 2

32.47 vsi-id

vsi-id

Syntax

vsi-id

Context

[Tree] (config>service>vpls>bgp-ad vsi-id)

Full Context

configure service vpls bgp-ad vsi-id

Description

Commands in this context configure the Virtual Switch Instance Identifier (VSI-ID).

Platforms

7705 SAR Gen 2

32.48 vsi-import

vsi-import

Syntax

vsi-import *policy-name* [*policy-name*]

no vsi-import

Context

[Tree] (config>service>vpls>bgp vsi-import)

Full Context

configure service vpls bgp vsi-import

Description

This command specifies the name of the VSI import policies to be used for BGP EVPN, BGP auto discovery, BGP VPLS, BGP VPWS, and BGP multi-homing if these features are configured in this VPLS service.

If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

The **no** form of this command removes the policy from the configuration.

Default

no vsi-import

Parameters

policy-name

Specifies up to five policy names, up to 32 characters.

Platforms

7705 SAR Gen 2

vsi-import

Syntax

vsi-import *policy-name* [*policy-name*]

no vsi-import

Context

[Tree] (config>service>epipe>bgp vsi-import)

Full Context

configure service epipe bgp vsi-import

Description

This command specifies the name of the VSI import policies to be used for BGP EVPN, BGP VPWS and BGP multi-homing if these features are configured in this Epipe service.

If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

The **no** form of this command removes the policy from the configuration.

Default

no vsi-import

Parameters

policy-name

Specifies up to five policy names, up to 32 characters.

Platforms

7705 SAR Gen 2

33 w Commands

33.1 wait

wait

Syntax

wait *seconds*

Context

[Tree] (bof wait)

Full Context

bof wait

Description

This command configures a pause, in seconds, at the start of the boot process which allows system initialization to be interrupted at the console.

When system initialization is interrupted the operator is allowed to manually override the parameters defined in the boot option file (BOF).

Only one **wait** command can be defined in the BOF.

Default

wait 3

Parameters

seconds

Specifies the time to pause at the start of the boot process, in seconds.

Values 1 to10

Platforms

7705 SAR Gen 2

33.2 watermarks

watermarks

Syntax

watermarks high *percentage-high* **low** *percentage-low*

no watermarks

Context

[Tree] (config>router>nat>outside>pool watermarks)

[Tree] (config>service>vprn>nat>outside>pool watermarks)

Full Context

configure router nat outside pool watermarks

configure service vprn nat outside pool watermarks

Description

This command configures the watermarks for ports or port-block utilization in a NAT pool.

For CGN and L2-aware NAT, the port and port-block watermarks are used to monitor ports or port-block utilization in a NAT pool. High and low thresholds are configured in percentages of total available port-blocks or ports per protocol in a pool.

For LSN44 pools with flexible port allocation, the watermarks represent the port utilization per outside IP address if paired address pooling is enabled. If an arbitrary address pooling is enabled, then the watermarks represent the port utilization per pool.

For port-block based pools where flexible port allocation is disabled, the watermarks represent port-block occupancy per pool. These watermarks cover combined initial and extended port-blocks in a NAT pool. If the extended port blocks are not enabled in L2-aware NAT, then the L2-aware pool contains only the initial port-blocks.

For the system to generate those events, the NAT event ID is configured as follows:

- 2001 is enabled for L2-aware pools which do not use l2-aware port-block-extensions
- 2003 is enabled for port-blocks based LSN pools
- 2044 is enabled for L2-aware pools which use l2aware port-block-extension
- 2046 is enabled for port based LSN pools for flexible-port-allocation

Event IDs are enabled via log event control configuration.

Default

no watermarks

Parameters

high percentage-high

Specifies the high percentage.

Values 1 to 100

low percentage-low

Specifies the low percentage.

Values 0 to 99

Platforms

7705 SAR Gen 2

watermarks

Syntax

watermarks high *percentage-high* **low** *percentage-low*

no watermarks

Context

[Tree] (config>service>nat>nat-policy>port-limits watermarks)

Full Context

configure service nat nat-policy port-limits watermarks

Description

This command configures the port usage watermarks for the NAT policy.

Default

no watermarks

Parameters

percentage-high

Specifies the high percentage.

Values 1 to 100

percentage-low

Specifies the low percentage.

Values 0 to 99

Platforms

7705 SAR Gen 2

watermarks

Syntax

watermarks high *percentage-high* **low** *percentage-low*

no watermarks

Context

[Tree] (config>service>nat>nat-policy>session-limits watermarks)

Full Context

configure service nat nat-policy session-limits watermarks

Description

This command configures the session watermarks for the NAT or residential firewall policy.

Default

no watermarks

Parameters

percentage-high

Specifies the high percentage.

Values 1 to 100

percentage-low

Specifies the low percentage.

Values 0 to 99

Platforms

7705 SAR Gen 2

33.3 weekday

weekday

Syntax

weekday {*weekday-number* [..*weekday-number*] | *day-name* [..*day-name*] | **all**}

no weekday

Context

[Tree] (config>system>cron>sched weekday)

Full Context

configure system cron schedule weekday

Description

This command specifies which days of the week that the schedule will fire on. Multiple days of the week can be specified. When multiple days are configured, each of them will cause the schedule to occur. If a weekday is configured without configuring the month, weekday, day-of-month, and minute, the event will not execute.

Using the **weekday** command as well as the **day-of month** command will cause the script to run twice. For example, consider that today is Monday January 1. If Tuesday January 5 is configured, the script will run on Tuesday (tomorrow) as well as January 5 (Friday).

The **no** form of this command removes the specified weekday from the configuration.

Default

no weekday

Parameters

weekday-number

Specifies a weekday number.

Values 1 to 7 (maximum 7 weekday-numbers)

day-name

Specifies a day by name.

Values sunday, monday, tuesday, wednesday, thursday, friday, saturday (maximum 7 weekday names)

all

Specifies all days of the week.

Platforms

7705 SAR Gen 2

33.4 weight

weight

Syntax

weight *weight*

no weight

Context

[Tree] (conf>router>segment-routing>sr-policies>policy>seg-list weight)

Full Context

configure router segment-routing sr-policies static-policy segment-list weight

Description

This command associates a weight value with a segment list of a statically-defined segment routing policy to achieve weighted ECMP behavior. Weight is an optional parameter.

When any segment-list in the active policy has a weight greater than 1, traffic matching the policy is load-balanced across the segment lists according to their relative weight values.

The **no** form of this command reverts to the default value.

Default

weight 1

Parameters

weight

Specifies the weight value.

Values 1 to 4294967295

Platforms

7705 SAR Gen 2

33.5 weight-down

weight-down

Syntax

[**no**] **weight-down** *lag-ports-down-weight*

Context

[Tree] (config>vrrp>policy>priority-event>lag-port-down weight-down)

Full Context

configure vrrp policy priority-event lag-port-down weight-down

Description

This command creates a context to configure an event set threshold within a lag-port-down priority control event. The weight-down command defines a sub-node within the lag-port-down event and is uniquely identified with the lag-ports-down-weight parameter. Each weight-down node within the same lag-port-down event node must have a unique lag-ports-down-weight value. Each weight-down node has its own priority command that takes effect whenever that node represents the current threshold. A single LAG can use either weight-based (**weight-down**) or port-based (**number-down**) thresholds. The weight-based thresholds are required for correct operation on mixed port-speed LAGs, but can be used for non mixed port-speed LAGs as well. The weights for the **weight-down** node are normalized from the **hash-weight** values of the LAG member ports to fit a 1 to 64 range for 64-link capable LAGs and a 1 to 32 range for other LAGs.

The total number of sub-nodes (uniquely identified by the lag-ports-down-weight parameter) allowed in the system is 2048.

A **weight-down** node is not required for each possible number of ports that could be down. The active threshold is always the closest lower threshold.

The **no** form of the command deletes the event set threshold. The threshold may be removed at any time. If the removed threshold is the current active threshold, the event set thresholds must be re-evaluated after removal.

Default

no weight-down

Parameters

lag-ports-down-weight

The total weight of LAG ports down to create a set event threshold. This is the active threshold when the weight of down ports in the LAG equals or exceeds *lag-ports-down-weight*, but does not equal or exceed the next highest configured *lag-ports-down-weight*.

Values 1 to 64 (for 64-link capable LAGs)

1 to 32 (for other LAGs)

Platforms

7705 SAR Gen 2

33.6 weighted-ecmp

weighted-ecmp

Syntax

[**no**] **weighted-ecmp**

Context

[Tree] (config>service>vprn>bgp-ipvpn>mpls>auto-bind-tunnel weighted-ecmp)

[Tree] (config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel weighted-ecmp)

[Tree] (config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel weighted-ecmp)

Full Context

configure service vprn bgp-ipvpn mpls auto-bind-tunnel weighted-ecmp

configure service epipe bgp-evpn mpls auto-bind-tunnel weighted-ecmp

configure service vpls bgp-evpn mpls auto-bind-tunnel weighted-ecmp

Description

This command enables weighted ECMP for packets using tunnels that a VPRN, VPLS, or Epipe automatically binds to. When weighted ECMP is enabled, packets are sprayed across LSPs in the ECMP according to the outcome of the hash algorithm and the configured load balancing weight of each LSP.

The **no** form of this command disables weighted ECMP for next hop tunnel selection.

Default

no weighted-ecmp

Platforms

7705 SAR Gen 2

weighted-ecmp

Syntax

[**no**] **weighted-ecmp**

Context

[Tree] (configure>service>vpls>bgp-evpn>ip-route-link-bw weighted-ecmp)

[Tree] (configure>service>vprn>bgp-evpn>mpls>evpn-link-bw weighted-ecmp)

Full Context

configure service vpls bgp-evpn ip-route-link-bandwidth weighted-ecmp

configure service vprn bgp-evpn mpls evpn-link-bandwidth weighted-ecmp

Description

This command enables the processing of the EVPN link bandwidth extended community when installing an ECMP set for an EVPN IP prefix route in the VPRN route table.

Flows to an IP prefix received with a weight and a zero ESI value are sprayed according to the weight. If the EVPN IP prefix route received with the weight has a non-zero ESI, the weight is divided into the number of PEs attached to the Ethernet Segment (and rounded up if the result is not an integer).

This command also enables the weighted ECMP functionality for BGP CEs where the weight is configured with the **evpn-link-bandwidth add-to-received-bgp** command.

The **no** form of this command disables EVPN link bandwidth extended community.

Default

no weighted-ecmp

Platforms

7705 SAR Gen 2

weighted-ecmp

Syntax

weighted-ecmp [**strict**]

no weighted-ecmp

Context

[Tree] (config>service>vprn weighted-ecmp)

Full Context

configure service vprn weighted-ecmp

Description

This command enables weighted load-balancing for IS-IS, OSPF, and static ECMP routes in the VPRN instance. Weighted ECMP can be performed when all next hops are configured with non-zero load-balancing weights. Weighted ECMP support for IS-IS, OSPF, and static ECMP routes applies to both IPv4 and IPv6.

The **no** form of this command restores regular ECMP spraying of packets to IS-IS, OSPF and static route destinations.

Default

no weighted-ecmp

Parameters

strict

Enables strict enforcement for a load balancing weight to be configured on each interface withing a wECMP interface bundle before the interface is taken into wECMP operation. However, when **strict** enforcement is not enabled, then, when **load-balancing-weight** is not configured on one or more interfaces within the wECMP interface bundle, the wECMP load-balancing falls back to classic ECMP operation and equally share the traffic load across the ECMP interface bundle. A special case is when none of the available paths or next-hops have a load balancing weight associated. Then, the load balancing falls back to classic ECMP.

Strict load balancing is only applied on IS-IS, OSPF, and static route entries.

Platforms

7705 SAR Gen 2

weighted-ecmp

Syntax

weighted-ecmp

Context

[Tree] (config>service>vprn weighted-ecmp)

Full Context

configure service vprn weighted-ecmp

Description

Note: This command is no longer supported and will be removed in a future release.

Platforms

7705 SAR Gen 2

weighted-ecmp

Syntax

[**no**] **weighted-ecmp**

Context

[Tree] (config>router>ldp weighted-ecmp)

Full Context

configure router ldp weighted-ecmp

Description

This command enables weighted ECMP on LDP using RSVP LSPs or SR-TE LSPs. LDP labeled packets are sprayed across the RSVP or SR-TE LSP ECMP in proportion to the configured **load-balancing-weight** of LSPs.

The **no** form of this command removes weighted ECMP.

Default

no weighted-ecmp

Platforms

7705 SAR Gen 2

weighted-ecmp

Syntax

weighted-ecmp [**strict**]

no weighted-ecmp

Context

[Tree] (config>router weighted-ecmp)

Full Context

configure router weighted-ecmp

Description

This command enables the weighted load-balancing, or weighted ECMP, over MPLS LSP.

When this command is enabled, packets of IGP, BGP, and static route prefixes resolved to a set of ECMP tunnel next-hops are sprayed proportionally to the weights configured for each MPLS LSP in the ECMP set.

Weighted load-balancing over MPLS LSP is supported in the following forwarding contexts:

- IGP prefix resolved to IGP shortcuts in RTM (**igp-shortcut** or **advertise-tunnel-link** enabled in the IGP instance).
- BGP prefix with the BGP next-hop resolved to IGP shortcuts in RTM (**igp-shortcut** or **advertise-tunnel-link** enabled in the IGP instance).
- Static route prefix resolved to an indirect next-hop, which itself is resolved to a set of equal-metric MPLS LSPs in TTM. The user can allow automatic selection or specify the names of the equal-metric MPLS LSPs in TTM to be used in the ECMP set.
- Static route prefix resolved to an indirect next-hop, which is resolved to IGP shortcuts in RTM.
- BGP prefix with a BGP next-hop resolved to a static route, which resolves to a set of tunnel next-hops toward an indirect next-hop in RTM or TTM.
- BGP prefix resolving to another BGP prefix, whose next-hop is resolved to a set of ECMP tunnel next-hops with a static route in RTM or TTM or to IGP shortcuts in RTM.

IGP computes the normalized weight for each prefix tunnel next-hop. IGP updates the route in RTM with the set of tunnel next-hops and normalized weights. RTM downloads the information to IOM for inclusion in the FIB.

If one or more LSPs in the ECMP set of a prefix do not have a weight configured, the regular ECMP spraying for the prefix will be performed.

The weight assigned to an LSP impacts only the forwarding decision, not the routing decision. In other words, it does not change the selection of the set of ECMP tunnel next-hops of a prefix when more next-hops exist than the value of the router **ecmp** option. Once the set of tunnel next-hops is selected, the LSP weight is used to modulate the amount of packets forwarded over each next-hop. It also does not change the hash routine, but only the spraying of the flows over the tunnel next-hops is modified to reflect the normalized weight of each tunnel next-hop.

The **no** form of this command resumes regular ECMP spraying of packets of IGP, BGP, and static route prefixes over MPLS LSP.

Default

no weighted-ecmp

Parameters

strict

Enables strict enforcement for a load balancing weight to be configured on each interface withing a wECMP interface bundle before the interface is taken into wECMP operation. However, when **strict** enforcement is not enabled, then, when **load-balancing-weight** is not configured on one or more interfaces within the wECMP interface bundle, the wECMP load-balancing falls back to classic ECMP operation and equally share the traffic load across the ECMP interface bundle. A special case is when none of the available paths or next-hops have a load balancing weight associated. Then, the load balancing falls back to classic ECMP.

Strict load balancing is only applied on IS-IS, OSPF, and static route entries.

Platforms

7705 SAR Gen 2

weighted-ecmp

Syntax

[**no**] **weighted-ecmp**

Context

[Tree] (config>service>sdp weighted-ecmp)

Full Context

configure service sdp weighted-ecmp

Description

This command enables weighted ECMP on an SDP. When weighted ECMP is enabled, packets from services using the SDP are sprayed across LSPs in the ECMP set to the SDP far end according to the outcome of the hash algorithm and the configured load-balancing weight of each LSP.

The **no** version of this command disables weighted ECMP for next-hop tunnel selection.

Default

no weighted-ecmp

Platforms

7705 SAR Gen 2

weighted-ecmp

Syntax

[**no**] **weighted-ecmp**

Context

[Tree] (config>router>bgp>next-hop-res weighted-ecmp)

Full Context

configure router bgp next-hop-resolution weighted-ecmp

Description

This command enables weighted ECMP for next-hop tunnel selection for 6PE. When weighted ECMP is enabled, the RSVP-TE tunnel used to forward 6PE packets to the ECMP next hop is chosen according to the outcome of the hash on the packet at the normalized load-balancing weight of the tunnel.

The **no** version of this command disables weighted ECMP for next-hop tunnel selection for 6PE.

Default

no weighted-ecmp

Platforms

7705 SAR Gen 2

33.7 wide-metrics-only

wide-metrics-only

Syntax

[**no**] **wide-metrics-only**

Context

[Tree] (config>service>vprn>isis>level wide-metrics-only)

Full Context

configure service vprn isis level wide-metrics-only

Description

This command enables the exclusive use of wide metrics in the LSPs for the level number. Narrow metrics can have values between 1 and 63. IS-IS can generate two TLVs, one for the adjacency and one for the

IP prefix. In order to support traffic engineering, wider metrics are required. When wide metrics are used, a second pair of TLVs are added, again, one for the adjacency and one for the IP prefix.

By default, both sets of TLVs are generated. When wide-metrics-only is configured, IS-IS only generates the pair of TLVs with wide metrics for that level.

The **no** form of this command reverts to the default value.

Platforms

7705 SAR Gen 2

wide-metrics-only

Syntax

[**no**] **wide-metrics-only**

Context

[Tree] (config>router>isis>level wide-metrics-only)

Full Context

configure router isis level wide-metrics-only

Description

This command enables the exclusive use of wide metrics in the LSPs for the level number. Narrow metrics can have values between 1 and 63. IS-IS can generate two TLVs, one for the adjacency and one for the IP prefix. In order to support traffic engineering, wider metrics are required. When wide metrics are used, a second pair of TLVs are added, again, one for the adjacency and one for the IP prefix.

By default, both sets of TLVs are generated. When wide-metrics-only is configured, IS-IS only generates the pair of TLVs with wide metrics for that level.

The **no** form of this command reverts to the default value.

Default

no wide-metrics-only

Platforms

7705 SAR Gen 2

33.8 width

width

Syntax
width *width*

Context
[Tree] (environment>terminal width)

Full Context
environment terminal width

Description
This command determines display terminal width.

Default
width 80

Parameters
width
Sets the width of the display terminal.

Values 1 to 512

Platforms
7705 SAR Gen 2

width

Syntax
width *width*

Context
[Tree] (config>system>management-interface>cli>md-cli>environment>console width)

Full Context
configure system management-interface cli md-cli environment console width

Description

This command configures the set number of columns displayed on the console.

Default

width 80

Parameters

width

Specifies the number of columns displayed in the console window.

Values 80 to 512

Platforms

7705 SAR Gen 2

33.9 window-size

window-size

Syntax

window-size *seconds*

no window-size

Context

[Tree] (config>port>ethernet>crc-monitor window-size)

Full Context

configure port ethernet crc-monitor window-size

Description

This command specifies sliding window size over which the Ethernet frames are sampled to detect signal fail or signal degrade conditions. The command is used jointly with the sf-threshold and the sd-threshold to configure the sliding window size.

The **no** version of this command reverts to the default value of 10 seconds.

Default

no window-size

Parameters

seconds

The size of the sliding window in seconds over which the errors are measured.

Values 5 to 60

Platforms
7705 SAR Gen 2

33.10 wrap-around

wrap-around

Syntax
[**no**] **wrap-around**

Context
[Tree] (config>filter>log wrap-around)

Full Context
configure filter log wrap-around

Description
This command configures a memory filter log to log until full or to store the most recent log entries (circular buffer).

Specifying **wrap-around** configures the memory filter log to store the most recent filter log entries (circular buffer). When the log is full, the oldest filter log entries are overwritten with new entries.

The **no** form of the command configures the memory filter log to accept filter log entries until full. When the memory filter log is full, filter logging for the log filter ID ceases.

Default
wrap-around

Platforms
7705 SAR Gen 2

33.11 wred-queue

wred-queue

Syntax
wred-queue [**policy** *slope-policy-name*] [**mode** *mode*] [**slope-usage** *slope-usage*]

no wred-queue

Context

[Tree] (config>qos>sap-egress>queue wred-queue)

Full Context

configure qos sap-egress queue wred-queue

Description

This command allows the configuration of WRED per queue with the following options:

- Native hardware WRED

  This uses the hardware per queue WRED capabilities of FP3- and higher-based hardware and is configured with the **native** keyword.

- Pool per queue WRED

  This implements each queue in its own pool and uses the WRED capabilities of the pool to provide WRED per queue. This is configured with the **pool-per-queue** keyword.

Native Hardware WRED

When the **wred-queue mode native** command is configured, the queue uses the WRED capabilities of FP3- and higher-based hardware. In this case, the out-of-profile and exceed-profile traffic map to the low and exceed WRED slopes specified within the slope policy, and the inplus-profile and in-profile traffic uses the MBS drop tail; this requires the **slope-usage** to be configured as **exceed-low**. The instantaneous queue depth is compared against the low and exceed slopes so the time average factor in the slope policy is ignored.

When a policy is not explicitly defined, the default slope policy is used.

When **native** mode is enabled for a queue, the **pool** and **drop-tail** commands are ignored; traffic mapped to a slope that is shut down will use the MBS drop tail.

This is only supported on FP3 hardware.

The **no** form of this command restores the queue default congestion control behavior to the queue.

Pool-per-queue WRED

When the **wred-queue mode pool-per-queue** command is defined and the queue ID is created, a buffer pool is created specifically for the queue and the queue obtains all buffers from that pool. The size of the pool is the same as the size of the queue. In this manner, the WRED slopes that operate based on the pool's buffer utilization are also reacting to the congestion depth of the queue.

The size of the buffer pool is dictated by the queue's MBS parameter. The size of the reserved CBS portion of the buffer pool is dictated by the queue's CBS parameter. The provisioning characteristics of the **mbs** and **cbs** commands are not changed.

In the case where this is applied with WRED queue support shutdown (**config>card>fp>egress>wred-queue-control>shutdown**), the queue will continue to map to its default pool. If the **no shutdown** command is executed in the **wred-queue-control** context, the queue is automatically moved to its own WRED pool.

Each pool created for a queue using the **wred-queue** command shares buffers with all other **wred-queue** enabled queues on the same forwarding plane. The WRED pool buffer management behavior is defined within the **config>card>fp>egress>wred-queue-control** CLI context.

The WRED slopes within the pool are defined by the slope policy associated with the queue. When a policy is not explicitly defined, the default slope policy is used. The slope policy enables, disables, and defines the relative geometry of the highplus, high, low, and exceed WRED slopes in the pool. The policy also specifies the time average factor used by the pool when calculating the weighted average pool depth.

As packets attempt to enter the egress queue, they are associated with the highplus, high, low, or exceed WRED slope based on the packet's profile. If the packet is inplus-profile, the highplus slope is used. If the packet is in-profile, the high slope is used. If the packet is out-of-profile, the low slope is used. If the packet is exceed-profile, the exceed slope is used. This mapping of packet profile to slope is enabled using the **slope-usage default** parameter. Each WRED slope performs a probability discard based on the current weighted average pool depth.

When **wred-queue** is enabled for a SAP egress queue, the queue **pool** and **drop-tail** commands are ignored; traffic mapped to a slope that is shut down will use the MBS drop tail.

The resource usage for the WRED queue pool-per-queue per forwarding plane can be seen in the **tools dump resource-usage card** [*slot-num*] **fp** [*fp-number*] output under *Dynamic Q2 Wred Pools*.

The **no** form of this command restores the generic buffer pool behavior to the queue. The WRED pool is removed from the system. The queue will be moved to the default buffer pool. The queue then uses the default congestion control behavior.

Default

no wred-queue

Parameters

slope-policy-name

Specifies an existing slope policy that is used to override the default WRED slope policy.

mode

Specifies whether the WRED per queue is using the native FP3- and higher-based hardware WRED capabilities or pool per queue.

Values **native** — uses the hardware per queue WRED capabilities of the FP3- and higher-based hardware and requires **slope-usage exceed-low**.

pool-per-queue — each queue uses its own pool and the WRED capabilities of the pool to provide WRED per queue. This is supported on both FP2- and higher-based hardware and requires **slope-usage default**.

Default native

slope-usage

Specifies congestion control to be used.

Values **default** — maps the inplus, in, out, and exceed-profile traffic to the highplus, high, low, and exceed WRED slopes, respectively; this is only supported for **pool-per-queue** mode.

exceed-low — maps the out and exceed-profile traffic to the low and exceed WRED slopes with the inplus and in-profile traffic using the MBS drop tail. This is only supported for **native** mode.

Default exceed-low

Platforms

7705 SAR Gen 2

wred-queue

Syntax

wred-queue [**policy** *slope-policy-name*] [**mode** {**native** | **pool-per-queue**}] [**slope-usage** {**default** | **exceed-low**}]

no wred-queue

Context

[Tree] (config>qos>qgrps>egr>qgrp>queue wred-queue)

Full Context

configure qos queue-group-templates egress queue-group queue wred-queue

Description

This command allows the configuration of WRED per queue with the following options:

- Native hardware WRED

  This uses the hardware per queue WRED capabilities of FP3- and higher-based hardware and is configured with the **native** keyword.

- Pool per queue WRED

  This implements each queue in its own pool and uses the WRED capabilities of the pool to provide WRED per queue. This is configured with the **pool-per-queue** keyword.

Native Hardware WRED

When the **wred-queue mode native** command is configured, the queue uses the WRED capabilities of FP3- and higher-based hardware. In this case, the out and exceed-profile traffic map to the low and exceed WRED slopes specified within the slope policy, and the inplus and in-profile traffic uses the MBS drop tail; this requires the **slope-usage** to be configured as **exceed-low**. The instantaneous queue depth is compared against the low and exceed slopes so the time average factor in the slope policy is ignored.

When a policy is not explicitly defined, the default slope policy is used.

When **native** mode is enabled for a queue, the drop-tail commands are ignored; traffic mapped to a slope that is shut down will use the MBS drop tail.

Native hardware WRED is supported on FP3- and higher-based hardware and is ignored on FP2 hardware.

The **no** form of this command restores the queue default congestion control behavior to the queue.

Pool-per-queue WRED

When the **wred-queue mode pool-per-queue** command is defined and the queue ID is created, a buffer pool is created specifically for the queue and the queue obtains all buffers from that pool. The size of the

pool is the same as the size of the queue. In this manner, the WRED slopes that operate based on the pool's buffer utilization are also reacting to the congestion depth of the queue.

The size of the buffer pool is dictated by the queue's **mbs** parameter. The size of the reserved CBS portion of the buffer pool is dictated by the queue's **cbs** parameter. The provisioning characteristics of the **mbs** and **cbs** commands are not changed.

In the case where this is applied with WRED queue support shut down (**config>card>fp>egress>wred-queue-control>shutdown**), the queue will continue to map to its default pool. If the **no shutdown** command is executed in the **wred-queue-control** context, the queue will be automatically moved to its own WRED pool.

Each pool created for a queue using the **wred-queue** command shares buffers with all other **wred-queue** enabled queues on the same forwarding plane. The WRED pool buffer management behavior is defined within the **config>card>fp>egress>wred-queue-control** CLI context.

The WRED slopes within the pool are defined by the slope policy associated with the queue. When a policy is not explicitly defined, the default slope policy is used. The slope policy enables, disables, and defines the relative geometry of the highplus, high, low, and exceed WRED slopes in the pool. The policy also specifies the time average factor used by the pool when calculating the weighted average pool depth.

As packets attempt to enter the egress queue, they are associated with the highplus, high, low, or exceed WRED slope based on the packet's profile. If the packet is inplus-profile, the highplus slope is used. If the packet is in-profile, the high slope is used. If the packet is out-of-profile, the low slope is used, and if the packet is exceed-profile, the exceed slope is used. This mapping of packet profile to slope is enabled using the **slope-usage default** parameter. Each WRED slope performs a probability discard based on the current weighted average pool depth.

When **wred-queue** is enabled for an egress queue group queue, the queue pool and drop-tail commands are ignored; traffic mapped to a slope that is shut down will use the MBS drop tail.

The resource usage for the wred-queue pool-per-queue per forwarding plane can be seen in the **tools dump resource-usage card** [*slot-num*] **fp** [*fp-number*] output under *Dynamic Q2 Wred Pools*.

The **no** form of this command restores the generic buffer pool behavior to the queue. The WRED pool is removed from the system. The queue will be moved to the default buffer pool. The queue then uses the default congestion control behavior.

Default

no wred-queue

Parameters

slope-policy-name

Specifies an existing slope policy that is used to override the default WRED slope policy.

mode {native | pool-per-queue}

Specifies whether the WRED per queue is using the native FP3- and higher-based hardware WRED capabilities or pool per queue.

Values **native** - Each queue uses the hardware per queue WRED capabilities of the FP3- and higher-based hardware and requires **slope-usage exceed-low**.

pool-per-queue - Each queue uses its own pool and the WRED capabilities of the pool to provide WRED per queue. This requires **slope-usage default**.

Default native

slope-usage {default | exceed-low}

Specifies the type of congestion control to be used.

Values **default** - Maps the inplus-profile, in-profile, out-of-profile, and exceed-profile traffic to the highplus, high, low, and exceed WRED slopes, respectively; this is only supported for **pool-per-queue** mode.

exceed-low - Maps the out-of-profile and exceed-profile traffic to the low and exceed WRED slopes, with the inplus-profile and in-profile traffic using the MBS drop tail. This option is only supported for **native** mode.

Default exceed-low

Platforms

7705 SAR Gen 2

33.12 write

write

Syntax

write {*user-name* | *broadcast*} *message*

Context

[Tree] (write)

Full Context

write

Description

This command sends a console message to a specific user or to all users with active console sessions.

Parameters

user-name

Specifies the name of a user, up to 32 characters, with an active console session to which to send a console message.

Values any valid CLI username

broadcast

Sends the *message-string* to all users logged into the router.

message

Specifies the message string to send. Allowed values are any string, up to 256 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, ?, space), the entire string must be enclosed within double quotes.

Platforms

7705 SAR Gen 2

33.13 write-algorithm

write-algorithm

Syntax

write-algorithm {**hash** | **hash2** | **custom** | **cleartext**}

no write-algorithm

Context

[Tree] (config>system>security>management-interface>classic-cli write-algorithm)

Full Context

configure system security management-interface classic-cli write-algorithm

Description

This command specifies how encrypted configuration secrets are output from the system. For example, how encrypted secrets are displayed in the output of the info command, and how they are written to the saved configuration file.

The **no** form of this command reverts to the default value.

Default

write-algorithm hash2

Parameters

hash

Specifies hash. Use this option to transport a phrase between modules and nodes. In this case the read-algorithm should be **hash** as well.

hash2

Specifies hash2 which is module-specific.

custom

Specifies the custom encryption to management interface.

cleartext

Specifies that the phrase is displayed as cleartext everywhere.

Platforms

7705 SAR Gen 2

34 x Commands

34.1 xc

xc

Syntax

xc [**detail**]

no xc

Context

[Tree] (debug>router>mpls>event xc)

Full Context

debug router mpls event xc

Description

This command debugs cross connect events.

The **no** form of the command disables the debugging.

Parameters

detail

Displays detailed information about cross connect events.

Platforms

7705 SAR Gen 2

34.2 xgig

xgig

Syntax

xgig {**lan** | **wan**}

Context

[Tree] (config>port>ethernet xgig)

Full Context

configure port ethernet xgig

Description

This command configures a 10 Gb/s interface to be in Local or Wide Area Network (LAN or WAN) mode. When configuring the port to be in WAN mode certain SONET/SDH parameters can be changed to reflect the SONET/SDH requirements for this port.

When the port is configured for LAN mode, all SONET/SDH parameters are pre-determined and not configurable.

Default

xgig lan

Parameters

lan

Sets the port to operate in LAN mode.

wan

Sets the port to operate in WAN mode.

Platforms

7705 SAR Gen 2

35 y Commands

35.1 yang-modules

yang-modules

Syntax

yang-modules

Context

[Tree] (config>system>management-interface yang-modules)

Full Context

configure system management-interface yang-modules

Description

Commands in this context configure YANG module parameters.

The **yang-modules** settings affect the data sent in a NETCONF <hello>, data populated in the RFC 6022 /netconf-state/schemas list, data returned in a <get-schema> request, and data populated in the RFC 8525 /yang-library. See the *7705 SAR Gen 2 System Management Guide*, sections "NETCONF Monitoring" and "YANG Library" for more information.

Platforms

7705 SAR Gen 2

36 z Commands

36.1 zone

zone

Syntax

zone {*std-zone-name* | *non-std-zone-name*} [*hh* [*:mm*]]

no zone

Context

[Tree] (config>system>time zone)

Full Context

configure system time zone

Description

This command sets the time zone and/or time zone offset for the device.

The SR OS supports system-defined and user-defined time zones. The system-defined time zones are listed in the Time Zones section.

For user-defined time zones, the zone and the UTC offset must be specified.

The **no** form of the command reverts to the default of Coordinated Universal Time (UTC). If the time zone in use was a user-defined time zone, the time zone will be deleted. If a **dst-zone** command has been configured that references the zone, the summer commands must be deleted before the zone can be reset to UTC.

Default

zone UTC 00

Parameters

std-zone-name

Specifies the standard time zone name. The standard name must be a system-defined zone in the Time Zones section. For zone names in the table that have an implicit summer time setting, for example MDT for Mountain Daylight Saving Time, the remaining **start-date**, **end-date** and **offset** parameters do not need to be provided unless it is necessary to override the system defaults for the time zone.

For system-defined time zones, a different offset cannot be specified. If a new time zone is needed with a different offset, the user must create a new time zone. Note that some system-defined time zones have implicit summer time settings which causes the

switchover to summer time to occur automatically; configuring the **config>system>time dst-zone** command is not required.

A user-defined time zone name is case-sensitive and can be up to 5 characters in length.

Values A user-defined value can be up to 5 characters or one of the following values: GMT, WET, CET, EET, EEST, MSK, MSD, AST, NST, EST, CST, MST, PST, HST, AKST, AWST, ACST, AEST, NZST, UTC

non-std-zone-name

Specifies the non-standard time zone name. The name can be up to 5 characters.

hh [:mm]

Specifies the hours and minutes offset from UTC time, expressed as integers. Some time zones do not have an offset that is an integral number of hours. In these instances, the *minutes-offset* must be specified. For example, the time zone in Pirlanngimpi, Australia UTC + 9.5 hours.

Values hours: -11 to 12 minutes: 0 to 59

Default hours: 0 minutes: 0

Platforms

7705 SAR Gen 2

Customer document and product support

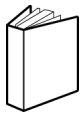

Customer documentation

Customer documentation welcome page

Technical support

Product support portal

Documentation feedback

Customer documentation feedback