



7210 Service Access System

Release 24.9.R1

7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide

3HE 20139 AAAB TQZZA
Edition: 01
September 2024

© 2024 Nokia.

Use subject to Terms available at: www.nokia.com/terms.

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.

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

© 2024 Nokia.

Table of contents

List of tables	11
List of figures	15
1 Getting started	16
1.1 About this guide.....	16
1.1.1 Document structure and content.....	17
1.2 7210 SAS modes of operation.....	17
1.3 7210 SAS port modes.....	19
1.4 7210 SAS router configuration process.....	21
1.5 Conventions.....	22
1.5.1 Precautionary and information messages.....	22
1.5.2 Options or substeps in procedures and sequential workflows.....	22
2 IP router configuration	24
2.1 Configuring IP router parameters.....	24
2.1.1 Interfaces.....	24
2.1.1.1 Secondary IPv4 addresses.....	24
2.1.1.2 Network interface.....	25
2.1.2 System interface.....	25
2.1.3 Router ID.....	25
2.1.4 Autonomous systems.....	25
2.1.5 Proxy ARP.....	26
2.1.6 Internet Protocol versions.....	26
2.1.6.1 IPv6 applications.....	28
2.1.6.2 IPv6 Provider Edge router over MPLS (6PE).....	30
2.1.7 Bidirectional Forwarding Detection.....	31
2.1.7.1 BFD control packet.....	32
2.1.7.2 Control packet format.....	32
2.1.7.3 Echo support.....	34
2.1.7.4 BFD IPv4 support on 7210 SAS platforms.....	34
2.1.7.5 BFD IPv6 support on 7210 SAS platforms.....	35
2.1.8 IGP-LDP and static route-LDP synchronization.....	36
2.1.9 IP fragmentation.....	37

2.2	Process overview.....	38
2.3	Configuration notes.....	38
2.4	Configuring an IP router with CLI.....	39
2.4.1	Router configuration overview.....	39
2.4.1.1	System interface.....	39
2.4.1.2	Network interface.....	40
2.4.2	Basic configuration.....	40
2.4.3	Common configuration tasks.....	40
2.4.3.1	Configuring a system name.....	40
2.4.3.2	Configuring interfaces.....	41
2.4.3.3	Configuring router advertisement.....	43
2.4.3.4	Configuring proxy ARP.....	44
2.4.3.5	ECMP considerations.....	45
2.4.3.6	Deriving the router ID.....	46
2.4.3.7	Configuring an autonomous system.....	47
2.4.3.8	Configuring static routes.....	47
2.4.4	Service management tasks.....	48
2.4.4.1	Changing the system name.....	48
2.4.4.2	Modifying interface parameters.....	48
2.4.4.3	Deleting a logical IP interface.....	49
2.5	IP router command reference.....	50
2.5.1	Command hierarchies.....	50
2.5.1.1	Configuration commands.....	50
2.5.1.2	Show commands.....	54
2.5.1.3	Clear commands.....	55
2.5.1.4	Debug commands.....	55
2.5.2	Command descriptions.....	55
2.5.2.1	Configuration commands.....	55
2.5.2.2	Show commands.....	126
2.5.2.3	Clear commands.....	160
2.5.2.4	Debug commands.....	166
3	VRRP.....	172
3.1	VRRP overview.....	172
3.1.1	VRRP components.....	173
3.1.1.1	Virtual router.....	173

3.1.1.2	IP address owner.....	173
3.1.1.3	Primary and secondary IP addresses.....	174
3.1.1.4	Virtual router master state.....	174
3.1.1.5	Virtual router backup.....	174
3.1.1.6	Owner and non-owner VRRP.....	174
3.1.2	Configurable parameters.....	175
3.1.2.1	Virtual router ID (VRID).....	175
3.1.2.2	Priority.....	175
3.1.2.3	IP addresses.....	176
3.1.2.4	Message interval and master inheritance.....	176
3.1.2.5	Skew time.....	177
3.1.2.6	Master down interval.....	177
3.1.2.7	Preempt mode.....	177
3.1.2.8	VRRP message authentication.....	178
3.1.2.9	Authentication data.....	179
3.1.2.10	Virtual MAC address.....	180
3.1.2.11	VRRP advertisement message IP address list verification.....	180
3.1.2.12	IPv6 virtual router instance operationally up.....	180
3.1.2.13	Policies.....	180
3.2	VRRP priority control policies.....	181
3.2.1	VRRP virtual router policy constraints.....	181
3.2.2	VRRP virtual router instance base priority.....	181
3.2.3	VRRP priority control policy delta in-use priority limit.....	181
3.2.4	VRRP priority control policy priority events.....	182
3.2.4.1	Priority event hold-set timers.....	182
3.2.4.2	Port down priority event.....	182
3.2.4.3	LAG degrade priority event.....	182
3.2.4.4	Host unreachable priority event.....	184
3.2.4.5	Route unknown priority event.....	184
3.3	VRRP non-owner accessibility.....	184
3.3.1	Non-owner access ping reply.....	185
3.3.2	Non-owner access Telnet.....	185
3.3.3	Non-owner access SSH.....	185
3.4	VRRP configuration process overview.....	185
3.5	Configuration notes.....	186
3.5.1	General.....	186

3.6	Configuring VRRP with CLI.....	187
3.7	VRRP configuration overview.....	187
3.7.1	Preconfiguration requirements.....	187
3.8	Basic VRRP configurations.....	187
3.8.1	VRRP policy.....	187
3.8.2	VRRP IES service parameters.....	188
3.8.3	VRRP router interface parameters.....	189
3.9	Common configuration tasks.....	190
3.9.1	Creating interface parameters.....	190
3.10	Configuring VRRP policy components.....	191
3.10.1	Configuring service VRRP parameters.....	191
3.10.1.1	Non-owner VRRP example.....	191
3.10.1.2	Owner service VRRP example.....	192
3.10.2	Configuring router interface VRRP parameters.....	192
3.10.2.1	Router interface VRRP non-owner.....	192
3.10.2.2	Router interface VRRP owner.....	193
3.11	VRRP configuration management tasks.....	193
3.11.1	Modifying a VRRP policy.....	193
3.11.2	Deleting a VRRP policy.....	194
3.11.3	Modifying service and interface VRRP parameters.....	194
3.11.3.1	Modifying non-owner parameters.....	194
3.11.3.2	Modifying owner parameters.....	194
3.11.3.3	Deleting VRRP on an interface or service.....	194
3.12	VRRP command reference.....	195
3.12.1	Command hierarchies.....	195
3.12.1.1	Configuration commands.....	195
3.12.1.2	Show commands.....	197
3.12.1.3	Monitor commands.....	198
3.12.1.4	Clear commands.....	198
3.12.1.5	Debug commands.....	198
3.12.2	Command descriptions.....	198
3.12.2.1	Configuration commands.....	198
3.12.2.2	Show commands.....	237
3.12.2.3	Monitor commands.....	250
3.12.2.4	Clear commands.....	252
3.12.2.5	Debug commands.....	254

4	Filter policies.....	256
4.1	Filter policy configuration overview.....	256
4.1.1	Service and network IP interface-based filtering.....	256
4.1.2	Filter policy entities.....	257
4.1.2.1	Applying filter policies.....	257
4.1.2.2	ACL on range SAPs.....	260
4.2	Creating and applying policies.....	261
4.2.1	Packet matching criteria.....	262
4.2.1.1	DSCP values.....	263
4.2.2	Ordering filter entries.....	266
4.2.3	Applying filters.....	267
4.2.3.1	Applying a filter to a SAP.....	267
4.2.3.2	Applying a filter to a network IP interface.....	268
4.3	Configuration notes.....	268
4.3.1	MAC filters.....	269
4.3.2	IP filters.....	269
4.3.3	IPv6 filters.....	270
4.3.3.1	Resource usage for ingress filter policies.....	270
4.3.3.2	Resource usage for egress filter policies.....	271
4.4	Configuring filter policies with CLI.....	273
4.5	Basic configuration.....	273
4.6	Common configuration tasks.....	275
4.6.1	Creating an IP filter policy.....	275
4.6.1.1	IP filter policy.....	275
4.6.1.2	IP filter entry.....	276
4.6.1.3	IP entry matching criteria.....	276
4.6.2	Creating an IPv6 filter policy.....	277
4.6.2.1	IPv6 filter policy.....	277
4.6.2.2	IPv6 filter entry.....	277
4.6.3	Creating a MAC filter policy.....	278
4.6.3.1	MAC filter policy.....	278
4.6.3.2	MAC filter entry.....	278
4.6.3.3	MAC entry matching criteria.....	279
4.6.3.4	Apply IP and MAC filter policies.....	279
4.6.3.5	Apply an IPv6 filter policy to VPLS.....	280

4.6.4	Applying filter policies to a network IP interface.....	280
4.6.4.1	Applying a filter policy to an IP interface.....	280
4.7	Filter management tasks.....	281
4.7.1	Renumbering filter policy entries.....	281
4.7.2	Modifying an IP filter policy.....	283
4.7.3	Modifying an IPv6 filter policy.....	284
4.7.4	Modifying a MAC filter policy.....	285
4.7.5	Detaching/deleting a filter policy.....	286
4.7.5.1	From an ingress SAP.....	286
4.7.5.2	From an egress SAP.....	286
4.7.5.3	From a network interface.....	286
4.7.5.4	From the filter configuration.....	287
4.7.6	Copying filter policies.....	287
4.8	Filter command reference.....	288
4.8.1	Command hierarchies.....	288
4.8.1.1	Configuration commands.....	288
4.8.1.2	Show commands.....	291
4.8.1.3	Clear commands.....	291
4.8.1.4	Monitor commands.....	291
4.8.2	Command descriptions.....	291
4.8.2.1	Configuration commands.....	291
4.8.2.2	Show commands.....	322
4.8.2.3	Clear commands.....	342
4.8.2.4	Monitor commands.....	344
5	Cflowd.....	348
5.1	Cflowd overview.....	348
5.1.1	Operation.....	348
5.1.1.1	Version 8.....	350
5.1.1.2	Version 9.....	351
5.1.1.3	Version 10.....	351
5.2	Cflowd configuration process overview.....	351
5.3	Configuration notes.....	352
5.4	Configuring cflowd with CLI.....	352
5.5	Cflowd configuration overview.....	352
5.5.1	Traffic sampling.....	352

5.5.2	Collectors.....	354
5.5.2.1	Aggregation.....	354
5.6	Basic cflowd configuration.....	354
5.7	Common configuration tasks.....	355
5.7.1	Global cflowd components.....	355
5.7.2	Configuring cflowd.....	356
5.7.3	Enabling cflowd.....	356
5.7.4	Configuring global cflowd parameters.....	357
5.7.5	Configuring cflowd collectors.....	357
5.7.5.1	Version 9 and Version 10 templates.....	358
5.7.6	Specifying cflowd options on an IP interface.....	367
5.7.6.1	Interface configurations.....	368
5.7.6.2	Service interfaces.....	368
5.7.7	Dependencies.....	368
5.8	Cflowd configuration management tasks.....	369
5.8.1	Modifying global cflowd components.....	369
5.8.2	Modifying cflowd collector parameters.....	370
5.9	Cflowd configuration command reference.....	370
5.9.1	Command hierarchies.....	370
5.9.1.1	Configuration commands.....	371
5.9.1.2	Show commands.....	371
5.9.1.3	Tools commands.....	371
5.9.1.4	Clear commands.....	372
5.9.2	Command descriptions.....	372
5.9.2.1	Global commands.....	372
5.9.2.2	Show commands.....	384
5.9.2.3	Tools commands.....	392
5.9.2.4	Clear commands.....	399
6	Common CLI command descriptions.....	400
6.1	Command descriptions.....	400
6.1.1	SAP syntax.....	400
	sap.....	400
7	Standards and protocol support.....	402
7.1	BGP.....	402

7.2	Ethernet.....	404
7.3	EVPN.....	405
7.4	Fast Reroute.....	405
7.5	Internet Protocol (IP) — General.....	406
7.6	IP — Multicast.....	408
7.7	IP — Version 4.....	410
7.8	IP — Version 6.....	410
7.9	IPsec.....	411
7.10	IS-IS.....	412
7.11	Management.....	413
7.12	MPLS — General.....	416
7.13	MPLS — GMPLS.....	417
7.14	MPLS — LDP.....	417
7.15	MPLS — MPLS-TP.....	417
7.16	MPLS — OAM.....	418
7.17	MPLS — RSVP-TE.....	418
7.18	OSPF.....	419
7.19	Pseudowire.....	420
7.20	Quality of Service.....	421
7.21	RIP.....	421
7.22	Timing.....	421
7.23	VPLS.....	423

List of tables

Table 1: Supported modes of operation and configuration methods.....	18
Table 2: Supported port modes by mode of operation.....	20
Table 3: 7210 SAS platforms supporting port modes.....	21
Table 4: Configuration process.....	22
Table 5: IPv6 header field descriptions.....	27
Table 6: BFD control packet field descriptions.....	33
Table 7: BFD IPv6 support matrix.....	35
Table 8: Default route preferences.....	67
Table 9: Output fields: router ARP.....	128
Table 10: Output fields: router ECMP.....	130
Table 11: Output fields: router BFD template.....	131
Table 12: Output fields: router BFD interface.....	132
Table 13: Output fields: router neighbor.....	134
Table 14: Output fields: router BFD session.....	136
Table 15: Output fields: router DHCP statistics.....	137
Table 16: Output fields: router DHCP summary.....	139
Table 17: Output fields: router ICMPv6.....	141
Table 18: Output fields: ICMPv6 interface.....	142
Table 19: Output fields: router interface.....	144
Table 20: Output fields: router interface detail.....	146
Table 21: Output fields: router route table.....	149

Table 22: Output fields: router advertisement.....	152
Table 23: Output fields: router static-ARP.....	154
Table 24: Output fields: static route.....	157
Table 25: Output fields: router status.....	158
Table 26: Output fields: router tunnel table.....	160
Table 27: Authentication data type.....	180
Table 28: LAG events.....	183
Table 29: host-unreachable operational states.....	229
Table 30: Route-unknown operational states.....	236
Table 31: Output fields: VRRP instance.....	240
Table 32: Output fields: VRRP policy.....	244
Table 33: Output fields: VRRP policy event.....	247
Table 34: Output fields: VRRP statistics.....	250
Table 35: Applying filter policies for 7210 SAS-T devices configured in network mode.....	257
Table 36: Applying filter policies for 7210 SAS-T (access-uplink mode).....	258
Table 37: Applying filter policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.....	258
Table 38: Applying filter policies for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE devices.....	259
Table 39: ACLs support in Epipe services on 7210 SAS-T network and access-uplink modes variants when using range SAPs.....	260
Table 40: ACLs support in Epipe services on 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp variants when using range SAPs.....	260
Table 41: ACLs support in VPLS services on 7210 SAS-T network and access-uplink mode variants when using range SAPs.....	260
Table 42: DSCP name to DSCP value table.....	264

Table 43: MAC match criteria exclusivity rules.....	269
Table 44: IP protocol IDs and descriptions.....	301
Table 45: 3-bit mask format.....	316
Table 46: 48-bit mask format.....	317
Table 47: 48-bit mask format.....	319
Table 48: Output fields: filter download-failed.....	323
Table 49: Output fields: filter IP.....	325
Table 50: Output fields: filter with filter ID specified.....	328
Table 51: Output fields: filter IP associations.....	330
Table 52: Output fields: filter counters.....	331
Table 53: Output fields: IP filter (no filter-id specified).....	333
Table 54: Output fields: filter (with filter-id specified).....	334
Table 55: Output fields: filter associations.....	336
Table 56: Output fields: filter counters.....	337
Table 57: Output fields: filter MAC.....	339
Table 58: Output fields: filter MAC associations.....	341
Table 59: Output fields: filter MAC counters.....	341
Table 60: Template-set.....	358
Table 61: Basic IPv4 template.....	359
Table 62: MPLS-IPv4 template.....	360
Table 63: Basic IPv6 template.....	361
Table 64: MPLS-IPv6 template.....	362
Table 65: Basic MPLS template.....	364

Table 66: MPLS-IP template.....	365
Table 67: Ethernet (L2-IP) flow template.....	366
Table 68: Cflowd configuration dependencies.....	369
Table 69: Output fields: cflowd collector.....	386
Table 70: Output fields: cflowd collector detailed.....	387
Table 71: Output fields: cflowd interface.....	389
Table 72: Output fields: cflowd status.....	391
Table 73: Output fields: tools dump cflowd cache.....	393
Table 74: Output fields: tools dump cflowd top-flows.....	396
Table 75: Output fields: tools dump cflowd top-protocols.....	398
Table 76: Formats of sap-id.....	400
Table 77: Port and encapsulation types.....	401

List of figures

Figure 1: IPv6 header format.....	27
Figure 2: IPv6 Internet exchange.....	28
Figure 3: IPv6 transit services.....	29
Figure 4: IPv6 services to enterprise customers and home users.....	29
Figure 5: IPv6 over IPv4 tunnels.....	30
Figure 6: Example of a 6PE topology within one AS.....	30
Figure 7: VRRP configuration.....	173
Figure 8: VRRP configuration and implementation flow.....	186
Figure 9: Creating and applying filter policies.....	262
Figure 10: Filtering process example.....	267
Figure 11: Applying an IP filter to an ingress interface.....	275
Figure 12: Basic cflowd steps.....	349
Figure 13: V5, V8, V9, V10, and flow processing.....	350
Figure 14: Cflowd configuration and implementation flow.....	351

1 Getting started

This chapter provides an overview of the document organization and content, and describes the terminology used in this guide.

1.1 About this guide



Note:

Unless explicitly noted otherwise, this guide uses 7210 SAS-Dxp to refer to the 7210 SAS-Dxp 12p, 7210 SAS-Dxp 16p, and 7210 SAS-Dxp 24p platforms.

This guide describes the logical IP routing interfaces and filtering support provided by the following 7210 SAS platforms, operating in one of the modes described in [Table 1: Supported modes of operation and configuration methods](#). If multiple modes of operation apply, they are explicitly noted in the topic.

- 7210 SAS-Mxp
- 7210 SAS-R6
- 7210 SAS-R12
- 7210 SAS-Sx/S 1/10GE
- 7210 SAS-Sx 10/100GE
- 7210 SAS-T

See [7210 SAS modes of operation](#) for information about the modes of operation supported by the 7210 SAS product family.



Note:

Unless explicitly noted otherwise, the phrase "Supported on all 7210 SAS platforms as described in this document" is used to indicate that the topic and CLI commands apply to all the 7210 SAS platforms in the following list, when operating in the specified modes only.

- network mode of operation
7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T
- standalone mode of operation
7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE
- standalone-VC mode of operation
7210 SAS-Sx/S 1/10GE

If the topic and CLI commands are supported on the 7210 SAS-T operating in the access-uplink mode, it is explicitly indicated, where applicable.

1.1.1 Document structure and content

This guide uses the following structure to describe logical IP routing interfaces and filtering content.



Note:

This guide generically covers Release 24.x.Rx content and may include some content that will be released in later maintenance loads. See the *7210 SAS Software Release Notes 24.x.Rx*, part number 3HE 20148 000x TQZZA, for information about features supported in each load of the Release 24.x.Rx software.

- This guide is organized into functional chapters and provides concepts and descriptions of the implementation flow. Each chapter describes a software area and provides CLI syntax and command usage to configure parameters for the functional area.
- Command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.
- See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for boot options to configure the [satellite](#) mode of operation on the router. See the 7750 SR software user guides for information about service and protocol provisioning, and operating the 7210 SAS router in [satellite](#) mode.
- Unless explicitly noted, the CLI commands and their configuration is similar for both [network](#) and [access-uplink](#) operating modes for features applicable to both modes of operation.

1.2 7210 SAS modes of operation

Unless explicitly noted, the phrase “mode of operation” and “operating mode” refers to the current operating mode of the 7210 SAS router. Each operating mode provides configuration access to a specific group of CLI commands.



Note:

Not all CLI commands are supported on all 7210 SAS platforms in all modes of operation. Users can only configure CLI commands supported by the current operating mode of the router. See the *7210 SAS Software Release Notes 24.x.Rx*, part number 3HE 20148 000x TQZZA, and to the appropriate 7210 SAS software user guide for information about features and capabilities supported by a 7210 SAS platform when operating in a specific mode.

The following modes of operation are supported by the 7210 SAS product family.

- **access-uplink**

In the access-uplink operating mode, the 7210 SAS router uplinks to the network using Layer 2 Ethernet VLAN switching (without IP/MPLS).

Platforms Supported: 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C, and 7210 SAS-T

- **network**

In the network operating mode, the 7210 SAS router uses IP/MPLS uplinks to the network. The IP routing protocols and MPLS functionality is available; see the appropriate 7210 SAS software user guide for more information about supported features.

Platforms Supported: 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C, 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T

- **satellite**

In the satellite operating mode, the 7210 SAS platform uses high-capacity uplinks (for example, 10GE ports on the 7210 SAS-Mxp and 100GE ports on the 7210 SAS-Sx 10/100GE) to connect to the 7750 SR host. The 7210 SAS router is managed by the 7750 SR host. There is no direct CLI access to the satellite node, and all services and protocols are configured on the host.

Platforms Supported: 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE

- **standalone**

In the standalone operating mode, the 7210 SAS platform supports IP/MPLS uplinks. It is operated and managed independently.

The functionality and features available on the standalone 7210 SAS platform are similar to the [network](#) operating mode. The standalone mode is primarily used to differentiate between a node being managed by the 7750 SR host (in the [satellite](#) operating mode), and a node managed independently (standalone operating mode).

Platforms Supported: 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE

- **standalone-VC**

In the standalone-VC operating mode, a set of 7210 SAS devices are stacked to provide larger 1GE/10GE port density and control-plane redundancy. The stack of nodes is provisioned and managed as a single chassis, and not as individual nodes.

The functionality and features available on the 7210 SAS platform are similar to the [network](#) operating mode, with additional capabilities, such as control-plane redundancy with non-stop routing and non-stop services.

Platforms Supported: 7210 SAS-Sx/S 1/10GE

For 7210 SAS platforms that support multiple explicit modes of operation ([Table 1: Supported modes of operation and configuration methods](#)), the operating mode must be configured in the Boot Option File (BOF) to ensure the router boots up in the specified mode. For example, the 7210 SAS-T supports access-uplink and network modes of operation, and the 7210 SAS-Sx/S 1/10GE supports satellite, standalone, and standalone-VC mode of operations. In some cases, the 7210 SAS router operates in a specific mode implicitly, and explicit configuration is not required.

See the appropriate *Basic System Configuration Guide* for boot options and information about how to boot the 7210 SAS platform in a specific operating mode.

The following table lists the supported modes of operation and the configuration methods for the 7210 SAS platforms. Unless explicitly noted otherwise, the operating mode is supported on all variants of the specific 7210 SAS platform.

Table 1: Supported modes of operation and configuration methods

7210 SAS platform	Mode of operation and configuration method				
	Network	Access-uplink	Standalone	Standalone-VC	Satellite
7210 SAS-D		Implicit	Implicit		
7210 SAS-Dxp		Implicit	Implicit		

7210 SAS platform	Mode of operation and configuration method				
	Network	Access-uplink	Standalone	Standalone-VC	Satellite
7210 SAS-K 2F1C2T		Implicit	Implicit		
7210 SAS-K 2F6C4T ¹	Port Mode Configuration ²	Port Mode Configuration ²	Implicit		
7210 SAS-K 3SFP+ 8C ¹	Port Mode Configuration ²	Port Mode Configuration ²	Implicit		
7210 SAS-Mxp	Implicit ³		Explicit BOF Configuration		Explicit BOF Configuration
7210 SAS-R6 ⁴	Implicit		Implicit		
7210 SAS-R12 ⁴	Implicit		Implicit		
7210 SAS-Sx/S 1/10GE	Implicit ³		Explicit BOF Configuration	Explicit BOF Configuration	Explicit BOF Configuration
7210 SAS-Sx 10/100GE	Implicit ³		Explicit BOF Configuration		Explicit BOF Configuration
7210 SAS-T	Explicit BOF Configuration	Explicit BOF Configuration	Implicit		

1.3 7210 SAS port modes

Unless explicitly noted, the phrase "port mode" refers to the current port configuration of the 7210 SAS node. The 7210 SAS platform supports the configuration of the following port modes.

- **access port mode**

Access ports are configured for customer-facing traffic if Service Access Points (SAPs) are required. The appropriate encapsulation type must be configured to distinguish the services on the port; services are configured on the port based on the encapsulation value.

Access ports can be configured on all the 7210 SAS platforms.

- **access-uplink port mode**

¹ By default, the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C boot up in the [network](#) mode of operation. These platforms also allow the use of [access-uplink port mode](#) (without explicit BOF configuration), which provides the option to use Layer 2 uplinks instead of IP/MPLS uplinks to the network core, similar to the 7210 SAS-K 2F1C2T router.

² See section [7210 SAS port modes](#) for information about port mode configuration

³ Implicitly operates in [network](#) mode when [standalone](#) mode of operation is configured

⁴ Supports MPLS uplinks only and implicitly operates in [network](#) mode

Access-uplink ports provide native Ethernet connectivity in service provider transport or in an infrastructure network. With this option, the encap-type can be configured to only QinQ. Access-uplink SAPs, which are QinQ SAPs, can only be configured on an access-uplink port to allow the operator to differentiate multiple services being carried over a single uplink port.

This is the default port mode of a 7210 SAS node in the [access-uplink](#) mode of operation.

- **network port mode**

Network ports are configured for network-facing traffic in the service provider transport or infrastructure network, and provide IP/MPLS uplinks.

This is the default port mode of a 7210 SAS node in the [network](#) or [standalone](#) mode of operation.

- **hybrid port mode**

Hybrid ports are configured for access and network facing traffic, and allow a single port to operate in both access and network modes.

Port modes available for configuration on a 7210 SAS node are determined by the current mode of operation of the router.



Note:

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C are unique; all port modes listed in [Table 2: Supported port modes by mode of operation](#) are available for configuration on the router, regardless of the current mode of operation.

The following table lists the port mode configuration support per 7210 SAS mode of operation.

Table 2: Supported port modes by mode of operation

Mode of operation	Supported port mode			
	Access	Network	Hybrid	Access-uplink
Access-Uplink	✓			✓
Network	✓	✓	✓	
Satellite ⁵				
Standalone	✓	✓	✓	
Standalone-VC	✓	✓	✓	

The following table lists the port mode configuration supported by the 7210 SAS product family. See the appropriate *Interface Configuration Guide* for more information about configuring the port modes for a specific platform.

⁵ Port modes are configured on the 7750 SR host and managed by the host.

Table 3: 7210 SAS platforms supporting port modes

Platform	Port mode			
	Access	Network	Hybrid	Access-uplink
7210 SAS-D	Yes	No	No	Yes
7210 SAS-Dxp	Yes	No	No	Yes
7210 SAS-K 2F1C2T	Yes	No	No	Yes
7210 SAS-K 2F6C4T	Yes	Yes	Yes	Yes
7210 SAS-K 3SFP+ 8C	Yes	Yes	Yes	Yes
7210 SAS-Mxp	Yes	Yes	Yes	No
7210 SAS-R6 IMM-b (IMMv2)	Yes	Yes	Yes	No
7210 SAS-R6 IMM-c 100GE (IMM-c 1QSFP28)	Yes	Yes	Yes	No
7210 SAS-R12 IMM-b	Yes	Yes	Yes	No
7210 SAS-R12 IMM-c 100GE (IMM-c 1QSFP28)	Yes	Yes	Yes	No
7210 SAS-Sx/S 1/10GE	Yes	Yes	Yes	No
7210 SAS-Sx 10/100GE	Yes	Yes	Yes	No
7210 SAS-T	Yes	Yes ⁶	Yes ⁷	Yes ⁸

1.4 7210 SAS router configuration process

The following table lists the tasks necessary to configure logical IP routing interfaces, virtual routers, IP and MAC-based filtering.

This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

⁶ Network ports are supported only if the node is operating in network mode.

⁷ Hybrid ports are supported only if the node is operating in network mode.

⁸ Access-uplink ports are supported only if the node is operating in access-uplink mode.

Table 4: Configuration process

Area	Task	Chapter
Router configuration	Configure router parameters, including router interfaces, addresses, router IDs, and autonomous systems	IP router configuration
	Configure VRRP parameters	VRRP
	Configure IP and MAC filters	Filter policies
Reference	List of IEEE, IETF, and other proprietary entities	Standards and protocol support

1.5 Conventions

This section describes the general conventions used in this guide.

1.5.1 Precautionary and information messages

The following information symbols are used in the documentation.



DANGER: Danger warns that the described activity or situation may result in serious personal injury or death. An electric shock hazard could exist. Before you begin work on this equipment, be aware of hazards involving electrical circuitry, be familiar with networking environments, and implement accident prevention procedures.



WARNING: Warning indicates that the described activity or situation may, or will, cause equipment damage, serious performance problems, or loss of data.



Caution: Caution indicates that the described activity or situation may reduce your component or system performance.



Note: Note provides additional operational information.



Tip: Tip provides suggestions for use or best practices.

1.5.2 Options or substeps in procedures and sequential workflows

Options in a procedure or a sequential workflow are indicated by a bulleted list. In the following example, at step 1, the user must perform the described action. At step 2, the user must perform one of the listed options to complete the step.

Example: Options in a procedure

1. User must perform this step.

2. This step offers three options. User must perform one option to complete this step:

- This is one option.
- This is another option.
- This is yet another option.

Substeps in a procedure or a sequential workflow are indicated by letters. In the following example, at step 1, the user must perform the described action. At step 2, the user must perform two substeps (a. and b.) to complete the step.

Example: Substeps in a procedure

1. User must perform this step.
2. User must perform all substeps to complete this action:
 - a. This is one substep.
 - b. This is another substep.

2 IP router configuration

This chapter provides information about commands required to configure basic router parameters.

2.1 Configuring IP router parameters

To provision services on a 7210 SAS router, logical IP routing interfaces must be configured to associate attributes, such as an IP address or the system with the IP interface.

A special type of IP interface is the system interface. A system interface must have an IP address with a 32-bit subnet mask. The system interface is used as the router identifier by higher-level protocols such as OSPF and BGP, unless overwritten by an explicit router ID.

2.1.1 Interfaces

7210 SAS routers use different types of interfaces for various functions. Interfaces must be configured with parameters, such as the interface type (system) and address. A port is not associated with a system interface. An interface can be associated with the system (loopback address).

2.1.1.1 Secondary IPv4 addresses

Secondary IPv4 addresses can be assigned to an IP interface to allow multiple IP subnets to be assigned to a single Ethernet LAN segment. This is useful in IPv4 address migration and the use of a single VLAN for multiple IP subnets in some network designs.

On the 7210 SAS, IPv4 secondary addresses are supported for the following IP interface types:

- IES IP interface
- VPRN IP interface
- Routed VPLS IP interface
- Network IP interface

The following optional functionality is supported with IPv4 secondary addresses:

- Use IPv4 secondary addresses with static routes.
- Advertise IPv4 secondary addresses through routing protocols such as OSPF or IS-IS.
- Use IPv4 secondary addresses with VRRP (IPv4).

The following restrictions apply to the use of IPv4 secondary addresses:

- An IP interface must be assigned a primary IP address before a secondary IP address can be used.
- Secondary IP addresses cannot be used for setting up OSPF or IS-IS neighbors.
- Secondary IP addresses cannot be used with MPLS protocols and PWs (such as RSVP, LDP, or BGP 3107) for specifying parameters such as path information. That is, MPLS protocols and PWs always use primary IPv4 addresses.

2.1.1.2 Network interface

A network interface (a logical IP routing interface) can be configured on a physical port.

2.1.2 System interface

The system interface is associated with the network entity (such as a specific router or switch), not a specific interface. The system interface is also referred to as the loop-back address. The system interface is associated during the configuration of the following entities as follows:

- the termination point of service tunnels
- the hops when configuring MPLS paths and LSPs
- the addresses on a target router for BGP and LDP peering

The system interface is used to preserve connectivity (when routing re-convergence is possible) when an interface fails or is removed. The system interface is also referred to as the loop-back address and is used as the router identifier. A system interface must have an IP address with a 32-bit subnet mask.

2.1.3 Router ID

The router ID, a 32-bit number, uniquely identifies the router within an autonomous system (AS). In protocols such as OSPF, routing information is exchanged between areas, groups of networks that share routing information. It can be set to be the same as the loop-back address. The router ID is used by both OSPF and BGP routing protocols in the routing table manager instance.

There are several ways to obtain the router ID. On each 7210 SAS router, the router ID can be derived in the following ways:

- Define the value in the **config>router** *router-id* context. The value becomes the router ID.
- Configure the system interface with an IP address in the **config>router>interface** *ip-int-name* context. If the router ID is not manually configured in the **config>router** *router-id* context, then the system interface acts as the router ID.
- If neither the system interface or router ID are implicitly specified, then the router ID is inherited from the last four bytes of the MAC address.
- The router can be derived on the protocol level.

2.1.4 Autonomous systems



Note:

BGP protocol (only selected families) is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Networks can be grouped into areas. An area is a collection of network segments within an autonomous system (AS) that have been administratively assigned to the same group. An area topology is concealed from the rest of the AS, which results in a significant reduction in routing traffic.

Routing in the AS takes place on two levels, depending on whether the source and destination of a packet reside in the same area (intra-area routing) or different areas (inter-area routing). In intra-area routing,

the packet is routed solely on information obtained within the area; no routing information obtained from outside the area can be used. This protects intra-area routing from the injection of bad routing information.

Routers that belong to more than one area are called area border routers. All routers in an AS do not have an identical topological database. An area border router has a separate topological database for each area it is connected to. Two routers, which are not area border routers, belonging to the same area, have identical area topological databases.

Autonomous systems share routing information, such as routes to each destination and information about the route or AS path, with other ASs using BGP. Routing tables contain lists of next hops, reachable addresses, and associated path cost metrics to each router. BGP uses the information and path attributes to compile a network topology.

2.1.5 Proxy ARP

Proxy ARP is the technique in which a router answers ARP requests intended for another node. The router appears to be present on the same network as the "real" node that is the target of the ARP and takes responsibility for routing packets to the "real" destination. Proxy ARP can help nodes on a subnet reach remote subnets without configuring routing or a default gateway. Typical routers only support proxy ARP for directly attached networks; the router is targeted to support proxy ARP for all known networks in the routing instance where the virtual interface proxy ARP is configured.

To support DSLAM and other edge like environments, proxy ARP supports policies that allow the provider to configure prefix lists that determine for which target networks proxy ARP will be attempted and prefix lists that determine for which source hosts proxy ARP will be attempted.

In addition, the proxy ARP implementation will support the ability to respond for other hosts within the local subnet domain. This is needed in environments such as DSL where multiple hosts are in the same subnet but cannot reach each other directly.

Static ARP is used when a 7210 SAS router needs to know about a device on an interface that cannot or does not respond to ARP requests. Therefore, the configuration can state that if it has a packet with 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.

2.1.6 Internet Protocol versions

The TiMOS implements IP routing functionality, providing support for IP version 4 (IPv4) and IP version 6 (IPv6). IP version 6 (RFC 1883, Internet Protocol, Version 6 (IPv6)) is a newer version of the Internet Protocol designed as a successor to IP version 4 (IPv4) (RFC-791, Internet Protocol). The changes from IPv4 to IPv6 effects the following categories:

- **Expanded addressing capabilities**

IPv6 increases the IP address size from 32 bits (IPv4) to 128 bits, to support more levels of addressing hierarchy, a much greater number of addressable nodes, and simpler auto-configuration of addresses. The scalability of multicast routing is improved by adding a scope field to multicast addresses. Also, a new type of address called an any cast address is defined that is used to send a packet to any one of a group of nodes.

- **Header format simplification**

Some IPv4 header fields have been dropped or made optional to reduce the common-case processing cost of packet handling and to limit the bandwidth cost of the IPv6 header.

- **Improved support for extensions and options**

Changes in the way IP header options are encoded allows for more efficient forwarding, less stringent limits on the length of options, and greater flexibility for introducing new options in the future.

- **Flow labeling capability**

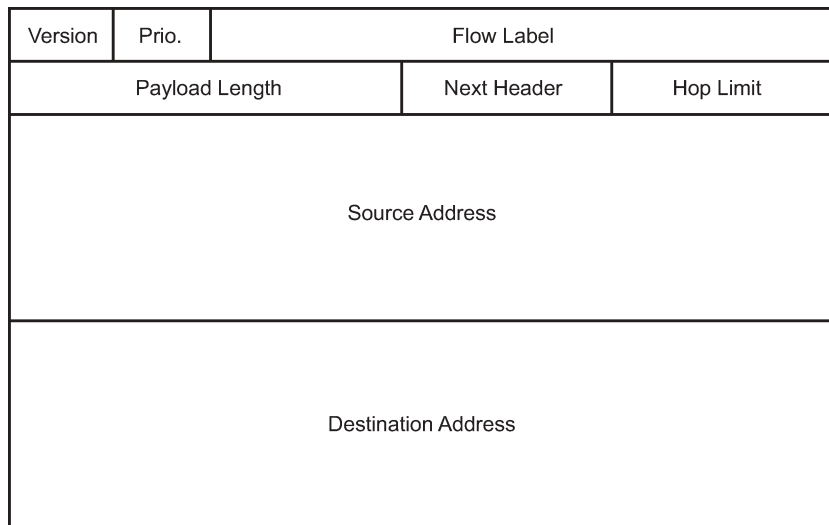
The capability to enable the labeling of packets belonging to particular traffic flows for which the sender requests special handling, such as non-default quality of service or “real-time” service was added in IPv6.

- **Authentication and privacy capabilities**

Extensions to support authentication, data integrity, and (optional) data confidentiality are specified for IPv6.

The following figure shows the IPv6 header format.

Figure 1: IPv6 header format



sw0706

The following table describes IPv6 header fields.

Table 5: IPv6 header field descriptions

Field	Description
Version	4-bit Internet Protocol version number = 6.
Prio.	4-bit priority value.
Flow Label	24-bit flow label.
Payload Length	6-bit unsigned integer. The length of payload, for example, the rest of the packet following the IPv6 header, in octets. If the value is zero, the payload length is carried in a jumbo payload hop-by-hop option.
Next Header	8-bit selector. Identifies the type of header immediately following the IPv6 header.

Field	Description
	This field uses the same values as the IPv4 protocol field.
Hop Limit	8-bit unsigned integer. Decremented by 1 by each node that forwards the packet. The packet is discarded if the hop limit is decremented to zero.
Source Address	128-bit address of the originator of the packet.
Destination Address	128-bit address of the intended recipient of the packet (possibly not the ultimate recipient if a routing header is present).

2.1.6.1 IPv6 applications

The IPv6 applications for 7210 SAS are:


- IPv6 inband management of the node using network port IPv6 IP interface
- IPv6 transit traffic (using network port IPv6 IP interfaces)

Examples of the IPv6 applications supported by the -TiMOS include:

- **IPv6 Internet exchange peering**

The following figure shows an IPv6 Internet exchange where multiple ISPs peer over native IPv6.


Figure 2: IPv6 Internet exchange



- **IPv6 transit services**

The following figure shows IPv6 transit provided by an ISP.

Figure 3: IPv6 transit services

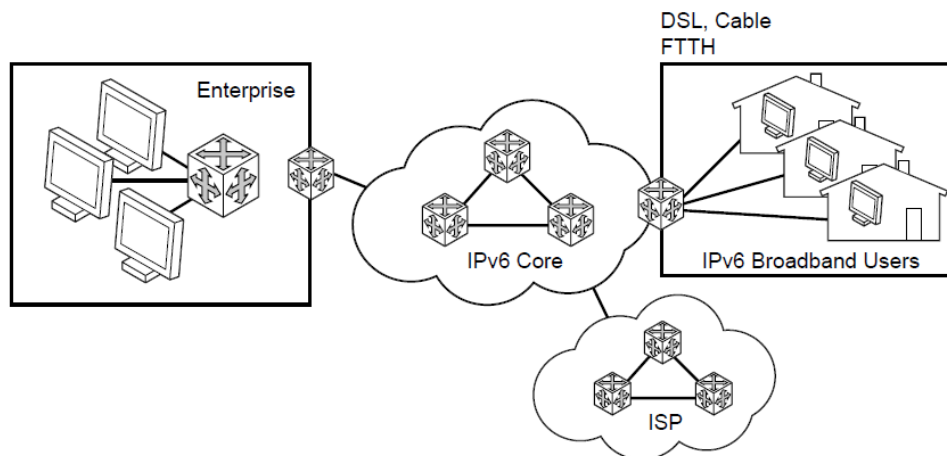


IPIPE_008

- **IPv6 services to enterprise customers and home users**

The following figure shows IPv6 connectivity to enterprise and home broadband users.

Figure 4: IPv6 services to enterprise customers and home users



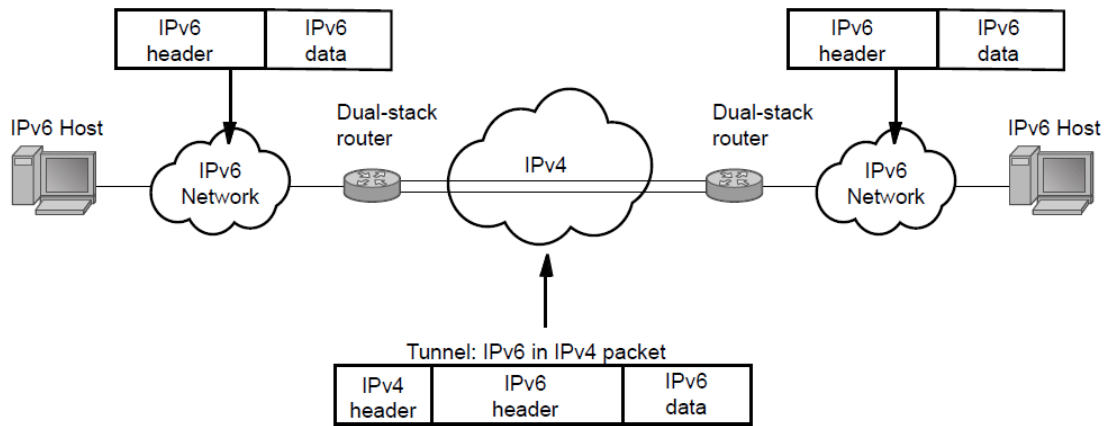
IPIPE_009

- **IPv6 over IPv4 relay services**

IPv6 over IPv4 tunnels are one of many IPv6 transition methods to support IPv6 in an environment where not only IPv4 exists but native IPv6 networks depend on IPv4 for greater IPv6 connectivity. The 7210 SAS supports dynamic IPv6 over IPv4 tunneling. The ipv4 source and destination address are taken from configuration, the source address is the ipv4 system address and the ipv4 destination is the next hop from the configured 6over4 tunnel.

IPv6 over IPv4 is an automatic tunnel method that gives a prefix to the attached IPv6 network. The following figure shows IPv6 over IPv4 tunneling to transition from IPv4 to IPv6.

Figure 5: IPv6 over IPv4 tunnels



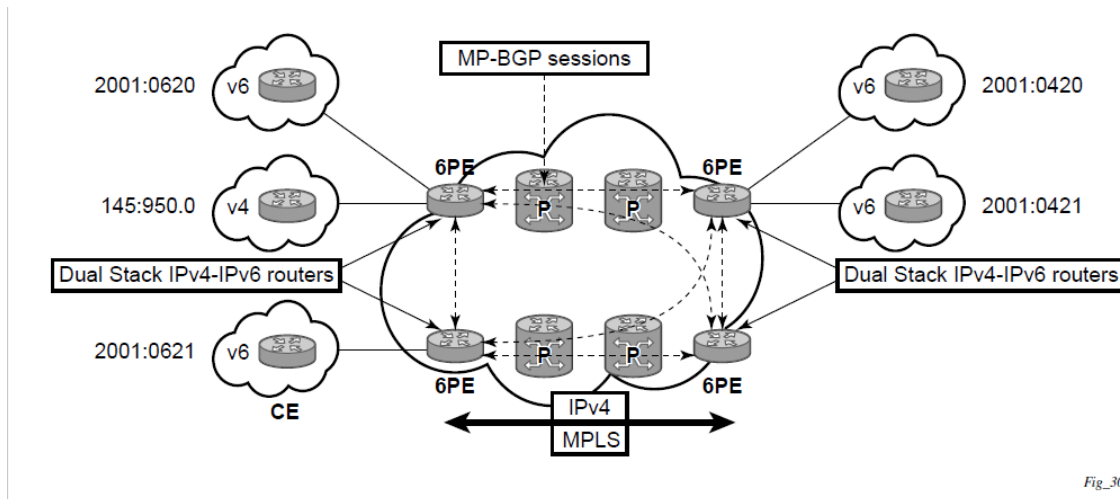
Fig_29a

2.1.6.2 IPv6 Provider Edge router over MPLS (6PE)

6PE allows IPv6 domains to communicate with each other over an IPv4 MPLS core network. This architecture requires no backbone infrastructure upgrades and no re-configuration of core routers, because forwarding is purely based on MPLS labels. 6PE is a cost effective solution for IPv6 deployment.

The following figure shows an example of a 6PE topology within one AS.

Figure 6: Example of a 6PE topology within one AS



Fig_30

2.1.6.2.1 6PE control plane support

The 6PE MP-BGP routers support:

- IPv4/IPv6 dual-stack
- MP-BGP can be used between 6PE routers to exchange IPv6 reachability information as follows:
 - The 6PE routers exchange IPv6 prefixes over MP-BGP sessions running over IPv4 transport. The MP-BGP AFI used is IPv6 (value 2).
 - An IPv4 address of the 6PE router is encoded as an IPv4-mapped IPv6 address in the BGP next-hop field of the IPv6 NLRI. By default, the IPv4 address that is used for peering is used. It is configurable through the route policies.
 - The 6PE router binds MPLS labels to the IPv6 prefixes it advertises. The SAFI used in MP-BGP is the SAFI (value 4) label. The router uses the IPv6 explicit null (value 2) label for all the IPv6 prefixes that it advertises and can accept an arbitrary label from its peers.
- LDP is used to create the MPLS full mesh between the 6PE routers and the IPv4 addresses that are embedded in the next-hop field are reachable by LDP LSPs. The ingress 6PE router uses the LDP LSPs to reach remote 6PE routers.

2.1.6.2.2 6PE data plane support

The ingress 6PE router can push two MPLS labels to send the packets to the egress 6PE router. The top label is an LDP label used to reach the egress 6PE router. The bottom label is advertised in MP-BGP by the remote 6PE router. Typically, the IPv6 explicit null (value 2) label is used but an arbitrary value can be used when the remote 6PE router is from a vendor other than Nokia.

The egress 6PE router pops the top LDP tunnel label. It sees the IPv6 explicit null label, which indicates an IPv6 packet is encapsulated. It also pops the IPv6 explicit null label and performs an IPv6 route lookup to find out the next hop for the IPv6 packet.

2.1.7 Bidirectional Forwarding Detection



Note:

This feature is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Bidirectional Forwarding Detection (BFD) is a light-weight, low-overhead, short-duration mechanism to detect failures in the path between two systems. If a system stops receiving BFD messages for a long enough period (based on configuration) it is assumed that a failure along the path has occurred and the associated protocol or service is notified of the failure.

The following are the advantages of implementing the BFD mechanism:

- used for activity detection over any media type
- can be used at any protocol layer
- proliferation of different methods can be avoided
- can be used with a wide range of detection times and overhead

BFD is implemented in asynchronous mode, in this mode periodic BFD control messages are used to test the path between the systems.

A path is declared operational when two-way communication has been established between both the systems. A separate BFD session is created for each communication path and data protocol between two systems.

BFD also supports the Echo function defined in *draft-ietf-bfd-base-04.txt*, Bidirectional Forwarding Detection. In this scenario one of the systems send a sequence of BFD echo packets to the other system which loops back the echo packets within the systems forwarding plane. If many of the echo packets are lost, the BFD session is declared as down.

2.1.7.1 BFD control packet

The base BFD specification does not specify the encapsulation type to be used for sending BFD control packets. Choice of the appropriate encapsulation-type to be implemented is based on the network and medium. The encapsulation for BFD over IPv4 and IPv6 networks is specified in *draft-ietf-bfd-v4v6-1hop-04.txt*, *BFD for IPv4 and IPv6 (Single Hop)*. This specification requires that BFD control packets be sent over UDP with a destination port number of 3784 and the source port number must be within the range 49152 to 65535.

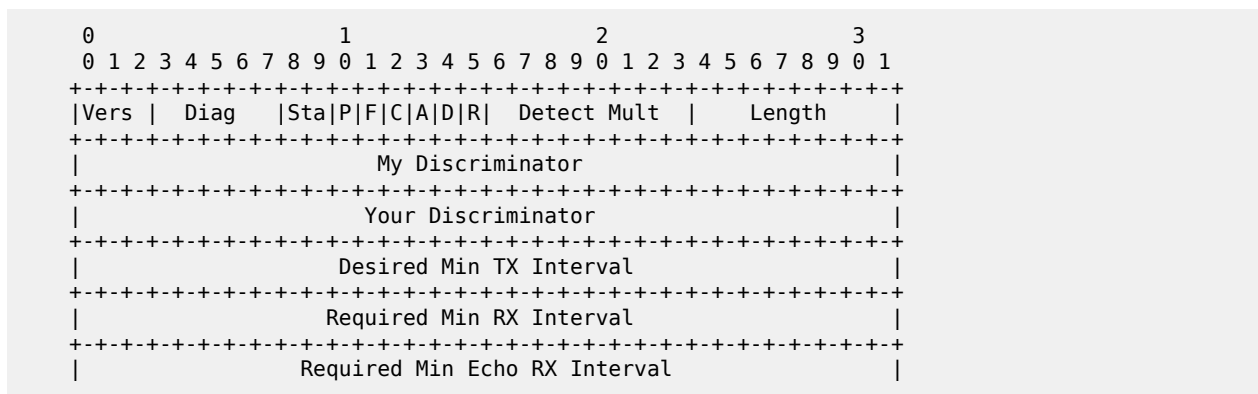


Note:

- The TTL of all transmitted BFD packets must have an IP TTL of 255
- If authentication is not enabled, all BFD packets received must have an IP TTL of 255.
- If authentication is enabled, the IP TTL should be 255. In case the IP TTL is not 255 the BFD packets are still processed, if packet passes the enabled authentication mechanism.
- If multiple BFD sessions exist between two nodes, the BFD discriminator is used to demultiplex the BFD control packet to the appropriate BFD session.

2.1.7.2 Control packet format

The BFD control packet has 2 sections, a mandatory section and an optional authentication section. The mandatory section is as follows.



The following table describes BFD control packet fields.

Table 6: BFD control packet field descriptions

Field	Description
Vers	The version number of the protocol. The initial protocol version is 0.
Diag	A diagnostic code specifying the local system reason for the last transition of the session from Up to some other state. Possible values are: 0-No diagnostic 1-Control detection time expired 2-Echo function failed 3-Neighbor signaled session down 4-Forwarding plane reset 5-Path down 6-Concatenated path down 7-Administratively down
H Bit	The "I Hear You" bit. This bit is set to 0 if the transmitting system either is not receiving BFD packets from the remote system, or is in the process of tearing down the BFD session for some reason. Otherwise, during normal operation, it is set to 1.
D Bit	The "demand mode" bit. (Not supported)
P Bit	The poll bit. If set, the transmitting system is requesting verification of connectivity, or of a parameter change.
F Bit	The final bit. If set, the transmitting system is responding to a received BFD control packet that had the poll (P) bit set.
Rsvd	Reserved bits. These bits must be zero on transmit and ignored on receipt.
Detect Mult	The "Detect time multiplier". In the Asynchronous mode, Detection time = Detect time Multiplier * transmit interval. If a BFD control packet is not received from the remote system within the detection time, implies that a failure has occurred.
Length	Length of the BFD control packet, in bytes.
My Discriminator	A unique, nonzero discriminator value generated by the transmitting system, used to demultiplex multiple BFD sessions between the same pair of systems.
Your Discriminator	The discriminator received from the corresponding remote system. This field reflects back the received value of my discriminator, or is zero if that value is unknown.
Desired Min TX Interval	This is the minimum interval, in microseconds, that the local system would like to use when transmitting BFD control packets.

Field	Description
Required Min RX Interval	This is the minimum interval, in microseconds, between received BFD control packets that this system is capable of supporting.
Required Min Echo RX Interval	This is the minimum interval, in microseconds, between received BFD echo packets that this system is capable of supporting. If this value is zero, the transmitting system does not support the receipt of BFD echo packets.

2.1.7.3 Echo support

In the BFD echo support scenario, the 7210 SAS loops back received BFD echo messages to the original sender based on the destination IP address in the packet.

The echo function is useful when the local router does not have sufficient CPU power to handle a periodic polling rate at a high frequency. As a result, it relies on the echo sender to send a high rate of BFD echo messages through the receiver node, which is only processed by the receiver forwarding path. This allows the echo sender to send BFD echo packets at any rate.

The 7210 SAS supports only response to echo requests and does not support sending of echo requests.

2.1.7.4 BFD IPv4 support on 7210 SAS platforms

BFD IPv4 support on 7210 SAS platforms is as follows.

BFD IPv4 in a VPRN service is supported for:

- OSPF2 PE-CE routing protocol
- static routes
- VRRP
- BGP for PE-CE protocol
- PIM for PE-CE protocol for ng-MVPN

BFD IPv4 in an IES service is supported for:

- OSPFv2
- IS-IS for IPv4 interfaces
- static routes
- VRRP

BFD IPv4 in the base routing instance is supported for:

- OSPFv2 on network IPv4 interfaces
- IS-IS on network IPv4 interfaces
- VRRP on network IPv4 interfaces
- MP-BGP for vpn-ipv4 and vpn-ipv6 family (only multi-hop)
- static routes
- RSVP-TE

- PIM
- TLDP
- interface LDP (link-level)

BFD IPv4 for MPLS-TP is supported for:

- BFD for MPLS-TP LSP linear protection, only on 7210 SAS-T, 7210 SAS-R6, and 7210 SAS-R12



Note:

- See [BFD IPv6 support on 7210 SAS platforms](#) for more information about BFD IPv6 support on 7210 SAS platforms.
- On the 7210 SAS-T, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp, BFD processing is supported in hardware, enabling faster detection (minimum timer supported is 10 ms). Hardware-based BFD sessions are supported only for an IP interface configured on a port. For IP interfaces configured over LAG and for BFD sessions using system IP addresses or loop-back IP addresses, CPM CPU-based sessions are supported with a minimum timer of 100ms.

2.1.7.5 BFD IPv6 support on 7210 SAS platforms



Note:

- BFD for IPv6 interfaces is supported only the 7210 SAS-Mxp.
- On the 7210 SAS-Mxp, BFD IPv6 processing is supported in hardware, enabling faster detection (the minimum timer supported is 10 ms). Hardware-based BFD sessions are supported only for an IP interface configured on a port.
- For IP interfaces using the IPv6 addresses of interfaces configured over LAG and for BFD sessions using system IPv6 addresses or loopback IPv6 addresses, CPM CPU-based sessions are supported with a minimum timer of 100 ms.

BFD IPv6 is supported on the 7210 SAS-Mxp in VPRN, IES, and R-VPLS services and network IPv6 interfaces. The following table indicates the support matrix for BFD IPv6 on the 7210 SAS-Mxp.

Table 7: BFD IPv6 support matrix

Service	Routing protocol	Hardware or central CPU-based BFD IPv6 session		IP address used by BFD IPv6 session	
		Hardware	Central	Link local	Global IPv6 address
Network Interface	OSPF	✓	✓	✓	
	IS-IS	✓	✓	✓	
	BGP	✓	✓		✓
	Static Route	✓	✓		✓

Service	Routing protocol	Hardware or central CPU-based BFD IPv6 session		IP address used by BFD IPv6 session	
		Hardware	Central	Link local	Global IPv6 address
	VRRP	✓	✓		✓
IES	OSPF	✓	✓	✓	
	IS-IS	✓	✓	✓	
	BGP	✓	✓		✓
	Static Route	✓	✓		✓
	VRRP	✓	✓		✓
R-VPLS	OSPF		✓	✓	
	IS-IS		✓	✓	
	BGP		✓		✓
	Static Route		✓		✓
	VRRP		✓		✓
VPRN	OSPF				
	IS-IS				
	BGP	✓	✓		✓
	Static Route	✓	✓		✓
	VRRP	✓	✓		✓

2.1.8 IGP-LDP and static route-LDP synchronization

With LDP, FECs learned from an interface do not necessarily link to that interface state. As long as the router that advertised the labels is reachable, the learned labels are stored in the incoming label map (ILM) table.

Although this feature gives LDP a lot of flexibility, it can also cause problems. For example, when an interface comes back up from a failure or from a shutdown state, the static routes bound to that interface are installed immediately. However, the LDP adjacency to the next hop might not be up, which means that the LDP SDP remains down. In this case, the MPLS traffic will be blackholed until the LDP adjacency comes up.

The same issue also applies to dynamic routes (OSPF and IS-IS).

To resolve this issue, the LDP synchronization timer enables synchronization of IGP or static routes to the LDP state.

With IGP, when a link is restored after a failure, IGP sets the link cost to infinity and advertises it. The value advertised in OSPF is 0xFFFF (65535). The value advertised in IS-IS regular metric is 0x3F (63) and in IS-IS wide-metric is 0xFFFFFE (16777214).

After IGP advertises the link cost, the LDP hello adjacency is brought up with the neighbor. The LDP synchronization timer is started by IGP from the time the LDP session to the neighbor is up over the interface. This synchronization timer allows time for the label-FEC bindings to be exchanged.

When the LDP synchronization timer expires, the link cost is restored and is readvertised. IGP will announce a new best next-hop and LDP will use it if the label binding for the neighbor FEC is available.

The preceding behavior is similar for static routes. If the static route is enabled for **ldp-sync**, the route is not enabled immediately after the interface to the next hop comes up. Routes are suppressed until the LDP adjacency with the neighbor comes up and the synchronization timer expires. The timer does not start until the LDP adjacency with the neighbor node is fully established.



Note:

- IGP-LDP synchronization is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.
- Static route-LDP synchronization is supported on all 7210 SAS platforms as described in this document, except platforms operating in access-uplink mode.

2.1.9 IP fragmentation



Note:

This feature is supported only on the 7210 SAS-Mxp.

The 7210 SAS does not support native IP fragmentation for IP packets that exceed the configured MTU. However, in situations where IP fragmentation is necessary, the CPM can be configured to extract oversized IPv4 packets from the datapath, fragment them using the system CPU, and insert the fragmented packets back into the datapath.

Run the **configure system ip allow-cpu-fragmentation** command to enable IP fragmentation. See the 7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide for more information about this command.



Caution:

- CPU-fragmented packets are subject to additional delay when compared to non-fragmented datapath forwarded packets.
- IP fragmentation in the CPM CPU competes for CPU cycles as a low-priority task. The number of fragmentations at any given time is limited to ensure system performance.

Even when IP fragmentation is enabled, packets that exceed the configured MTU are dropped if the Do not Fragment (DF) bit is set in the IP header.

IP fragmentation on the 7210 SAS is supported in the following contexts:

- network IP interface
- IES including R-VPLS



Note:

On R-VPLS interfaces, the decision to fragment a packet is based on the egress port MTU and the fragment size is determined by the service MTU.

2.2 Process overview

The following items are components to configure basic router parameters:

- **interface**

A logical IP routing interface. When created, attributes like an IP address, port, link aggregation group or the system can be associated with the IP interface.

- **address**

The address associates the device system name with the IP system address. An IP address must be assigned to each IP interface.

- **system interface**

This creates an association between the logical IP interface and the system (loop-back) address. The system interface address is the circuit-less address (loop-back) and is used by default as the router ID for protocols such as OSPF and BGP.

- **router ID**

(Optional) The router ID specifies the router's IP address.

- **autonomous system**

(Optional) An autonomous system (AS) is a collection of networks that are subdivided into smaller, more manageable areas.

2.3 Configuration notes

The following information describes router configuration guidelines:

- A system interface and associated IP address should be specified.
- Boot options file (BOF) parameters must be configured before configuring router parameters.
- IPv4 and IPv6 route table lookup entries are shared. Before adding routes for IPv6 destinations, route entries in the routed lookup table needs to be allocated for IPv6 addresses. This can be done using the **config system resource-profile max-ipv6-routes** CLI command. This command allocates route entries for /64 IPv6 prefix route lookups. The system does not allocate any IPv6 route entries by default and user needs to allocate some resources before using IPv6. For the command to take effect the node must be rebooted after making the change. See the 7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide for more information.
- A separate route table (or a block in the route table) must be used for IPv6 /128-bit prefix route lookup. A limited number of IPv6 /128 prefixes route lookup entries are supported. The software enables lookups in this table by default, so no user configuration is required to enable IPv6 /128-bit route lookup.
- IPv6 interfaces are allowed to be created without allocating IPv6 route entries; only IPv6 hosts on the same subnet are reachable.

- In 7210 SAS, the FIB is shared among all routing instances (Base instance, management instance, and VPRN service instances).
- Software shuts down control protocols (for example, OSPF) if the routing FIB (either IPv4 FIB or IPv6 FIB) size limit is exceeded. Users must ensure through proper network design that the FIB size is not exceeded. Use the available tools (that is, route policies) to ensure that all the features that share the IPv4/IPv6 FIB do not install routes more than the available FIB size.

2.4 Configuring an IP router with CLI

This section provides information to configure an IP router.

2.4.1 Router configuration overview

On a 7210 SAS, an interface is a logical named entity. An interface is created by specifying an interface name under the **config>router** context. This is the global router configuration context where objects like static routes are defined. An IP interface name can be up to 32 alphanumeric characters, must start with a letter, and is case-sensitive; for example, the interface name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed.

To create an interface on a 7210 SAS, the basic configuration tasks are as follows:

- Assign a name to the interface.
- Associate an IP address with the interface.
- Associate the interface with a network interface or the system interface.
- Associate the interface with a system or a loop-back interface.
- Configure appropriate routing protocols.

A system interface and network interface should be configured.

2.4.1.1 System interface

The system interface is associated with the network entity (such as a specific 7210 SAS IP router), not a specific interface. The system interface is also referred to as the loop-back address. The system interface is associated during the configuration of the following entities:

- termination point of service tunnels
- hops when configuring MPLS paths and LSPs
- addresses on a target router for BGP and LDP peering

The system interface is used to preserve connectivity (when routing re-convergence is possible) when an interface fails or is removed. The system interface is used as the router identifier. A system interface must have an IP address with a 32-bit subnet mask.

2.4.1.2 Network interface



Note:

Network interfaces are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

A network interface can be configured on a physical port or LAG on a physical or logical port.

2.4.2 Basic configuration

The most basic router configuration must have the following:

- system name
- system address

Output example

The following is a sample router configuration output.

```
A:ALA-A> config# info
. . .
#-----
# Router Configuration
#-----
router
  interface "system"
    address 10.10.10.103/32
  exit
  interface "to-104"
    address 10.0.0.103/24
    port 1/1/1
  exit
  exit
  autonomous-system 12345

router-id 10.10.10.103
...
  exit
  isis
  exit
...
#-----
A:ALA-A> config#
```

2.4.3 Common configuration tasks

The following sections describe basic system tasks.

2.4.3.1 Configuring a system name

Use the **system** command to configure a name for the device. The name is used in the prompt string. Only one system name can be configured. If multiple system names are configured, the last one configured overwrites the previous entry.

If special characters are included in the system name string, such as spaces, #, or ?, the entire string must be enclosed in double quotes. Use the following syntax to configure the system name.

```
config# system
      name system-name
```

Example

```
config# system
config>system# name ALA-A
ALA-A>config>system# exit all
ALA-A#
```

Output example

The following is a sample system name configuration output.

```
A:ALA-A>config>system# info
#-----
# System Configuration
#-----
      name "ALA-A"
      location "Mt.View, CA, NE corner of FERG 1 Building"
      coordinates "37.390, -122.05500 degrees lat."
      snmp
      exit
      . . .
      exit
-----
```

2.4.3.2 Configuring interfaces

The following command sequences create a system and a logical IP interface. The system interface cannot be deleted.

2.4.3.2.1 Configuring a system interface

Use the following syntax to configure a system interface.

```
config>router
      interface interface-name
      address {[ip-address/mask]|[ip-address] [netmask]} [broadcast {all-ones | host-ones}]
```

Output example

The following is a sample IP configuration output showing network interface information.

```
A:ALA-A>config>router# info
#-----
# IP Configuration
#-----
      interface "system"
      address 10.10.0.4/32
      exit
      interface "to-ALA-2"
```

```

        address 10.10.24.4/24
        port 1/1/1
        egress
            filter ip 10
        exit
    exit
...
#-----
A:ALA-A>config>router#

```

2.4.3.2.2 Configuring IPv6 parameters

Output example

The following is a sample interface configuration output showing the IPv6 default configuration when IPv6 is enabled on the interface.

```

*A:dut-d>config>router>if>ipv6# info detail
-----
        icmp6
            packet-too-big 100 10
            param-problem 100 10
            redirects 100 10
            time-exceeded 100 10
            unreachablees 100 10
        exit
        address 2001:db8::1/64
        no dad-disable
        no reachable-time
        no neighbor-limit
        no qos-route-lookup
        no local-proxy-nd
        no tcp-mss
-----

```

Use the following syntax to configure IPv6 parameters on a router interface.

```

config>router# interface interface-name
port port-name
ipv6
    address {ipv6-address/prefix-length} [eui-64]
    icmp6
        packet-too-big [number seconds]
        param-problem [number seconds]
        redirects [number seconds]
        time-exceeded [number seconds]
        unreachablees [number seconds]
    neighbor ipv6-address mac-address

```

Output example

The following is a sample configuration output showing interface information.

```

A:ALA-49>config>router>if# info
-----
        address 10.11.10.1/64
        port 1/1/10
        ipv6
            address 2001:db8::1/64

```

```
exit
-----
A:ALA-49>config>router>if#
```

2.4.3.3 Configuring router advertisement

To configure the router to originate router advertisement messages on an interface, the interface must be configured under the router-advertisement context and be enabled (**no shutdown**). All other router advertisement configuration parameters are optional.

Use the following syntax to enable router advertisement and configure router advertisement parameters.

```
config>router# router-advertisement
  interface ip-int-name
  current-hop-limit number
  managed-configuration
  max-advertisement-interval seconds
  min-advertisement-interval seconds
  mtu mtu-bytes
  other-stateful-configuration
  prefix ipv6-prefix/prefix-length
  autonomous
  on-link
  preferred-lifetime {seconds | infinite}
  valid-lifetime {seconds | infinite}
  reachable-time milli-seconds
  retransmit-time milli-seconds
  router-lifetime seconds
  no shutdown
  use-virtual-mac
```

Output example

The following is a sample router advertisement configuration output.

```
*A:sim131>config>router>router-advert# info
-----
  interface "n1"
    prefix 2001:db8::/64
    exit
    use-virtual-mac
    no shutdown
  exit
-----
*A:sim131>config>router>router-advert# interface n1
*A:sim131>config>router>router-advert>if# prefix 2001:db8::/64
*A:sim131>config>router>router-advert>if>prefix# info detail
-----
  autonomous
  on-link
  preferred-lifetime 604800
  valid-lifetime 2592000
-----
*A:tahi>config>router>router-advert>if>prefix#
```

2.4.3.4 Configuring proxy ARP

To configure proxy ARP, you can configure the following:

- a prefix list in the **config>router>policy-options>prefix-list** context
- a route policy statement in the **config>router>policy-options>policy-statement** context and apply the specified prefix list
 - In the policy statement **entry>to** context, specify the host source addresses for which ARP requests can or cannot be forwarded to non-local networks, depending on the specified action.
 - In the policy statement **entry>from** context, specify network prefixes that ARP requests will or will not be forwarded to depending on the action if a match is found. See the 7210 SAS-Mxp, R6, R12, S, Sx, T Routing Protocols Guide for more information about route policies.
- apply the policy statement to the **proxy-arp** configuration in the **config>router>interface** context

```
config>router# policy-options
begin
  commit
  prefix-list name
    prefix ip-prefix/mask [exact | longer | through
    length | prefix-length-range length1-length2]
```

Use the following syntax to configure the policy statement specified in the **proxy-arp-policy policy-statement** command.

```
config>router# policy-options
begin
  commit
  policy-statement name
    default-action {accept | next-entry | next-policy | reject}
    entry entry-id
      action {accept | next-entry | next-policy | reject}
      to
        prefix-list name [name...(upto 5 max)]
      from
        - prefix-list name [name...(upto 5 max)]
```

Output example

The following is a sample prefix list and policy statement configuration output.

```
A:ALA-49>config>router>policy-options# info
-----
  prefix-list "prefixlist1"
    prefix 10.20.30.0/24 through 32
  exit
  prefix-list "prefixlist2"
    prefix 10.10.10.0/24 through 32
  exit
  ...
  policy-statement "ProxyARPolicy"
    entry 10
      from
        prefix-list "prefixlist1"
      exit
      to
        prefix-list "prefixlist2"
```

```
        exit
        action reject
    exit
    default-action accept
    exit
exit
...
-----
A:ALA-49>config>router>policy-options#
```

Use the following syntax to configure proxy ARP.

```
config>router>interface interface-name
local-proxy-arp
proxy-arp-policy policy-name [policy-name...(upto 5 max)]
remote-proxy-arp
```

Output example

The following is a sample proxy ARP configuration output.

```
A:ALA-49>config>router>if# info
-----
address 192.0.2.59/24
local-proxy-arp
proxy-arp
    policy-statement "ProxyARPolicy"
exit
-----
A:ALA-49>config>router>if#
```

2.4.3.5 ECMP considerations



Note:

- IP ECMP is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.
- LDP LSR ECMP is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.
- LDP LER ECMP is not supported on any 7210 SAS platforms.

When ECMP is enabled and multiple equal-cost next-hops exit for the IGP route, the packets for this route is sprayed based on hashing routine currently supported for IPv4 packets.

When the preferred RTM entry corresponds to a regular IP route, spraying will be performed across regular IP next-hops for the prefix.

2.4.3.5.1 Configuration notes

The following information describes ECMP configuration guidelines:

- Users must allocate resources using the **config system resource-profile router ecmp max-ecmp-routes** command (on the 7210 SAS-R6 and R12, the command is **config>system>global-resource-profile>router>ecmp>max-ecmp-routes**) before ECMP can be enabled using the **config router**

ecmp command. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for more information about the resource profile command.

- LDP LER ECMP (including LDP over RSVP) is not supported. LDP LSR ECMP is supported on specific platforms. Check the release notes and see the *7210 SAS-Mxp, R6, R12, S, Sx, T MPLS Guide* for information about the platforms that support it and to learn more about it. IPv4 ECMP and LDP LSR ECMP share common set of resources in the hardware. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for information about resource allocation for IPv4 ECMP and LDP LSR ECMP.
- IPv6 ECMP is not supported. Only a single IPv6 route for a IPv6 destination is programmed in the IPv6 FIB. IPv6 routing and IPv6 IP interfaces cannot be used if IPv4 ECMP is in use (these features are mutually exclusive).

2.4.3.6 Deriving the router ID

The router ID defaults to the address specified in the system interface command. If the system interface is not configured with an IP address, then the router ID inherits the last four bytes of the MAC address. The router ID can also be manually configured in the **config>router router-id** context. On the BGP protocol level, a BGP router ID can be defined in the **config>router>bgp router-id** context and is only used within BGP.

If a new router ID is configured, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized the new router ID is used. An interim period of time can occur when different protocols use different router IDs. To force the new router ID, issue the **shutdown** and **no shutdown** commands for each protocol that uses the router ID, or restart the entire router.

Use the following syntax to configure the router ID.

```
config>router
  router-id router-id
  interface ip-int-name
    address {ip-address/mask | ip-address netmask} [broadcast all-ones | host-ones]
```

Output example

The following is a sample router ID configuration output.

```
A:ALA-4>config>router# info
#-----
# IP Configuration
#-----
      interface "system"
        address 10.10.0.4/32
      exit
      .
      .
      .
      router-id 10.10.0.4
#-----
A:ALA-4>config>router#
```

2.4.3.7 Configuring an autonomous system

Configuring an autonomous system is optional. Use the following syntax to configure an autonomous system.

```
config>router
  autonomous-system as-number
```

Output example

The following is a sample autonomous system configuration output.

```
A:ALA-A>config>router# info
#-----
# IP Configuration
#-----
    interface "system"
      address 10.10.10.103/32
    exit
  interface "to-104"
    address 10.0.0.103/24
    port 1/1/1
    exit
  exit
  autonomous-system 100
  router-id 10.10.10.103
#-----
A:ALA-A>config>router#
```

2.4.3.8 Configuring static routes

The 7210 SAS supports both static routes and dynamic routing to next-hop addresses.

See the *7210 SAS-Mxp, R6, R12, S, Sx, T Routing Protocols Guide* for information about configuring OSPF, RIP, IS-IS, and BGP routing.

Only one next-hop IP address can be specified per IP interface for static routes.

Use the following syntax to create static route entries.

```
config>router
  static-route {ip-prefix/prefix-length} |
  {ip-prefix netmask} [preference preference] [metric metric] [tag tag] [enable | disable] next-
  hop {ip-int-name | ip-address} [bfd-enable] [ldp-sync]
```

Example

```
config>router# static-route 192.168.250.0/24 preference 5 metric 1 enable next-hop
10.200.10.3 ldp-sync
config>router# exit
```



Note:

If **ldp-sync** is enabled on a static route, the LDP synchronization timer must also be configured on the associated interface, using the **config router if ldp-sync-timer** command.

2.4.4 Service management tasks

This section describes the following service management tasks.

2.4.4.1 Changing the system name

The **system** command sets the name of the device and is used in the prompt string. Only one system name can be configured. If multiple system names are configured, the last one configured overwrites the previous entry.

Use the following syntax to change the system name.

```
config# system
      name system-name
```

Example

The following shows the command usage to change the system name.

```
A:ALA-A>config>system# name tgif
A:TGIF>config>system#
```

Output example

The following is a sample system name change configuration output.

```
A:ALA-A>config>system# name TGIF
A:TGIF>config>system# info
#-----
# System Configuration
#-----
      name "TGIF"
      location "Mt.View, CA, NE corner of FERG 1 Building"
      coordinates "37.390, -122.05500 degrees lat."
      synchronize
      snmp
      exit
      security
      snmp
      community "private" rwa version both
      exit
      exit
      . . .
#-----
A:TGIF>config>system#
```

2.4.4.2 Modifying interface parameters

Starting at the **config>router** level, navigate down to the router interface context.

Example

The following shows the command usage to modify an IP address.

```
A:ALA-A>config>router# interface "to-sr1"
```



```
A:ALA-A>config>router>if# shutdown
A:ALA-A>config>router>if# no address
A:ALA-A>config>router>if# address 10.0.0.25/24
A:ALA-A>config>router>if# no shutdown
```

Example

The following shows the command usage to modify a port.

```
A:ALA-A>config>router# interface "to-sr1"
A:ALA-A>config>router>if# shutdown
A:ALA-A>config>router>if# no port
A:ALA-A>config>router>if# port 1/1/2
A:ALA-A>config>router>if# no shutdown
```

Output example

The following is a sample interface configuration output.

```
A:ALA-A>config>router# info
#-----
# IP Configuration
#-----
      interface "system"
        address 10.0.0.103/32
      exit
      interface "to-sr1"
        address 10.0.0.25/24
        port 1/1/2
      exit
      router-id 10.10.0.3
#-----
A:ALA-A>config>router#
```

2.4.4.3 Deleting a logical IP interface

The **no** form of the **interface** command typically removes the entry, but all entity associations must be shut down or deleted before an interface can be deleted.

1. Before loop-back IP interface can be deleted, it must first be administratively disabled with the **shutdown** command.
2. After the interface has been shut down, it can then be deleted with the **no interface** command.

```
config>router
no interface ip-int-name
```

Example

```
config>router# interface test-interface
config>router>if# shutdown
config>router>if# exit
config>router# no interface test-interface
config>router#
```

2.5 IP router command reference

2.5.1 Command hierarchies

- Configuration commands
 - Router commands
 - Router BFD commands
 - Router interface commands
 - Router interface IPv6 commands
 - Router advertisement commands
- Show commands
- Clear commands
- Debug commands

2.5.1.1 Configuration commands

2.5.1.1.1 Router commands

```
config
- router [router-name]
- aggregate ip-prefix/ip-prefix-length [summary-only] blackhole
- no aggregate ip-prefix/ip-prefix-length
- autonomous-system autonomous-system
- no autonomous-system
- ecmp max-ecmp-routes
- no ecmp
- mpls-labels
- static-label-range static-range
- no static-label-range
- sr-labels start start-value end end-value
- no sr-labels
- router-id ip-address
- no router-id
- sgt-qos (See Note below)
- application dscp-app-name dscp {dscp-value | dscp-name}
- application dot1p-app-name dot1p dot1p-priority
- no application
- dscp dscp-name fc fc-name
- no dscp dscp-name
- [no] static-route {ip-prefix/prefix-length | ip-prefix netmask}
[preference preference] [metric metric] [tag tag] [enable | disable] next-hop gateway [bfd-
enable] [{cpe-check cpe-ip-address [interval seconds] [drop-count count] [log]}] [ldp-sync]
- [no] static-route {ip-prefix/prefix-length|ip-prefix netmask} [preference preference]
[metric metric] [tag tag] [enable|disable] indirect ip-address [{cpe-check cpe-ip-address
[interval seconds] [drop-count count] [log]}]
- [no] static-route {ip-prefix/prefix-length|ip-prefix netmask} [preference preference]
[metric metric] [enable|disable] black-hole
```

```

- [no]static-route {ip-prefix/prefix-length | ip-prefix netmask}
[preference preference][metric metric] [enable | disable] indirect ip-address {prefix-list
prefixlist-name [all | none]}
- [no] static-route {ip-prefix/prefix-length | ip-prefix netmask}
[preference preference][metric metric] [tag tag] [enable | disable] next-hop ip-int-name | ip-
address{prefix-list prefix-list-name [all | none]}
- [no] static-route {ip-prefix/prefix-length | ip-prefix netmask}
[preference preference][metric metric] [tag tag] [enable | disable] black-hole {prefix-list
prefix-listname [all | none]}
- [no] triggered-policy
  
```



Note:

For information about the self-generating traffic remarking **sgt-qos** commands, refer to the “Self-Generated Traffic Commands (for 7210 SAS-Mxp)” section in the *7210 SAS-Mxp, R6, R12, S, Sx, T Quality of Service Guide*.

2.5.1.1.2 Router BFD commands



Note:

Router BFD commands are only supported on 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T (network mode).

```

config
- router
- bfd
- abort
- begin
- bfd-template [32 chars max]
- no bfd-template
- echo-receive milliseconds
- no echo-receive
- multiplier [3...20]
- no multiplier
- receive-interval milliseconds
- no receive-interval
- transmit-interval milliseconds
- no transmit-interval
- commit
  
```

2.5.1.1.3 Router interface commands

```

config
- router [router-name]
- if-attribute
- admin-group group-name value group-value
- no admin-group group-name
- srlg-group group-name value group-value
- no srlg-group group-name
- [no] interface ip-int-name unnumbered mpls-tp
- accounting-policy policy-id
- no accounting-policy
- address {ip-address/mask | ip-address netmask} [broadcast {all-ones | host-
ones}]
- no address
- arp-timeout seconds
- no arp-timeout
  
```

```

- bfd transmit-interval [receive receive-interval] [multiplier multiplier] [echo-
receive echo-interval] [type iom-hw]
- no bfd
- cflowd-parameters
  - sampling {unicast|multicast} type {interface} [direction {ingress-only}]
  - no sampling {unicast|multicast}
- delayed-enable
- no delayed-enable
- description long-description-string
- no description
- egress
  - filter ip ip-filter-id
  - filter ipv6 ipv6-filter-id
  - no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
- icmp
  - [no] mask-reply
  - redirects [number seconds]
  - no redirects
  - ttl-expired [number seconds]
  - no ttl-expired
  - unreachable [number seconds]
  - no unreachable
- if-attribute
  - [no] admin-group group-name [group-name ... (up to 5 max)]
  - no admin-group
  - [no] srlg-group group-name [group-name ... (up to 5 max)]
  - no srlg-group
- ingress
  - filter ip ip-filter-id
  - no filter
  - no filter ipv6 ipv6-filter-id
  - no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
- ldp-sync-timer seconds
- no ldp-sync-timer
- [no] local-proxy-arp
- [no] loopback
- mac ieee-mac-addr
- no mac
- [no] ntp-broadcast
- port port-name
- no port
- [no] proxy-arp-policy policy-name [policy-name...(upto 5 max)]
- qos network-policy-id
- no qos
- [no] remote-proxy-arp
- secondary {ip-address/mask | ip-address netmask} [broadcast {all-ones | host-
ones}] [igp-inhibit]
  - no secondary {ip-address/mask | ip-address netmask}
  - [no] shutdown
  - static-arp ip-address ieee-address unnumbered
  - no static-arp unnumbered
  - static-arp ieee-mac-addr unnumbered
  - no static-arp unnumbered
  - tos-marking-state {trusted | untrusted}
  - no tos-marking-state
  - no unnumbered
  - unnumbered [ip-int-name | ip-address]
- route-next-hop-policy
  - abort
  - begin
  - commit
  - [no] template name
    - description description-string
    - no description
    
```

```
- [no] exclude-group ip-admin-group-name
- include-group ip-admin-group-name [pref preference]
- no include-group ip-admin-group-name
- nh-type ip
- no nh-type
- protection-type {link | node}
- no protection-type
- [no] srlg-enable
```

2.5.1.1.4 Router interface IPv6 commands

```
config
- router [router-name]
  - [no] interface ip-int-name
    - [no] ipv6
      - address ipv6-address/prefix-length [eui-64] [preferred]
      - no address ipv6-address/prefix-length
      - icmp6
        - packet-too-big [number seconds]
        - no packet-too-big
        - param-problem [number seconds]
        - no param-problem
        - redirects [number seconds]
        - no redirects
        - time-exceeded number seconds]
        - no time-exceeded
        - unreachables [number seconds]
        - no unreachables
      - link-local-address ipv6-address [preferred]
      - [no] local-proxy-nd
      - neighbor ipv6-address [mac-address]
      - no neighbor ipv6-address
      - proxy-nd-policy policy-name [ policy-name...(up to 5 max)]
      - no proxy-nd-policy
```

2.5.1.1.5 Router advertisement commands

```
config
- router
  - [no] router-advertisement
    - [no] interface ip-int-name
      - current-hop-limit number
      - no current-hop-limit
      - [no] managed-configuration
      - max-advertisement-interval seconds
      - no max-advertisement-interval
      - min-advertisement-interval seconds
      - no min-advertisement-interval
      - mtu mtu-bytes
      - no mtu
      - [no] other-stateful-configuration
      - [no] prefix ipv6-prefix/prefix-length
        - [no] autonomous
        - [no] on-link
        - preferred-lifetime {seconds | infinite}
        - no preferred-lifetime
        - valid-lifetime{seconds | infinite}
        - no valid-lifetime
```

```

- reachable-time milli-seconds
- no reachable-time
- retransmit-time milli-seconds
- no retransmit-time
- router-lifetime seconds
- no router-lifetime
- use-virtual-mac
- no use-virtual-mac
- [no] shutdown
  
```

2.5.1.2 Show commands

```

show
- router router-instance
  - aggregate [family] [active]
  - arp [ip-int-name | ip-address/mask | mac ieee-msac-address | summary] [local |
dynamic | static | managed]
  - bfd
    - bfd-template template-name
    - interface [interface-name] [family] detail
    - interface summary
    - session [src ip-address [dst ip-address] | [detail][ipv4]]
    - session [type type] [ipv4]
    - session [summary]
    - session lsp-name lsp Name [link-type {cc-only|cc-cv}] detail
  - dhcp
    - statistics [interface ip-int-name | ip-address]
    - summary
  - ecmp
  - fib slot-number [ip-prefix/prefix-length] [longer]
  - interface [{ip-address | ip-int-name] [detail}] | [summary]
  - interface [ip-address | ip-int-name] [detail]
  - interface [ip-address | ip-int-name]
  - icmp6
    - interface [interface-name]
  - interface [{ip-address | ip-int-name] [detail] [family]} | [summary] | [exclude-
services]
  - interface [family] [detail]
  - interface ip-address | ip-int-name> statistics
  - neighbor [family] [ip-address | ip-int-name | mac ieee-mac-address | summary]
  [dynamic|static|managed]
  - policy [name | prefix-list [name] | admin]
  - route-table [family] [ip-prefix [prefix-length] [longer|exact]] [protocol protocol-
name | [summary]
  - rtr-advertisement [interface interface-name] [prefix ipv6-prefix[/prefix-length]
  [conflicts]
  - sgt-qos (See Note below)
    - application [app-name] [dscp | dot1p]
    - dscp-map [dscp-name]
  - static-arp [ip-address | ip-int-name | mac ieee-mac-addr]
  - static-route [family] [{ip-prefix /mask] [ip-prefix /prefix-length] |
  [preference preference] | [next-hop ip-address] tag tag] | [detail]
  - status
  - tunnel-table [ip-address[/mask]] | [protocol protocol | sdp sdp-id] [summary]
  
```



Note:

For information about the self-generating traffic remarking **sgt-qos** commands, refer to the "Self-Generated Traffic Commands (for 7210 SAS-Mxp" section in the *7210 SAS-Mxp, R6, R12, S, Sx, T Quality of Service Guide*.

2.5.1.3 Clear commands

```
clear
- router [router-instance]
- arp {all | ip-addr | interface {ip-int-name | ip-addr}}
- bfd
  - session src-ip ip-address dst-ip ip-address
  - statistics src-ip ip-address dst-ip ip-address
- statistics alldhcp
  - statistics [ip-int-name | ip-address]
- icmp6 all
- icmp6 global
- icmp6 interface interface-name
- neighbor {all | ipv6-address}
- neighbor interface [ip-int-name | ipv6-address]
- router-advertisement all
- router-advertisement [interface interface-name]
```

2.5.1.4 Debug commands

```
debug
- router router-instance
- ip
  - [no] arp
  - icmp
  - no icmp
  - interface [ip-int-name]
  - no interface
  - [no] interface [ip-int-name | ip-address]
  - neighbor [ip-int-name]
  - packet [ip-int-name | ip-address] [headers] [protocol-id]
  - no packet [ip-int-name | ip-address]
  - route-table [ip-prefix/prefix-length] [longer]
  - no route-table
```

2.5.2 Command descriptions

- [Configuration commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

2.5.2.1 Configuration commands

- [Generic commands](#)

- [Router global commands](#)
- [Router BFD commands](#)
- [Router interface commands](#)
- [Route next-hop policy commands](#)
- [Router interface filter commands](#)
- [Router interface ICMP commands](#)
- [Interface attribute commands](#)
- [Router interface IPv6 commands](#)
- [Router advertisement commands](#)

2.5.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

config>router>interface

config>router>router-advertisement

Platforms

Supported on all 7210 SAS platforms as described in this document

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 administratively enables an entity.

Default

no shutdown

description

Syntax

description *description-string*

no description

Context

config>router>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a text description stored in the configuration file for a configuration context.

The **no** form of this command removes the description string from the context.

Parameters

description-string

Specifies the description character string. Allowed values are any string of up to 80 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.

2.5.2.1.2 Router global commands

router

Syntax

router

Context

config

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure router parameters, and interfaces.

aggregate

Syntax

aggregate *ip-prefix/ip-prefix-length* [**summary-only**] **blackhole**

no aggregate *ip-prefix/ip-prefix-length*

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates an aggregate route.

Use this command to group a number of routes with common prefixes into a single entry in the routing table. This reduces the number of routes that need to be advertised by this router and reduces the number of routes in the routing tables of downstream routers.

Both the original components and the aggregated route (source protocol aggregate) are offered to the Routing Table Manager (RTM). Subsequent policies can be configured to assign protocol-specific characteristics (BGP, IS-IS or OSPF), such as the route type or OSPF tag to aggregate routes.

Multiple entries with the same prefix but a different mask can be configured; for example, routes are aggregated to the longest mask. If one aggregate is configured as 10.0./16 and another as 10.0.0./24, then route 10.0.128/17 would be aggregated into 10.0/16, and route 10.0.0.128/25 would be aggregated into 10.0.0/24. If multiple entries are made with the same prefix and the same mask, the previous entry is overwritten.

The **no** form of this command removes the 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)

ipv6-prefix x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:d.d.d.d

x - 0 to FFFF (hexadecimal)

d - 0 to 255 (decimal)

ip-prefix-length

Specifies the mask associated with the network address expressed as a mask length.

Values

ipv4-prefix-length - 0 to 32

ipv6-prefix-length - 0 to 128

summary-only

Specifies an optional parameter that 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.

black-hole

Specifies that the route is a blackhole route. If the destination address on a packet matches this static route, it will be silently discarded.

autonomous-system

Syntax

autonomous-system autonomous-system

no autonomous-system

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the autonomous system (AS) number for the router. A router can only belong to one AS. An ASN 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.

If the ASN is changed on a router with an active BGP instance, the new ASN 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.

Parameters

autonomous-system

Specifies the autonomous system number expressed as a decimal integer.

Values 1 to 4294967295

cflowd-parameters

Syntax

cflowd-parameters

Context

config>router>interface

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure traffic sampling for the interface.

sampling

Syntax

sampling {unicast|multicast} type {interface}[direction {ingress-only}]

no sampling {unicast|multicast}

Context

config>router>interface>cflowd-parameters

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables traffic sampling for the interface. See [Configuration notes](#) for more information.

The **no** form of this command disables traffic sampling for the interface.

Default

no sampling

Parameters

unicast

Keyword to enable unicast sampling.

multicast

Keyword to enable multicast sampling.

type

Keyword to configure the cflowd sampling type.

interface

Keyword to configure interface cflowd sampling type.

direction

keyword to configure the direction of the cflowd analysis.

ingress-only

Keyword to configure the ingress direction only for cflowd analysis.

ecmp

Syntax

ecmp max-ecmp-routes

no ecmp

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables ECMP and configures the number of routes for path sharing; for example, the value 2 means two equal cost routes will be used for cost sharing. ECMP can only be used for routes learned with the same preference and same protocol. See the description on preferences in the **static-route** command. When more ECMP routes are available at the best preference than configured in **max-ecmp-routes**, then the lowest next-hop IP address algorithm is used to select the number of routes configured in **max-ecmp-routes**.



Note:

- For the 7210 SAS-T (network mode), 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp: Before enabling ECMP, user must allocate appropriate amount of resources using the command **configure>system>resource-profile>router>ecmp>max-ecmp-routes**. The value specified with this command must be less than or equal to the value specified with the command **configure>system>resource-profile>router>ecmp> max-ecmp-routes**. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for more information.
- For 7210 SAS-R6 and 7210 SAS-R12: Before enabling ECMP, user must allocate appropriate amount of resources using the **configure>system>global-resource-profile>router>ecmp>max-ecmp-routes** command. The value specified with this command must be less than or equal to the value specified with the **configure>system>global-resource-profile>router>ecmp>max-ecmp-routes** command. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for more information.

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, then IGP chooses the next-hop based on lowest router-ID while static-route chooses the next-hop based on lowest next-hop ip address.

For more information, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide*.

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 one yields the same result as entering *no ecmp*.

Values 0 to 16

mpls-labels

Syntax

mpls-labels

Context

config>router

Platforms

7210 SAS-Mxp

Description

Commands in this context configure global parameters related to MPLS labels.

Default

N/A

static-label-range

Syntax

static-label-range *static-range*

no static-label-range

Context

config>router>mpls-labels

Platforms

7210 SAS-Mxp

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

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 thus computed as {32+ static-range-1}.

Values 0 to 131040

sr-labels

Syntax

sr-labels start *start-value* **end** *end-value*

no sr-labels

Context

config>router>mpls-labels

Platforms

7210 SAS-Mxp

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 start-value

Specifies the start label value in the SRGB.

Values 18432 to 131071

Default none

end end-value

Specifies the end label value in the SRGB.

Values 18432 to 131071

Default None

router-id

Syntax

router-id *ip-address*

no router-id

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

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 **no** form of this command reverts to the default value.

Default

system interface address (also the loopback address) if a system interface address is not configured, use the last 32 bits of the chassis MAC address

Parameters

router-id

Specifies the 32 bit router ID, expressed in dotted decimal notation or as a decimal value.

static-route

Syntax

```
[no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [tag tag] [enable | disable] next-hop gateway [bfd-enable] [{cpe-check cpe-ip-address [interval seconds] [drop-count count] [log]]} [ldp-sync]
```

```
[no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [tag tag] [enable | disable] indirect ip-address [{cpe-check cpe-ip-address [interval seconds] [drop-count count] [log]]}
```

```
[no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [enable | disable] black-hole
```

```
[no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [enable | disable] indirect ip-address {prefix-list prefixlist-name [all | none]}
```

```
[no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [tag tag] [enable | disable] next-hop ip-int-name | ip-address {prefix-list prefix-list-name [all | none]}
```

```
[no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [tag tag] [enable | disable] black-hole {prefix-list prefix-listname [all | none]}
```

Context

```
config>router
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates static route entries for both the network and access routes.

When configuring a static route, either **next-hop** or **black-hole** must be configured to indicate the type of static route. Different types of static routes can be applied to the same IP prefix. If a static route that is forwarding traffic goes down, the default route will be used instead. The **preference** parameter is used to specify the order in which the routes are applied. If a blackhole static route has the same reference as another route with the same prefix, the blackhole route takes a lower precedence.

If a CPE connectivity check target address is already being used as the target address in a different static route, then **cpe-check** parameters must match. If they do not, the new configuration command will be rejected.

If a **static-route** command is issued with no **cpe-check** target but the destination prefix/netmask and next hop address matches a static route that did have an associated **cpe-check**, the **cpe-check** test will be removed from the associated static route.

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 as necessary to uniquely identify the static route must be entered.

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)
ipv6-prefix	x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
	x - 0 to FFFF (hexadecimal)
	d - 0 to 255 (decimal)

prefix-length

Specifies the mask associated with the network address expressed as a mask length.

Values

ipv4-prefix-length	0 to 32
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)

Values

ipv6-address	x:x:x:x:x:x:x[-interface]
	x:x:x:x:x:d.d.d.d[-interface]
	x - 0 to FFFF (hexadecimal)
	d - 0 to 255 (decimal)

netmask

Specifies the subnet mask, in dotted decimal notation.

Values a.b.c.d (network bits all 1 and host bits all 0)

prefix-list prefix-list-name [all | none]

Specifies the prefix-list to be considered.

preference preference

Specifies the preference of this static route versus the routes from different sources such as BGP or OSPF, expressed as a decimal integer. When modifying the preference of an existing static route, the metric will not be changed unless specified. This parameter is also used to prioritize static routes applied to the same prefix. If a blackhole static route has the

same preference as another route with the same prefix, the blackhole route takes a lower precedence. Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is according to the route preference defaults listed in the following table.

Table 8: Default route preferences

Route type	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
OSPF External	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

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 route to use is determined by the next hop with the lowest address.

Values 1 to 255

Default 5

metric *metric*

Specifies the cost metric for the static route, expressed as a decimal integer. 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.
- If there are multiple static routes with equal preferences and metrics the route with the lowest next hop will be installed.

If there are multiple routes with different preferences then the lower preference route will be installed.

Values 0 to 65535

Default 1

black-hole

Specifies the route as a blackhole route. If the destination address on a packet matches this static route, it will be silently discarded.

The **black-hole** keyword and the **next-hop** keyword are mutually exclusive. If an identical command is entered (with the exception of the **next-hop** keyword), then this static route will be replaced with the newly entered command, and unless specified, the respective defaults for preference and metric will be applied.

next-hop gateway

Specifies the directly connected next hop IP address used to reach the destination. If the next hop is over an unnumbered interface, the *ip-int-name* of the unnumbered interface (on this node) can be configured.

The **next-hop** keyword and the **black-hole** keywords are mutually exclusive. If an identical command is entered (with the exception of the **black-hole** keyword), then this static route will be replaced with the newly entered command, and unless specified, the respective defaults for preference and metric will be applied.

The *gateway* configured here can be either on the network side or the access side on this node. This address must be associated with a network directly connected to a network configured on this node.

Values *ip-int-name* 32 chars max (must start with a letter)

tag tag

Specifies a 32-bit integer tag to be added to the static route. The tag is used in route policies to control distribution of the route into other protocols.

Values 1 to 4294967295

Default 5

enable

Specifies that static routes can be administratively enabled or disabled. Use the **enable** parameter to reenabling a disabled static route. To enable a static route, it must be uniquely identified by the IP address, mask, and any other parameter that is required to identify the exact static route. The administrative state is maintained in the configuration file.

Default enable

disable

Specifies that static routes can be administratively enabled or disabled. Use the **disable** parameter to disable a static route while maintaining the static route in the configuration. To enable a static route, it must be uniquely identified by the IP address, mask, and any other parameter that is required to identify the exact static route. The administrative state is maintained in the configuration file.

Default enable

indirect ip-address

Specifies that the route is indirect and specifies the next-hop IP address used to reach the destination. The configured *ip-address* is not directly connected to a network configured on

this node. The destination can be reachable via multiple paths. The indirect address can be resolved either via a dynamic routing protocol or by another static route.

If a static route is configured with the same destination address, subnet mask, and indirect next-hop IP address as a previously configured static route, the newly configured route replaces the previous one, and unless specified, the respective defaults for **preference** and **metric** will be applied.

The *ip-address* configured for the **indirect** parameter must be on the network side of this node and be at least one hop away from the node.

Values *ip-address* a.b.c.d

bfd-enable

Specifies that the state of the static route will be associated to a BFD session between the local system and the configured next hop. This keyword cannot be configured if the next hop is **indirect** or **blackhole** keywords are specified. Supported only in Network mode.

cpe-check cpe-ip-address

Specifies the IP address of the target CPE device. ICMP pings will be sent to this target IP address. This parameter must be configured to enable the CPE connectivity feature for the associated static route. The target-ip-address cannot be in the same subnet as the static route subnet to avoid possible circular references. This option is mutually exclusive with BFD support on a specific static route.

Default no cpe-check enabled

seconds

Specifies the interval, in seconds, between ICMP pings to the target IP address.

Values 1 to 255

Default 1

count

Specifies the number of consecutive ping-replies that must be missed to declare the CPE down and to deactivate the associated static route.

Values 1 to 255

Default 3

ldp-sync

Specifies that the LDP synchronization feature is extended to a static route. When an interface comes back up after a failure, it is possible that a preferred static route using the interface as the next hop for a specific prefix is enabled before the LDP adjacency to the peer LSR comes up on this interface. When this happens, traffic on an SDP that uses the static route for the far-end address is blackholed until the LDP session comes up and the FECs exchanged. When LDP synchronization is enabled, activation of the static route is delayed until the LDP session comes up over the interface and the **ldp-sync-timer** configured on that interface has expired (see [ldp-sync-timer](#)).

triggered-policy

Syntax

triggered-policy
no triggered-policy

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables route policy reevaluation.

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 7210 SAS Mrouter, the consequences could be dramatic. It would be more effective to control changes on a peer-by-peer basis.

If the **triggered-policy** command is enabled, and a specific peer is established, and you want the peer to remain up, 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. This keeps the peer up, and the change made to a route policy is applied only to that peer or group of peers.

2.5.2.1.3 Router BFD commands



Note:

For more information about the protocols and platforms that support BFD, see [Bidirectional Forwarding Detection](#).

abort

Syntax

abort

Context

config>router>bfd

Platforms

7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T (network mode)

Description

This command discards the changes to the BFD template configuration.

```
begin
```

Syntax

```
begin
```

Context

```
config>router>bfd
```

Platforms

7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T (network mode)

Description

Commands in this context configure a BFD template.

```
bfd-template
```

Syntax

```
bfd-template [32 chars max]
```

```
no bfd-template
```

Context

```
config>router>bfd
```

Platforms

7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T (network mode)

Description

This command creates or edits a BFD template.

A BFD template defines the set of configurable parameters used by a BFD session. These parameters include the transmit and receive timers 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 behavior.

Default

```
no bfd-template
```

Parameters

32 chars max

Specifies a text string name for the template, up to 32 characters, in printable 7-bit ASCII, enclosed in double quotes.

transmit-interval

Syntax

transmit-interval *milli-seconds*

no transmit-interval

Context

config>router>bfd>bfd-template

Platforms

7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T (network mode)

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, this timer is used for CC packets.

The **no** form of this command removes the transit timer interval from the configuration.

Default

no transmit-interval

Parameters

milli-seconds

Specifies the transmit interval.

Values 10 ms to 100,000 ms in 1 ms intervals

Default 10 ms

receive-interval

Syntax

receive-interval *milli-seconds*

no receive-interval

Context

config>router>bfd>bfd-template

Platforms

7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T (network mode)

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, this timer is used for CC packets.

The **no** form of this command removes the receive timer interval from the configuration.

Default

no receive-interval

Parameters

milli-seconds

Specifies the receive timer interval.

Values 10 to 100,000 ms in 1 ms intervals

Default 10 ms

echo-receive

Syntax

echo-receive *milli-seconds*

no echo-receive

Context

config>router>bfd>bfd-template

Platforms

7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T (network mode)

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 removes the minimum echo receive interval from the configuration.

Default

no echo-receive

Parameters

milli-seconds

Specifies the echo receive interval.

Values 100 ms to 100,000 ms in 1 ms increments

Default 100 ms

multiplier

Syntax

multiplier [3...20]

no multiplier

Context

config>router>bfd>bfd-template

Platforms

7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T (network mode)

Description

This command specifies the detect multiplier used for a BFD session. If a BFD control packet is not received for a period of *multiplier x receive-interval*, the session is declared down.

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

Default

multiplier 3

Parameters

3...20

Specifies the multiplier, in integer notation.

Values 3 to 20

commit

Syntax

commit

Context

config>router>bfd

Platforms

7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T (network mode)

Description

This command saves the changes made to the BFD template configuration. Executing this command is required for all BFD commands to take effect and become persistent after a system reboot.

2.5.2.1.4 Router interface commands

interface

Syntax

[no] interface *ip-int-name*

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a logical system or a loopback IP routing or unnumbered MPLS-TP interface. When 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**. 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.

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 7210 SAS IP 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 **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.



Note:

- MPLS-TP unnumbered interfaces are only supported on 7210 SAS-T (network operating mode), 7210 SAS-R6, and 7210 SAS-R12.
- IP unnumbered interfaces are supported on all 7210 SAS platforms as described in this document, except for those operating in access-uplink mode.

- See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for information about allocating addresses toward IP subnets using the **config>system>resource-profile>max-ip-subnets** CLI command.
- Before using IPv6, resources for IPv6 routes must be allocated. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for information about the **config>system>resource-profile>max-ipv6-routes** CLI command.

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** 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.

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 of type 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. Either a unicast, multicast or broadcast remote MAC address may be configured using the **static-arp** command. Only static ARP is supported. This option is supported only on 7210 SAS-T network mode, 7210 SAS-Sx 1/10GE, 7210 SAS-Sx 10/100GE, 7210 SAS-R6 and 7210 SAS-R12.

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an accounting policy. 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.

Parameters

acct-policy-id

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

Values 1 to 99

address

Syntax

address {*ip-address/mask* | *ip-address netmask*} [**broadcast** {**all-ones** | **host-ones**}]

no address

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns an IP address, IP subnet, and broadcast address format to an IP system IP interface. Only one IP address can be associated with an IP interface.

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.

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. The **no** form of this command can only be performed when the IP interface is administratively shut down. Shutting down the IP interface will operationally stop any protocol interfaces or MPLS LSPs that explicitly reference that IP address. When a new IP address is defined, the IP interface can be administratively enabled (**no shutdown**), which reinitializes the protocol interfaces and MPLS LSPs associated with that IP interface.

If a new address is entered while another address is still active, the new address will be rejected.

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 a.b.c.d (no multicast/broadcast address)

/

Specifies 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-length* 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-addr* 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. Allowed values are integers in the range 1 to 32.

Values 1 to 32 (mask length of 32 is reserved for system IP addresses)

netmask

Specifies the subnet netmask, in dotted-decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-addr* 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 a.b.c.d (network bits all 1 and host bits all 0)

broadcast {all-ones | host-ones}

Specifies an optional **broadcast** parameter that 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 **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.

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-addr* and the *mask-length* or *mask* with all the host bits set to binary 1. 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.

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.

Values **all-ones, host-ones**

Default **host-ones**

arp-timeout

Syntax

arp-timeout *seconds*

no arp-timeout

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

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

14400 seconds (4 hours)

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 0 specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 to 65535

bfd

Syntax

bfd *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier*] [**echo-receive** *echo-interval*] [**type** *iom-hw*]

no bfd

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

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) are notified of the fault.



Note:

- These hardware sessions cannot be used for IP interfaces configured over a LAG or for BFD-over-IP interfaces with a system IP address or loopback address. LAG-based IP interfaces always use the CPM-based centralized CPU sessions on the 7210 SAS-R6 and 7210 SAS-R12, and CPU-based sessions on the 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), 7210 SAS-Sx 10/100GE, 7210 SAS-T, and 7210 SAS-Mxp with a minimum timer support of 100 ms. The user cannot configure centralized CPU sessions on the 7210 SAS-R6 and 7210 SAS-R12, and CPU-based sessions on the 7210 SAS-T for port-based IP interfaces.
- For more information about protocols and platforms that support BFD, see [Bidirectional Forwarding Detection](#).

The **no** form of this command removes BFD from the router interface, regardless of the RSVP.

Default

no bfd

Parameters

transmit-interval

Specifies the transmit interval, in milliseconds, for the BFD session.

Values 10 to 100000

Default 100

receive *receive-interval*

Specifies the receive interval, in milliseconds, for the BFD session.

Values 10 to 100000

Default 100

multiplier *multiplier*

Specifies the multiplier for the BFD session.

Values 3 to 20

Default 3

echo-receive *echo-interval*

Specifies the minimum echo receive interval, in milliseconds, for the session.

Values 100 to 100000

Default 100

type *iom-hw*

Specifies the use of IMM-based hardware BFD sessions on IMMs on:

- the 7210 SAS-R6 and 7210 SAS-R12
- hardware sessions on the 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), and 7210 SAS-Sx 10/100GE

The user must explicitly set this keyword when configuring a BFD on an IP interface that is configured on a port.

delayed-enable

Syntax

delayed-enable *seconds*

no delayed-enable

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a delay to make the interface operational by the specified number of seconds

The value is used whenever the system attempts to bring the interface operationally up.

Parameters

seconds

Specifies a delay, in seconds, to make the interface operational.

Values 1 to 1200

ldp-sync-timer

Syntax

ldp-sync-timer *seconds*

no ldp-sync-timer

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the IGP-LDP synchronization timer. This timer enables synchronization of IGP and LDP, and synchronization of static routes and LDP. When a link is restored after a failure, IGP sets the link cost to infinity and advertises it; if it is a static route, the route activation is delayed until this timer expires. The supported IGPs are OSPF and IS-IS. 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 command 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 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.

The LDP hello adjacency is brought up with the neighbor. The LDP synchronization timer is started by IGP from the time the LDP session to the neighbor is up over the interface. This synchronization timer allows time for the label-FEC bindings to be exchanged.

When the LDP synchronization timer expires, the link cost is restored and is re-advertised. IGP will announce a new best next-hop and LDP will use it if the label binding for the neighbor FEC is available.

The preceding behavior is similar for static routes. If the static route is enabled for **ldp-sync** (see [static-route](#)), the route is not enabled immediately after the interface to the next hop comes up. Routes are suppressed until the LDP adjacency with the neighbor comes up and the synchronization timer expires. The timer does not start until the LDP adjacency with the neighbor node is fully established.

If the user changes the cost of an interface, the new value is advertised at the next flooding of link attributes by IGP. However, if the LDP synchronization timer is still running, the new cost value will only be advertised after the timer expires. Also, if the currently advertised cost is different, the new cost value will 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>interface>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. That is, if the timer is still running, it will continue using 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 failed then restored interface. In this case, it will only consider this interface for forwarding after IGP re-advertised its actual cost value.

The LDP Sync Timer State is not always synchronized across to the standby CPM, so after an activity switch the timer state might not be same as it was on the previously active CPM.

The **no** form of this command disables IGP-LDP synchronization and deletes the configuration.



Note:

- IGP-LDP synchronization is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.
- Static route-LDP synchronization is supported on all 7210 SAS platforms as described in this document, except platforms operating in access-uplink mode.

For more information, see [IGP-LDP and static route-LDP synchronization](#).

Default

no ldp-sync-timer

Parameters

seconds

Specifies the time interval for the IGP-LDP synchronization timer, in seconds.

Values 1 to 1800

local-proxy-arp

Syntax

[no] local-proxy-arp

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables local proxy ARP on the interface.

Default

no local-proxy-arp

loopback

Syntax

[no] loopback

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the interface as a loopback interface.

mac

Syntax

mac *ieee-mac-addr*

no mac

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

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 reverts the MAC address of the IP interface to the default value.

Default

IP interface has a system-assigned MAC address

Parameters

ieee-mac-addr

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.

ntp-broadcast

Syntax

[no] ntp-broadcast

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables SNTP broadcasts received on the IP interface. This parameter is only valid when the SNTP **broadcast-client** global parameter is configured.

The **no** form of this command disables SNTP broadcast received on the IP interface.

Default

no ntp-broadcast

port

Syntax

port *port-name*

no port

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

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 reattempted. The *port-id* can be in one of the following forms:

- Ethernet Interfaces

If the card in the slot has MDAs, *port-id* is in the *slot_number/MDA_number/port_number* format; for example, **1/1/3** specifies port 3 of the MDA installed in MDA slot 1 on the card installed in chassis slot 1.

The encapsulation type is an property of a Ethernet network port. The port in this context can be tagged with either IEEE 802.1Q (referred to as dot1q) encapsulation or null encapsulation. Dot1q encapsulation supports multiple logical IP interfaces on a specific network port and Null encapsulation supports a single IP interface on the 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.

Parameters

port-name

Specifies the physical port identifier to associate with the IP interface.

Values

<i>port-name</i>	<i>port-id</i> [:encap-val]
encap-val	- 0 for null
	- 0 to 4094 for dot1q
<i>port-id</i>	- slot/mda/port[.channel]
lag-id	- lag-<id>
lag	- keyword
ID	- 1 to 200

proxy-arp-policy

Syntax

[no] **proxy-arp-policy** *policy-name* [*policy-name...*(up to 5 max)]

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables and configures 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 specific protocol, set of protocols, or a particular neighbor. The policy name is configured in the **config>router>policy-options** context.

Use proxy ARP so the 7210 SAS responds to ARP requests on behalf of another device. Static ARP is used when a 7210 SAS needs to know about a device on an interface that cannot or does not respond to

ARP requests. Therefore, the 7210 SAS 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 of 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 policy name must already be defined.

qos

Syntax

qos *network-policy-id*

no qos

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command associates a network QoS policy of the type "ip-interface" with an 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.

The network QoS policy of the type ip-interface allows the user to configure an ingress and an egress component. The ingress component allows user to map the EXP bits in the MPLS packets received on the IP interface to one of the eight forwarding classes, and to rate-limit the traffic per FC using ingress policers and meters. The egress component allows the user to optionally enable the marking of EXP bits in MPLS packets by configuring the MPLS EXP values for each of the forwarding classes.

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

Default

2

Parameters

network-policy-id

Specifies an existing network policy ID to associate with the IP. interface.

Values 2 to 65535

remote-proxy-arp

Syntax

[no] remote-proxy-arp

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables remote proxy ARP on the interface.

Default

no remote-proxy-arp

secondary

Syntax

secondary {*ip-address/mask* | *ip-address netmask*} [**broadcast** {**all-ones** | **host-ones**}] [**igp-inhibit**]

no secondary {*ip-address/mask* | *ip-address netmask*}

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns up to 64 secondary IP addresses to the interface, including the primary IP address. Each address can be configured in an IP address, IP subnet, or broadcast address format.

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 a.b.c.d

/

Specifies 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 *netmask* 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 *netmask* parameter indicates the complete mask that will be used in a logical 'AND' function to derive the local subnet of the IP address. A netmask of 255.255.255.255 is reserved for system IP addresses.

Values a.b.c.d (network bits all 1 and host bits all 0)

broadcast {all-ones | host-ones}

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 configured as **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

Values **all-ones** — 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 *mask* or *netmask* with all of the host bits set to binary 1. This is the default broadcast address used by an IP interface.

Default **host-ones**

igp-inhibit

Specifies that the secondary IP address should not be recognized as a local interface by the running IGP.

static-arp

Syntax

static-arp *ip-addr ieee-mac-addr* **unnumbered**

no static-arp unnumbered

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

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 particular 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 an IP router needs to know about a device on an interface that cannot or does not respond to ARP requests. Therefore, the static ARP configuration can state that if it has a packet that has a certain IP address to send it to the corresponding ARP address.



Note:

- When used within the context for an MPLS-TP unnumbered interface, the **unnumbered** parameter is only supported on 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T (network operating mode).
- When used within the context for an MPLS IP unnumbered interface, the **unnumbered** parameter is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

The **no** form of this command removes a static ARP entry.

Parameters

ip-addr

Specifies the IP address for the static ARP in IP address dotted-decimal notation.

ieee-mac-addr

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. When this command is configured, it overrides any dynamic ARP. This parameter is only supported on 7210 SAS-T network mode, 7210 SAS-R6, and 7210 SAS-R12.

static-arp

Syntax

static-arp *ieee-mac-addr unnumbered*

no static-arp

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a static Address Resolution Protocol (ARP) entry associating an unnumbered interface 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 an unnumbered interface.

If an entry for a particular unnumbered interface already exists and a new MAC address is configured for the interface, 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 the node needs to know about a device on an interface that cannot or does not respond to ARP requests. Therefore, the node 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 node responds to ARP requests on behalf of another device.

The **no** form of this command removes a static ARP entry.

Parameters

ieee-mac-addr

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.

ip-addr

Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. When this command is configured, it overrides any dynamic ARP.

tos-marking-state

Syntax

```
tos-marking-state {trusted | untrusted}  
no tos-marking-state
```

Context

```
config>router>interface
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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 when 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 specified. 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

trusted

Parameters

trusted

Specifies the default, which prevents the ToS field from being 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.

unnumbered

Syntax

unnumbered [*ip-address* | *ip-int-name*]

no unnumbered

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

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-address | *ip-int-name*

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. Nokia recommends using the system IP address as it is not associated with a particular interface and is therefore always reachable. The system IP address is the default if no *ip-address* or *ip-int-name* is configured.

2.5.2.1.5 Route next-hop policy commands

route-next-hop-policy

Syntax

route-next-hop-policy

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure route next-hop policies.

abort

Syntax

abort

Context

config>router>route-next-hop-policy

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command discards the changes that have been made to route next-hop templates during the current session.

begin

Syntax

begin

Context

config>router>route-next-hop-policy

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context edit route next-hop templates. Use the **commit** command to save edits made during the current session. Use the **abort** command to discard edits made during the current session.

commit

Syntax

commit

Context

config>router>route-next-hop-policy

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command saves the changes that have been made to route next-hop templates during the current session.

template

Syntax

[no] template *name*

Context

config>router>route-next-hop-policy

Platforms

Supported on all 7210 SAS platforms as described in this document

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 a LFA backup next-hop for a subset of prefixes which 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 ISIS interface in the global routing 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 interfaces.

The commands within the route next-hop policy template use the begin-commit-abort model.

The following are the steps needed to create and modify the template.

1. To create a template, the user enters the name of the new template directly under the **route-next-hop-policy** context.
2. To delete a template which is not in use, the user enters the **no** form of the template command under the **route-next-hop-policy** context.

3. 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** command 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.
4. The user is allowed to create or delete a template instantly when in the editing mode without the need to enter the **commit** command. Also, if the **abort** command is executed, it will have no effect on the prior deletion or creation of a template.

When the **commit** command is executed, IS-IS or OSPF will reevaluate the templates. If there are any net changes, IS-IS or OSPF will schedule a new LFA SPF to recompute the LFA next-hop for the prefixes associated with these templates.

The **no** form of this command deletes the specified template.

Parameters

name

Specifies the name of the template, up to 32 characters maximum.

description

Syntax

description *description-string*

no description

Context

config>router>route-next-hop-policy>template

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command is used to configure the description of the next-hop template.

Parameters

description-string

Specifies the description of the next-hop template, up to 80 characters maximum.

exclude-group

Syntax

[no] **exclude-group** *ip-admin-group-name*

Context

config>router>route-next-hop-policy>template

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command 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-group** and **exclude-group** configurations, the **exclude-group** configuration takes precedence. In other words, the exclude-group 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 of this command deletes the admin group exclusion constraint from the route next-hop policy template.

Parameters

ip-admin-group-name

Specifies the name of the admin group to be excluded, up to 32 characters maximum.

include-group

Syntax

include-group *ip-admin-group-name* [**pref** *preferences*]

no include-group *ip-admin-group-name*

Context

config>router>route-next-hop-policy>template

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command 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 an **include-group** configuration but also belongs to other groups which are not part of any **include-group** configuration 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 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 specific admin group name, then it is supposed to be the least preferred, or numerically the highest preference value.

When evaluating multiple **include-group** configurations 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.

If the same group name is part of both **include-group** and **exclude-group** configurations, the **exclude-group** configuration takes precedence. In other words, the exclude-group 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 of this command deletes the admin group constraint from the route next-hop policy template.

Parameters

ip-admin-group-name

Specifies the name of the admin group to be included, up to 32 characters maximum.

preferences

Specifies the relative preference of a group, with 1 corresponding to the highest preference and 255 corresponding to the lowest preference.

Values 1 to 255

nh-type

Syntax

nh-type ip

no nh-type

Context

config>router>route-next-hop-policy>template

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the next-hop type for the route next-hop policy template.

The user can select IP backup next-hop.

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 of this command deletes the next-hop type constraint from the route next-hop policy template.

Parameters

ip

Specifies that IP backup next-hop is preferred.

protection-type

Syntax

```
protection-type {link | node}  
no protection-type
```

Context

```
config>router>route-next-hop-policy>template
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the protection type for the route next-hop policy template.

The user can select if link protection or node protection is preferred in the selection of a 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 of this command deletes the protection type constraint from the route next-hop policy template.

Parameters

link

Specifies that link protection is preferred.

node

Specifies that node protection is preferred.

srlg-enable

Syntax

```
[no] srlg-enable
```

Context

```
config>router>route-next-hop-policy>template
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the SRLG constraint for 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 from 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 of this command deletes the SRLG constraint from the route next-hop policy template.

2.5.2.1.6 Router interface filter commands

egress

Syntax

egress

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure egress network filter policies for the IP interface. If an egress filter is not defined, no filtering is performed.

ingress

Syntax

ingress

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure ingress network filter policies for the IP interface. If an ingress filter is not defined, no filtering is performed.

filter

Syntax

filter ip *ip-filter-id*

filter ipv6 *ipv6-filter-id*

no filter

Context

config>router>if>ingress

config>router>if>egress

Platforms

Supported on all 7210 SAS platforms as described in this document

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* and *ipv6-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.



Note:

For more information about service and IP interface support for different ACL match criteria per platform, see the tables in the [Applying filter policies](#) section.

The **no** form of this command removes the filter policy association with the IP interface.

Parameters

ip-filter-id

Specifies 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 65535

ipv6-filter-id

Specifies the ID for the IPv6 filter policy expressed as a decimal integer. The filter policy must already exist within the **config>filter>ip** context.

Values 1 to 65535

2.5.2.1.7 Router interface ICMP commands

icmp

Syntax

icmp

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context 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.

mask-reply

Syntax

[no] mask-reply

Context

config>router>if>icmp

Platforms

Supported on all 7210 SAS platforms as described in this document

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

redirects

Syntax

redirects [*number seconds*]

no redirects

Context

config>router>if>icmp

Platforms

Supported on all 7210 SAS platforms as described in this document

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 specific time interval.

The **no** form of this command disables the generation of ICMP redirects on the router interface.

Parameters

number

Specifies 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

Default 100

seconds

Specifies the time frame, in seconds, used to limit the *number* of ICMP redirect messages that can be issued

Values 1 to 60

Default 10

tll-expired

Syntax

tll-expired [*number seconds*]

no ttl-expired

Context

```
config>router>if>icmp
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the rate that Internet Control Message Protocol (ICMP) Time To Live (TTL) expired messages are issued by the IP interface.

The **no** form of this command disables the generation of TTL expired messages.

Parameters

number

Specifies 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 1000

Default 100

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

Default 10

unreachables

Syntax

```
unreachables [number seconds]
```

```
no unreachables
```

Context

```
config>router>if>icmp
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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 specific time interval.

The **no** form of this command disables the generation of ICMP destination unreachables on the router interface.

Parameters

number

Specifies the maximum number of ICMP unreachable messages to send, expressed as a decimal integer. The *seconds* parameter must also be specified.

Values 10 to 1000

Default 100

seconds

Specifies 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

Default 10

2.5.2.1.8 Interface attribute commands

if-attribute

Syntax

if-attribute

Context

config>router

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure or apply IP interface attributes such as administrative group (admin-group) or Shared Risk Loss Group (SRLG).

admin-group

Syntax

admin-group *group-name* **value** *group-value*

no admin-group *group-name*

Context

config>router>if-attribute

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines an administrative group (admin-group) which can be associated with an IP or MPLS interface.

Admin groups, also known as affinity, are used to tag IP and MPLS interfaces which share a specific characteristic with the same identifier. For example, an admin group identifier could represent all links which connect to core routers, all links which have bandwidth higher than 10G, or all links which are dedicated to a specific service.

The user first configures locally on each router the name and identifier of each admin group. A maximum of 32 admin groups can be configured per system.

The user then configures the admin group membership of an interface. The user can apply admin groups to a 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 admin group name. CSPF will compute a path which satisfies the admin group include and exclude constraints.

When applied to network IP interfaces, the interfaces can be included or excluded in the route next-hop selection by inferring the admin group name in a route next-hop policy template applied to an interface or a set of prefixes.

The following provisioning rules are applied to admin group configuration. The system will reject the creation of an admin group if it reuses the same name or group value as an existing group.



Note:

Only 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.

Parameters

group-name

Specifies the name of the administrative group. The association of the group name and value should be unique within an IP/MPLS domain, up to 32 characters maximum.

group-value

Specifies the value associated with the group. The association of the group name and value should be unique within an IP/MPLS domain.

Values 0 to 31

srlg-group

Syntax

srlg-group *group-name* **value** *group-value*

no srlg-group *group-name*

Context

config>router>if-attribute

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines a Shared Risk Loss Group (SRLG) which can be associated with an IP or MPLS interface.

SRLG is used to tag IP or MPLS interfaces that 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 that 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 a network IP or MPLS interface. A maximum of 64 SRLGs can be applied to a specific interface.

When SRLGs are applied to MPLS interfaces, CSPF at LER will exclude the SRLGs of interfaces used by the LSP primary path when computing the path of the secondary path. CSPF at a 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 are applied to 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 reuses 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 reuses the same group value but with a different name than an existing group.



Note:

Only the SRLGs 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.

Parameters

group-name

Specifies the name of the administrative group. The association of the group name and value should be unique within an IP/MPLS domain, up to 32 characters maximum.

group-value

Specifies the value associated with the group. The association of the group name and value should be unique within an IP/MPLS domain.

Values 0 to 4294967295

admin-group

Syntax

[no] **admin-group** *group-name* [*group-name* ... (up to 5 max)]

no admin-group

Context

config>router>interface>if-attribute

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the admin group membership of an interface. The user can apply admin groups to a network IP or MPLS interface.

Each single operation of the **admin-group** command allows a maximum of 5 groups to be specified at a time. However, a maximum of 32 groups can be added to a specific interface through multiple operations. When an admin group is bound to one or more interfaces, its value cannot be changed until all bindings are removed.

The configured admin group membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.



Note:

Only 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.

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.

Parameters

group-name

Specifies the name of an admin-group, up to 32 characters maximum.

srlg-group

Syntax

[no] **srlg-group** *group-name* [*group-name* ... (up to 5 max)]

no srlg-group

Context

config>router>interface>if-attribute

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the SRLG membership of an interface. The user can apply SRLGs to a network IP or MPLS interface.

An interface can belong to a maximum of 64 SRLG groups. However, each single operation of the **srlg-group** command allows a maximum of 5 groups to be specified at a time. When an SRLG group is bound to one or more interfaces, its value cannot be changed until all bindings are removed.

The configured SRLG membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.



Note:

Only the SRLGs 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.

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 an SRLG, up to 32 characters maximum.

2.5.2.1.9 Router interface IPv6 commands

ipv6

Syntax

[no] **ipv6**

Context

config>router>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures IPv6 for a router interface.

The **no** form of this command disables IPv6 on the interface.

Default

no ipv6

address

Syntax

address {*ipv6-address/prefix-length*} [**eui-64**]

no address {*ipv6-address/prefix-length*}

Context

config>router>if>ipv6

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns an IPv6 address to the interface.

Parameters

ip-prefix

Specifies the IPv6 address on the interface in dotted-decimal notation.

Values

ipv6-address

x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:d.d.d

x - 0 to FFFF (hexadecimal)

d - 0 to 255 (decimal)

prefix-length

Specifies the mask associated with the network address expressed as a mask length.

Values

ipv6-prefix-length 0 to 128

-

eui-64

Specifies that 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.

icmp6

Syntax

icmp6

Context

```
config>router>if>ipv6
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure ICMPv6 parameters for the interface.

packet-too-big

Syntax

packet-too-big [*number seconds*]

no packet-too-big

Context

```
config>router>if>ipv6>icmp6
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the rate for ICMPv6 packet-too-big messages.

Parameters

number

Specifies that the number of packet-too-big messages issued per the time frame specified in the *seconds* parameter will be limited.

Values 10 to 1000

seconds

Specifies 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

param-problem

Syntax

param-problem [*number seconds*]

no param-problem

Context

config>router>if>ipv6>icmp6

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the rate for ICMPv6 **param-problem** messages.

Parameters

number

Specifies that the number of **param-problem** messages issued per the time frame specified in the *seconds* parameter will be limited.

Values 10 to 1000

seconds

Specifies the time frame, in seconds, that is used to limit the number of param-problem messages issued per time frame.

Values 1 to 60

redirects

Syntax

redirects [*number seconds*]

no redirects

Context

config>router>if>ipv6>icmp6

Platforms

Supported on all 7210 SAS platforms as described in this document

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

100 10

Parameters

number

Specifies that the number of redirects issued per the time frame specified in the *seconds* parameter will be limited.

Values 10 to 1000

seconds

Specifies the time frame, in seconds, that is used to limit the number of redirects issued per time frame.

Values 1 to 60

time-exceeded

Syntax

time-exceeded [*number seconds*]

no time-exceeded

Context

config>router>if>ipv6>icmp6

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures rate for ICMPv6 time-exceeded messages.

Parameters

number

Specifies that the number of time-exceeded messages issued per the time frame specified in *seconds* parameter will be limited.

Values 10 to 1000

seconds

Specifies the time frame, in seconds, that is used to limit the number of time-exceeded messages issued per time frame.

Values 1 to 60

unreachables

Syntax

unreachables [*number seconds*]

no unreachables

Context

config>router>if>ipv6>icmp6

Platforms

Supported on all 7210 SAS platforms as described in this document

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

100 10 (when IPv6 is enabled on the interface)

Parameters

number

Specifies the number destination unreachable ICMPv6 messages to issue in the time frame specified in *seconds* parameter.

Values 10 to 1000

seconds

Sets the time frame, in seconds, to limit the number of destination unreachable ICMPv6 messages issued per time frame.

Values 1 to 60

link-local-address

Syntax

link-local-address *ipv6-address* [**preferred**]

no link-local-address

Context

config>router>if>ipv6

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the link local address.

local-proxy-nd

Syntax

[**no**] **local-proxy-nd**

Context

config>router>if>ipv6

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables local proxy neighbor discovery on the interface.

The **no** form of this command disables local proxy neighbor discovery.

proxy-nd-policy

Syntax

proxy-nd-policy *policy-name* [*policy-name...*(up to 5 max)]

no proxy-nd-policy

Context

config>router>if>ipv6

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configure a proxy neighbor discovery policy for the interface.

Parameters

policy-name

Specifies the neighbor discovery policy name. Allowed values are any string of 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. The specified policy name must already be defined.

neighbor

Syntax

neighbor [*ipv6-address*] [*mac-address*]

no neighbor [*ipv6-address*]

Context

config>router>if>ipv6

Platforms

Supported on all 7210 SAS platforms as described in this document

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

Specifies the IPv6 address assigned to a router interface.

Values

ipv6-address - x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:d.d.d.d

x - 0 to FFFF (hexadecimal)

d - 0 to 255 (decimal)

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.

2.5.2.1.10 Router advertisement commands

router-advertisement

Syntax

[no] router-advertisement

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the configuration of router advertisement properties.

The **no** form of this command disables all IPv6 interfaces. However, the **no interface** *ip-int-name* command disables a specific interface.

Default

disabled

interface

Syntax

[no] interface *ip-int-name*

Context

config>router>router-advertisement

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures router advertisement properties on a specific interface. The interface must already exist in the **config>router>interface** context.

Parameters

ip-int-name

Specifies the interface name. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

current-hop-limit

Syntax

current-hop-limit *number*

no current-hop-limit

Context

config>router>router-advertisement>if

Platforms

Supported on all 7210 SAS platforms as described in this document

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

64

Parameters

number

Specifies the hop limit.

Values 0 to 255. A value of zero means there is an unspecified number of hops.

managed-configuration

Syntax

[no] managed-configuration

Context

config>router>router-advertisement>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets the managed address configuration flag. This flag indicates that DHCPv6 is available for address configuration in addition to any address autoconfigured using stateless address autoconfiguration.

Default

no managed-configuration

max-advertisement-interval

Syntax

[no] **max-advertisement-interval** *seconds*

Context

config>router>router-advertisement>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum interval between sending router advertisement messages.

Default

600

Parameters

seconds

Specifies the maximum interval, in seconds, between sending router advertisement messages.

Values 4 to 1800

min-advertisement-interval

Syntax

[no] **min-advertisement-interval** *seconds*

Context

config>router>router-advertisement>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the minimum interval between sending ICMPv6 neighbor discovery router advertisement messages.

Default

200

Parameters

seconds

Specifies the minimum interval, in seconds, between sending ICMPv6 neighbor discovery router advertisement messages.

Values 3 to 1350

mtu

Syntax

[no] *mtu mtu-bytes*

Context

config>router>router-advertisement>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the MTU for the nodes to use to send packets on the link.

Default

no mtu

Parameters

mtu-bytes

Specifies the MTU for the nodes to use to send packets on the link.

Values 1280 to 9212

other-stateful-configuration

Syntax

[no] *other-stateful-configuration*

Context

config>router>router-advertisement>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

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

prefix

Syntax

[no] prefix [ipv6-prefix/prefix-length]

Context

config>router>router-advertisement>if

Platforms

Supported on all 7210 SAS platforms as described in this document

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

ip-prefix

Specifies the IP prefix for the prefix list entry, in dotted decimal notation.

Values

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv6-prefix	x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
	x - 0 to FFFF (hexadecimal)
	d - 0 to 255 (decimal)

prefix-length

Specifies that a route must match the most significant bits and have a prefix length.

Values	
ipv4-prefix-length	0 to 32
ipv6-prefix-length	0 to 128

autonomous

Syntax

[no] autonomous

Context

config>router>router-advertisement>if>prefix

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies whether the prefix can be used for stateless address autoconfiguration.

Default

enabled

on-link

Syntax

[no] on-link

Context

config>router>router-advertisement>if>prefix

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies whether the prefix can be used for on-link determination.

Default

enabled

preferred-lifetime

Syntax

[no] preferred-lifetime {seconds | infinite}

Context

config>router>router-advertisement>if

Platforms

Supported on all 7210 SAS platforms as described in this document

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

604800

Parameters

seconds

Specifies the remaining length of time, in seconds, that this prefix will continue to be preferred.

infinite

Specifies that the prefix will always be preferred. A value of 4,294,967,295 represents infinity.

valid-lifetime

Syntax

valid-lifetime {seconds | infinite}

Context

config>router>router-advertisement>if

Platforms

Supported on all 7210 SAS platforms as described in this document

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

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.

reachable-time

Syntax

reachable-time *milli-seconds*

no reachable-time

Context

config>router>router-advertisement>if

Platforms

Supported on all 7210 SAS platforms as described in this document

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

milli-seconds

Specifies the length of time the router should be considered reachable.

Values 0 to 3600000

retransmit-time

Syntax

retransmit-timer *milli-seconds*

no retransmit-timer

Context

config>router>router-advertisement>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the retransmission frequency of neighbor solicitation messages.

Default

no retransmit-time

Parameters

milli-seconds

Specifies how often the retransmission should occur, in milliseconds.

Values 0 to 1800000

router-lifetime

Syntax

router-lifetime *seconds*

no router-lifetime

Context

config>router>router-advertisement>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets the router lifetime.

Default

1800

Parameters

seconds

Specifies the length of time (relative to the time the packet is sent), in seconds, that the prefix is valid for route determination.

Values 0, 4 to 9000. A value of 0 means that the router is not a default router on this link.

use-virtual-mac

Syntax

[no] use-virtual-mac

Context

config>router>router-advertisement>if

Platforms

Supported on all 7210 SAS platforms as described in this document

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

2.5.2.2 Show commands

aggregate

Syntax

aggregate [family] [active]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays aggregate routes.

Parameters

active

When the active keyword is specified, inactive aggregates are filtered out.

family

Specifies the router IP interface family to display.

arp

Syntax

arp [*ip-int-name* | *ip-address/mask* | **mac** *ieee-mac-address* | **summary**] [**local** | **dynamic** | **static**]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the router ARP table sorted by IP address. If no command line options are specified, all ARP entries are displayed.

Parameters

ip-address/mask

Displays ARP entries associated with the specified IP address and mask.

ip-int-name

Displays ARP entries associated with the specified IP interface name.

mac ieee-mac-addr

Displays ARP entries associated with the specified MAC address.

summary

Displays an abbreviate list of ARP entries.

[local | dynamic | static]

Displays ARP information associated with the keyword.

Output

The following output is an example of router ARP table information, and [Table 9: Output fields: router ARP](#) describes the output fields.

Sample output

```
*B:7710-Red-RR# show router arp
=====
ARP Table (Router: Base)
=====
IP Address      MAC Address      Expiry      Type      Interface
-----
10.20.1.24      00:16:4d:23:91:b8 00h00m00s 0th      system
10.10.4.11      00:03:fa:00:d0:c9 00h57m03s Dyn[I]    to-core-sr1
10.10.4.24      00:03:fa:41:8d:20 00h00m00s 0th[I]    to-core-sr1
=====
```

```

No. of ARP Entries: 3
=====
A:ALA-A# show router ARP 10.10.0.3
=====
ARP Table
=====
IP Address      MAC Address      Expiry      Type Interface
-----
10.10.0.3      04:5d:ff:00:00:00 00:00:00    Oth  system
=====
A:ALA-A#

A:ALA-A# show router ARP to-ser1
=====
ARP Table
=====
IP Address      MAC Address      Expiry      Type Interface
-----
10.10.13.1     04:5b:01:01:00:02 03:53:09    Dyn  to-ser1
=====
A:ALA-A#
    
```

Table 9: Output fields: router ARP

Label	Description
IP Address	The IP address of the ARP entry
MAC Address	The MAC address of the ARP entry
Expiry	The age of the ARP entry
Type	Dyn — The ARP entry is a dynamic ARP entry Inv — The ARP entry is an inactive static ARP entry (invalid) Oth — The ARP entry is a local or system ARP entry Sta — The ARP entry is an active static ARP entry
*Man	The ARP entry is a managed ARP entry
Int	The ARP entry is an internal ARP entry
[!]	The ARP entry is in use
Interface	The IP interface name associated with the ARP entry
No. of ARP Entries	The number of ARP entries displayed in the list

bfd

Syntax

bfd

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context display bidirectional forwarding detection (BFD) information.



Note:

For more information about the protocols and platforms that support BFD, see [Bidirectional Forwarding Detection](#).

ecmp

Syntax

ecmp

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the ECMP settings for the router.



Note:

Weighted ECMP is not supported on 7210 SAS platforms, though it appears in the show output.

Output

The following output is an example of ECMP settings information, and [Table 10: Output fields: router ECMP](#) describes the output fields.

Sample output

```
*A:dut-d>show>router# ecmp
```

```
=====  
Router ECMP
```

```
=====
Instance      Router Name      ECMP      Max-ECMP-      Weight ECMP
              Rtes
-----
1             Base             False     n/a             False
=====
*A:dut-d>show>router#
```

Table 10: Output fields: router ECMP

Label	Description
Instance	The router instance number.
Router Name	The name of the router instance.
ECMP	False ECMP is disabled for the instance.
	True ECMP is enabled for the instance.
Max-ECMP-Rtes	Displays the maximum amount of routes to be considered for ECMP.

bfd-template

Syntax

bfd-template *template-name*

Context

show>router>bfd

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays BFD template information.

Output

The following output is an example of BFD template information, and [Table 11: Output fields: router BFD template](#) describes the output fields.

Sample output

```
*A:SASR1# show router bfd bfd-template
```

```
=====
Bfd Templates Summary
```

```

=====
Template Name          Tmpl Type  Tx Tim*  Rx Tim*  Mult  Echo Rx
                    Int
-----
my-bfd-template      iomHw     1000     1000     3     100
=====
    
```

* indicates that the corresponding row element may have been truncated.
 *A:SASR1# show router bfd session

Table 11: Output fields: router BFD template

Label	Description
Templates Name	Displays the name of the template.
Template Type	Displays the type of the template.
TX time Interval	Displays the interval, in milliseconds, between the transmitted BFD messages to maintain the session
RX time Interval	Displays the expected interval, in milliseconds, between the received BFD messages to maintain the session
Multiplier	Displays the integer used by BFD to declare when the neighbor is down.
Echo Receive Interval	Displays the echo receive interval, in milliseconds.

interface

Syntax

interface [*interface-name*]

Context

show>router>bfd

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays interface information.

Output

The following output is an example of BFD interface information, and [Table 12: Output fields: router BFD interface](#) describes the output fields.

Sample output

```

*A:7210-SAS>show>router>bfd# interface
=====
    
```

```

BFD Interface
=====
Interface name          Tx Interval    Rx Interval    Multiplier
-----
F_Port                 100           100           3
F_Lag                  300           300           3
C_Lag                  300           300           3
-----
No. of BFD Interfaces: 3
=====
*A:7210-SAS>show>router>bfd#

*A:7210-SAS>show>router>bfd# interface C_Lag

=====
BFD Interface
=====
Interface name          Tx Interval    Rx Interval    Multiplier
-----
C_Lag                  300           300           3
-----
No. of BFD Interfaces: 1
=====
*A:7210-SAS>show>router>bfd#

```

Table 12: Output fields: router BFD interface

Label	Description
TX Interval	Displays the interval, in milliseconds, between the transmitted BFD messages to maintain the session
RX Interval	Displays the expected interval, in milliseconds, between the received BFD messages to maintain the session
Multiplier	Displays the integer used by BFD to declare when the neighbor is down.

neighbor

Syntax

neighbor [*ip-int-name* | *ip-address* | **mac** *ieee-mac-address* | **summary**] [**dynamic**|**static**|**managed**]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information about the IPv6 neighbor cache.

Parameters

ip-int-name

Specifies the IP interface name.

ip-address

Specifies the address of the IPv6 interface address.

mac ieee-mac-address

Specifies the MAC address.

summary

Displays summary neighbor information.

dynamic

Specifies that the IPv6 neighbor entry is a dynamic neighbor entry.

static

Specifies that the IPv6 neighbor entry is an active static neighbor entry.

managed

Specifies that the IPv6 neighbor entry is a managed neighbor entry.

Output

The following output is an example of router neighbor information, and [Table 13: Output fields: router neighbor](#) describes the output fields.

Sample output

```
*A:Dut-A>config>router# show router neighbor
=====
Neighbor Table (Router: Base)
=====
IPv6 Address          State      Interface      Type      RTR
  MAC Address          Expiry
-----
fe80::203:faff:fe78:5c88
  00:00:1b:00:00:01    REACHABLE    -              Static    No
fe80::203:faff:fe81:6888
  e4:81:84:24:1d:6c    STALE        01h12m35s     Dynamic   Yes
-----
No. of Neighbor Entries: 2
=====
*A:Dut-A>config>router# show router neighbor dynamic
=====
Neighbor Table (Router: Base)
=====
IPv6 Address          State      Interface      Type      RTR
  MAC Address          Expiry
-----
fe80::203:faff:fe78:5c88
  e4:81:84:24:1d:6c    STALE        01h12m27s     Dynamic   Yes
-----
No. of Neighbor Entries: 1
=====
*A:Dut-A>config>router#
*A:Dut-A>config>router# show router neighbor static
```

```

=====
Neighbor Table (Router: Base)
=====
IPv6 Address      Interface
MAC Address      State      Expiry      Type      RTR
-----
fe80::203:faff:fe78:5c88
00:00:1b:00:00:01    REACHABLE    -            A_to_B2_17
                               Static      No
-----
No. of Neighbor Entries: 1
=====
*A:Dut-A>config>router# show router neighbor ma
mac      managed
*A:Dut-A>config>router# show router neighbor managed
=====
Neighbor Table (Router: Base)
=====
IPv6 Address      Interface
MAC Address      State      Expiry      Type      RTR
    
```

Table 13: Output fields: router neighbor

Label	Description
IPv6 Address	Displays the IPv6 address
Interface	Displays the name of the IPv6 interface name
MAC Address	Specifies the link-layer address
State	Displays the current administrative state
Exp	Displays the number of seconds until the entry expires
Type	Displays the type of IPv6 interface
Interface	Displays the interface name
Rtr	Specifies whether a neighbor is a router
Dynamic	The IPv6 neighbor entry is a dynamic neighbor entry
Static	The IPv6 neighbor entry is an active static neighbor entry
Managed	The IPv6 neighbor entry is a managed neighbor entry
Mtu	Displays the MTU size

session

Syntax

session [src ip-address [dst ip-address] | detail]

Context

show>router>bfd

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays session information.

Parameters

ip-address

Displays the interface information associated with the specified IP address.

Values ipv4-address a.b.c.d (host bits must be 0)

Output

The following output is an example of BFD session information, and [Table 14: Output fields: router BFD session](#) describes the output fields.

Sample output

```
*A:SASR1# show router bfd session
=====
Legend:  wp = Working path  pp = Protecting path
=====
BFD Session
=====
Interface/Lsp Name          State      Tx Intvl  Rx Intvl  Multipl
Remote Address/Info        Protocols  Tx Pkts   Rx Pkts   Type
-----
wp::unnumberedLSP          Up (3)    1000      1000      3
  4294967295::0.0.0.43     mplsTp    131       130      iom-hw
pp::unnumberedLSP          Up (3)    1000      1000      3
  4294967295::0.0.0.43     mplsTp    130       130      iom-hw
wp::numberedLSP            Up (3)    1000      1000      3
  4294967295::0.0.0.43     mplsTp    136       131      iom-hw
pp::numberedLSP            Up (3)    1000      1000      3
  4294967295::0.0.0.43     mplsTp    138       130      iom-hw
-----
No. of BFD sessions: 100
=====
* indicates that the corresponding row element may have been truncated.
*A:SASR1#

*A:7210-SAS>show>router>bfd# session
=====
BFD Session
=====
Interface          State      Tx Intvl  Rx Intvl  Mult
Remote Address     Protocol  Tx Pkts   Rx Pkts
-----
F_Port            Up (3)    100       100       3
  10.1.1.1         ospf2    801259    801275
```

```

F_Lag           Up (3)           300           300           3
  10.1.1.3      ospf2           267087        267093
C_Lag           Up (3)           300           300           3
  10.1.1.2      ospf2           267005        266996
-----
No. of BFD sessions: 3
=====
*A:7210-SAS>show>router>bfd#
    
```

Table 14: Output fields: router BFD session

Label	Description
State	Displays the administrative state for this BFD session.
Protocol	Displays the active protocol.
Tx Intvl	Displays the interval, in milliseconds, between the transmitted BFD messages to maintain the session
Tx Pkts	Displays the number of transmitted BFD packets.
Rx Intvl	Displays the expected interval, in milliseconds, between the received BFD messages to maintain the session
Rx Pkts	Displays the number of received packets.
Mult	Displays the integer used by BFD to declare when the neighbor is down.

statistics

Syntax

statistics interface [*ip-int-name*|*ip-address*]

Context

show>router>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays DHCP statistics information.

Parameters

ip-int-name

Displays statistics for the specified IP interface.

ip-address

Displays statistics for the specified IP address.

Output

The following output is an example of DHCP statistics information, and [Table 15: Output fields: router DHCP statistics](#) describes the output fields.

Sample output

```
*A:7210SAS>show>router>dhcp# statistics
=====
DHCP Global Statistics, service 1
=====
Rx Packets                : 416554
Tx Packets                : 206405
Rx Malformed Packets     : 0
Rx Untrusted Packets     : 0
Client Packets Discarded : 0
Client Packets Relayed   : 221099
Client Packets Snooped   : 0
Client Packets Proxied (RADIUS) : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded : 0
Server Packets Relayed   : 195455
Server Packets Snooped   : 0
DHCP RELEASEs Spoofed   : 0
DHCP FORCERENEWs Spoofed : 0
=====
*A:7210SAS>show>service>id>dhcp#
```

Table 15: Output fields: router DHCP statistics

Label	Description
Received Packets	The number of packets received from the DHCP clients. Includes DHCP packets received from both DHCP client and DHCP server.
Transmitted Packets	The number of packets transmitted to the DHCP clients. Includes DHCP packets transmitted from both DHCP client and DHCP server.
Received Malformed Packets	The number of corrupted/invalid packets received from the DHCP clients. Includes DHCP packets received from both DHCP client and DHCP server
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients. In this case, a frame is dropped due to the client sending a DHCP packet with Option 82 filled in before "trust" is set under the DHCP interface command.
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded.

Label	Description
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded.
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped.
Server Packets Discarded	The number of packets received from the DHCP server that were discarded.
Server Packets Relayed	The number of packets received from the DHCP server that were forwarded.
Server Packets Snooped	The number of packets received from the DHCP server that were snooped.

summary

Syntax

summary

Context

show>router>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays DHCP configuration summary information.

Output

The following output is an example of DHCP summary information, and [Table 16: Output fields: router DHCP summary](#) describes the output fields.

Sample output

```
A:7210SAS# show router dhcp summary
DHCP Summary, service 1
=====
Interface Name      Arp    Used/   Info   Admin
  SapId/Sdp         Populate Provided   Option  State
-----
egr_1               No     0/0     Replace Up
i_1                 No     0/0     Replace Up
-----
Interfaces: 2
=====
*A:7210SAS>show>service>id>dhcp#
```

Table 16: Output fields: router DHCP summary

Label	Description
Interface Name	Name of the router interface.
Arp Populate	Specifies whether ARP populate is enabled. 7210 SAS does not support ARP populate.
Used/Provided	7210 SAS does not maintain lease state.
Info Option	Indicates whether Option 82 processing is enabled on the interface.
Admin State	Indicates the administrative state.

fib

Syntax

fib *slot-number* [*ip-prefix/prefix-length* [**longer**]]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the active FIB entries for a specific IOM.

Parameters

ip-prefix/prefix-length

Displays FIB entries only matching the specified *ip-prefix* and length.

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 (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x - 0 to FFFF (hexadecimal) d - 0 to 255 (decimal)
ipv6-prefix-length	0 to 128

slot-number

Displays FIB entries only matching the specified slot number.

longer

Displays FIB entries matching the *ip-prefix/mask* and routes with longer masks.

icmp6

Syntax

icmp6

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays Internet Control Message Protocol Version 6 (ICMPv6) statistics. ICMP generates error messages (for example, ICMP destination unreachable messages) to report errors during processing and other diagnostic functions. ICMPv6 packets can be used in the neighbor discovery protocol and path MTU discovery.

Output

The following output is an example of ICMP6 information, and [Table 17: Output fields: router ICMPv6](#) describes the output fields.

Sample output

```
A:SR-3>show>router>auth# show router icmp6
=====
Global ICMPv6 Stats
=====
Received
Total                : 14          Errors                : 0
Destination Unreachable : 5          Redirects            : 5
Time Exceeded         : 0          Pkt Too Big         : 0
Echo Request          : 0          Echo Reply           : 0
Router Solicits       : 0          Router Advertisements : 4
Neighbor Solicits     : 0          Neighbor Advertisements : 0
-----
Sent
Total                : 10          Errors                : 0
Destination Unreachable : 0          Redirects            : 0
Time Exceeded         : 0          Pkt Too Big         : 0
Echo Request          : 0          Echo Reply           : 0
Router Solicits       : 0          Router Advertisements : 0
Neighbor Solicits     : 5          Neighbor Advertisements : 5
=====
A:SR-3>show>router>auth#
```

Table 17: Output fields: router ICMPv6

Label	Description
Total	The total number of all messages.
Destination Unreachable	The number of message that did not reach the destination.
Time Exceeded	The number of messages that exceeded the time threshold.
Echo Request	The number of echo requests.
Router Solicits	The number of times the local router was solicited.
Neighbor Solicits	The number of times the neighbor router was solicited.
Errors	The number of error messages.
Redirects	The number of packet redirects.
Pkt Too big	The number of packets that exceed appropriate size.
Echo Reply	The number of echo replies.
Router Advertisements	The number of times the router advertised its location.
Neighbor Advertisements	The number of times the neighbor router advertised its location.

interface

Syntax

interface [*interface-name*]

Context

show>router>icmp6

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays interface ICMPv6 statistics.

Parameters

interface-name

Displays entries associated with the specified IP interface name.

Output

[Table 18: Output fields: ICMPv6 interface](#) describes the ICMPv6 interface output fields.

Sample output

Table 18: Output fields: ICMPv6 interface

Label	Description
Total	The total number of all messages.
Destination Unreachable	The number of message that did not reach the destination.
Time Exceeded	The number of messages that exceeded the time threshold.
Echo Request	The number of echo requests.
Router Solicits	The number of times the local router was solicited.
Neighbor Solicits	The number of times the neighbor router was solicited.
Errors	The number of error messages.
Redirects	The number of packet redirects.
Pkt Too big	The number of packets that exceed appropriate size.
Echo Reply	The number of echo replies.
Router Advertisements	The number of times the router advertised its location.
Neighbor Advertisements	The number of times the neighbor router advertised its location.

interface

Syntax

interface {[*ip-address* | *ip-int-name*] [**detail**]}

interface {[*ip-address* | *ip-int-name*] [**detail**] [**family**]} | [**summary**] | [**exclude-services**]

interface *family* [**detail**]

interface [*ip-address* | *ip-int-name*]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the router IP interface table sorted by interface index.

Parameters

ip-address

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 (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
	x: [0 — FFFF]H
	d: [0 — 255]D

ip-int-name

Displays the interface information associated with the specified IP interface name.

detail

Displays detailed IP interface information.

family

Specifies the router IP interface family to display.

Values	ipv4 — Displays the peers that are IPv6-capable.
	ipv6 — Displays the peers that are IPv6-capable.

Output

The following outputs are examples of router interface information. The associated tables describe the output fields.

- Standard output: [Sample output, Table 19: Output fields: router interface](#)
- Detailed output: [Sample output — detailed, Table 20: Output fields: router interface detail](#)

Sample output

```
*A:SASR1>config>router# show router interface
=====
Interface Table (Router: Base)
=====
Interface-Name      Adm      Opr(v4/v6)  Mode      Port/SapId
IP-Address          PfxState
-----
if1                  Up        Up/Down     Network   1/1/8:1
 10.1.1.1/24         n/a
if1-1                Up        Up/Down     Network   2/1/1:1
 Unnumbered If[system] n/a
if2                  Up        Up/Down     Unnumb*   5/1/1:1
 Unnumbered If[system] n/a
if2-1                Up        Up/Down     Network   6/1/1:1
 10.2.2.1/24         n/a
```

```

system                               Up           Up/Down     Network system
10.100.100.1/32                       n/a
-----
Interfaces : 5
=====
* indicates that the corresponding row element may have been truncated.
*A:SASR1>config>router#

A:ALU-7210# show router interface
=====
Interface Table (Router: Base)
=====
Interface-Name      Adm         Opr         Mode      Port/SapId
IP-Address                               PfxState
-----
system              Up          Up          Network   system
192.0.2.169/32     n/a
-----
Interfaces : 1
=====
A:ALU-7210#
    
```

Table 19: Output fields: router interface

Label	Description
Interface-Name	The IP interface name.
Type	n/a No IP address has been assigned to the IP interface, so the IP address type is not applicable.
	Pri The IP address for the IP interface is the Primary address on the IP interface.
IP-Address	The IP address and subnet mask length of the IP interface.
	n/a Indicates no IP address is assigned to the IP interface.
Adm	Down The IP interface is administratively disabled.
	Up The IP interface is administratively enabled.
Opr	Down The IP interface is operationally disabled.
	Up The IP interface is operationally enabled.

Label	Description
Mode	Network The IP interface is a network/core IP interface.
Port	The physical network port associated with the IP interface.

Sample output — detailed

```
A:SIM7# show router interface tosim6 detail
=====
Interface Table (Router: Base)
=====
Interface
-----
If Name       : tosim6
Admin State   : Up                               Oper State    : Up
Protocols     : None
IP Addr/mask  : 10.0.0.7/24                       Address Type  : Primary
IGP Inhibit   : Disabled                         Broadcast Address: Host-ones
-----
Details
-----
If Index      : 5                               Virt. If Index : 5
Last Oper Chg: 01/09/2009 03:30:15             Global If Index : 4
SAP Id        : 1/1/2:0.*
TOS Marking   : Untrusted                       If Type       : IES
SNTP B.Cast   : False                           IES ID        : 100
MAC Address   : 2e:59:01:01:00:02               Arp Timeout   : 14400
IP MTU        : 1500                            Arp Timeout   : 14400

ICMP Details
Redirects     : Number - 100                     Time (seconds) - 10
Unreachables : Number - 100                     Time (seconds) - 10
TTL Expired  : Number - 100                     Time (seconds) - 10
=====
A:SIM7#

*A:ALU_SIM11>show>router>ldp# interface detail

=====
LDP Interfaces (Detail)
=====
Interface "a"
-----
Admin State      : Up                               Oper State      : Up
Hold Time        : 15                               Hello Factor    : 3
Keepalive Timeout : 30                             Keepalive Factor : 3
Transport Addr   : System                           Last Modified   : 07/06/2010 10:36:59
Active Adjacencies : 1
Tunneling        : Disabled
Lsp Name         : None

=====
*A:ALU_SIM11>show>router>ldp#

*A:Dut-C# show router 1 mvpn
=====
MVPN 1 configuration data
=====
```

```

signaling      : Bgp          auto-discovery   : Enabled
UMH Selection  : Highest-IP    intersite-shared : Enabled
vrf-import    : N/A
vrf-export    : N/A
vrf-target    : target:1:1
C-Mcast Import RT : target:10.20.1.3:2

ipmsi         : pim-asm 224.0.0.0
admin status  : Up          three-way-hello  : N/A
hello-interval : N/A        hello-multiplier : 35 * 0.1
tracking support : Disabled    Improved Assert  : N/A

spmsi         : pim-ssm 224.0.0.0/32
join-tlv-packing : N/A
data-delay-interval: 3 seconds
data-threshold : 224.0.0.0/4 --> 1 kbps

=====
    
```

Table 20: Output fields: router interface detail

Label	Description
If Name	The IP interface name.
Admin State	Down — The IP interface is administratively disabled.
	Up — The IP interface is administratively enabled.
Oper State	Down — The IP interface is operationally disabled.
	Up — The IP interface is operationally enabled.
IP Addr/mask	The IP address and subnet mask length of the IP interface. Not Assigned — Indicates no IP address has been assigned to the IP interface.
If Index	The interface index of the IP router interface.
Virt If Index	The virtual interface index of the IP router interface.
Last Oper Change	The last change in operational status.
Global If Index	The global interface index of the IP router interface.
If Type	Network — The IP interface is a network/core IP interface.
SNTP B.cast	Displays if the broadcast-client global parameter is configured.
QoS Policy	The QoS policy ID associated with the IP interface.
MAC Address	The MAC address of the interface.
Arp Timeout	The ARP timeout for the interface, in seconds, which is the time an ARP entry is maintained in the ARP cache without being refreshed.

Label	Description
ICMP Mask Reply	False — The IP interface will not reply to a received ICMP mask request. True — The IP interface will reply to a received ICMP mask request.
Arp Populate	Displays whether ARP is enabled or disabled.

policy

Syntax

policy [*name* | **prefix-list** *name* | **admin**]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays policy-related information.

Parameters

name

Specifies an existing policy-statement name.

prefix-list *name*

Specifies a prefix list name to display the route policy entries.

admin

Specifies the admin keyword to display the entities configured in the config>router>policy-options context.

route-table

Syntax

route-table [*ip-address*[*mask*] [**longer** | **exact**]] | [**summary**]

route-table [*family* [*ip-prefix* [*prefix-length*] [**longer** | **exact**]] | [**protocol** *protocol-name* | [**summary**]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the active routes in the routing table.

If no command line arguments are specified, all routes are displayed, sorted by prefix.

Parameters

family

Specifies the type of routing information to be distributed by this peer group.

- Values**
- ipv4** — Displays only those BGP peers that have the IPv4 family enabled and not those capable of exchanging IP-VPN routes.
 - ipv6** — Displays the BGP peers that are IPv6 capable.

ip-prefix[/*prefix-length*]

Displays routes only matching the specified *ip-address* and length.

Values

ipv4-address:	a.b.c.d (host bits must be set to 0)
ipv4-prefix-length:	0 to 32

Values

ipv6 address:	x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6 prefix-length:	1 to 128

longer

Displays routes matching the *ip-prefix/mask* and routes with longer masks.

exact

Displays the exact route matching the *ip-prefix/mask* masks.

summary

Displays a route table summary information.

Output

The following outputs are examples of route table information. The associated tables describe the output fields.

- Standard output: [Sample output, Table 21: Output fields: router route table](#)
- Summary output: [Sample output — summary](#)

Sample output

```

B:ALA-B# show router route-table 10.10.0.0 exact
=====
Route Table (Router: Base)
=====
Dest Address Next Hop Type Proto Age Metric Pref
-----
 10.10.0.0/16 Black Hole Remote Static 00h03m17s 1 5
-----
No. of Routes: 1
=====
B:ALA-B#
    
```

Table 21: Output fields: router route table

Label	Description
Dest Address	The route destination address and mask.
Next Hop	The next hop IP address for the route destination.
Type	Local — The route is a local route. Remote — The route is a remote route.
Protocol	The protocol through which the route was learned.
Age	The route age, in seconds, for the route.
Metric	The route metric value for the route.

Sample output — summary

```

A:ALA-A# show router route-table summary
=====
Route Table Summary
=====
                Active          Available
-----
Static          1              1
Direct         6              6
-----
Total
=====
A:ALA-A#
    
```

rtr-advertisement

Syntax

rtr-advertisement [**interface** *interface-name*] [**prefix** *ipv6-prefix[/prefix-length]*]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays router advertisement information.

If no command line arguments are specified, all routes are displayed, sorted by prefix.

Parameters

interface-name

Specifies the name of the interface. Maximum of 32 characters.

ipv6-prefix[/prefix-length]

Displays routes only matching the specified *ip-address* and length.

Values

ipv6	ipv6-prefix[/pref*:	x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:d.d.d.d
		x: [0 to FFFF]H
		d: [0 to 255]D
	prefix-length:	1 to 128

Output

The following output is an example of router advertisement information, and [Table 22: Output fields: router advertisement](#) describes the output fields.

Sample output

```
A:7210SAS# show router rtr-advertisement
=====
Router Advertisement
=====
-----
Interface: interfaceNetworkNonDefault
-----
Rtr Advertisement Tx : 8           Last Sent           : 00h01m28s
Nbr Solicitation Tx  : 83          Last Sent           : 00h00m17s
Nbr Advertisement Tx : 74          Last Sent           : 00h00m25s
Rtr Advertisement Rx : 8           Rtr Solicitation Rx : 0
Nbr Advertisement Rx : 83          Nbr Solicitation Rx : 74
-----
Max Advert Interval  : 601           Min Advert Interval : 201
Managed Config      : TRUE             Other Config        : TRUE
Reachable Time       : 00h00m00s400ms  Router Lifetime     : 00h30m01s
Retransmit Time      : 00h00m00s400ms  Hop Limit           : 63
Link MTU              : 1500
-----
Prefix: 211::/120
Autonomous Flag      : FALSE           On-link flag        : FALSE
Preferred Lifetime   : 07d00h00m     Valid Lifetime      : 30d00h00m
-----
Prefix: 231::/120
```

```

Autonomous Flag      : FALSE      On-link flag        : FALSE
Preferred Lifetime   : 49710d06h   Valid Lifetime      : 49710d06h

Prefix: 241::/120
Autonomous Flag      : TRUE        On-link flag        : TRUE
Preferred Lifetime   : 00h00m00s   Valid Lifetime      : 00h00m00s

Prefix: 251::/120
Autonomous Flag      : TRUE        On-link flag        : TRUE
Preferred Lifetime   : 07d00h00m   Valid Lifetime      : 30d00h00m
-----
Advertisement from: FE80::200:FF:FE00:2
Managed Config      : FALSE        Other Config        : FALSE
Reachable Time       : 00h00m00s0ms Router Lifetime     : 00h30m00s
Retransmit Time      : 00h00m00s0ms Hop Limit           : 64
Link MTU             : 0
-----
Interface: interfaceServiceNonDefault
-----
Rtr Advertisement Tx : 8           Last Sent           : 00h06m41s
Nbr Solicitation Tx  : 166         Last Sent           : 00h00m04s
Nbr Advertisement Tx : 143         Last Sent           : 00h00m05s
Rtr Advertisement Rx : 8           Rtr Solicitation Rx : 0
Nbr Advertisement Rx : 166         Nbr Solicitation Rx : 143
-----
Max Advert Interval  : 601         Min Advert Interval : 201
Managed Config      : TRUE        Other Config        : TRUE
Reachable Time       : 00h00m00s400ms Router Lifetime     : 00h30m01s
Retransmit Time      : 00h00m00s400ms Hop Limit           : 63
Link MTU             : 1500

Prefix: 23::/120
Autonomous Flag      : FALSE      On-link flag        : FALSE
Preferred Lifetime   : infinite    Valid Lifetime      : infinite

Prefix: 24::/120
Autonomous Flag      : TRUE        On-link flag        : TRUE
Preferred Lifetime   : 00h00m00s   Valid Lifetime      : 00h00m00s

Prefix: 25::/120
Autonomous Flag      : TRUE        On-link flag        : TRUE
Preferred Lifetime   : 07d00h00m   Valid Lifetime      : 30d00h00m
-----
Advertisement from: FE80::200:FF:FE00:2
Managed Config      : FALSE        Other Config        : FALSE
Reachable Time       : 00h00m00s0ms Router Lifetime     : 00h30m00s
Retransmit Time      : 00h00m00s0ms Hop Limit           : 64
Link MTU             : 0

Prefix: 2::/120
Autonomous Flag      : TRUE        On-link flag        : TRUE
Preferred Lifetime   : 07d00h00m   Valid Lifetime      : 30d00h00m

Prefix: 23::/120
Autonomous Flag      : TRUE        On-link flag        : TRUE
Preferred Lifetime   : 07d00h00m   Valid Lifetime      : 30d00h00m

Prefix: 24::/119
Autonomous Flag      : TRUE        On-link flag        : TRUE
Preferred Lifetime   : 07d00h00m   Valid Lifetime      : 30d00h00m

Prefix: 25::/120
Autonomous Flag      : TRUE        On-link flag        : TRUE
Preferred Lifetime   : 07d00h00m   Valid Lifetime      : infinite
    
```

```

Prefix: 231::/120
Autonomous Flag      : TRUE           On-link flag       : TRUE
Preferred Lifetime  : 07d00h00m      Valid Lifetime     : 30d00h00m
-----
...
A:7210SAS#
    
```

Table 22: Output fields: router advertisement

Label	Description
Rtr Advertisement Tx/ Last Sent	The number of router advertisements sent and time since they were sent
Nbr Solicitation Tx	The number of neighbor solicitations sent and time since they were sent
Nbr Advertisement Tx	The number of neighbor advertisements sent and time since they were sent
Rtr Advertisement Rx	The number of router advertisements received and time since they were received
Nbr Advertisement Rx	The number of neighbor advertisements received and time since they were received
Max Advert Interval	The maximum interval between sending router advertisement messages
Managed Config	True — Indicates that DHCPv6 has been configured
	False — Indicates that DHCPv6 is not available for address configuration
Reachable Time	The time, in milliseconds, that a node assumes a neighbor is reachable after receiving a reachability confirmation
Retransmit Time	The time, in milliseconds, between retransmitted neighbor solicitation messages
Link MTU	The MTU number the nodes use for sending packets on the link
Rtr Solicitation Rx	The number of router solicitations received and time since they were received
Nbr Solicitation Rx	The number of neighbor solicitations received and time since they were received
Min Advert Interval	The minimum interval between sending ICMPv6 neighbor discovery router advertisement messages
Other Config	True — Indicates there are other stateful configurations
	False — Indicates there are no other stateful configurations

Label	Description
Router Lifetime	Displays the router lifetime, in seconds
Hop Limit	Displays the current hop limit

static-arp

Syntax

static-arp [*ip-addr* | *ip-int-name* | **mac** *ieee-mac-addr*]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the router static ARP table sorted by IP address. If no options are present, all ARP entries are displayed.

Parameters

ip-addr

Displays only static ARP entries associated with the specified IP address.

ip-int-name

Displays only static ARP entries associated with the specified IP interface name.

mac *ieee-mac-addr*

Displays only static ARP entries associated with the specified MAC address.

Output

The following output is an example of static ARP table information, and [Table 23: Output fields: router static-ARP](#) describes the output fields.

Sample output

```
A:ALA-A# show router static-arp
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
10.200.0.253    00:00:5a:40:00:01 00:00:00 Sta  to-ser1
10.200.1.1      00:00:5a:01:00:33 00:00:00 Inv  to-ser1a
-----
No. of ARP Entries: 1
=====
A:ALA-A#
```

```

A:ALA-A# show router static-arp 10.200.1.1
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
10.200.1.1      00:00:5a:01:00:33 00:00:00 Inv  to-ser1
=====

A:ALA-A#

A:ALA-A# show router static-arp to-ser1
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
10.200.0.253    00:00:5a:40:00:01 00:00:00 Sta  to-ser1
=====

A:ALA-A#

A:ALA-A# show router static-arp mac 00:00:5a:40:00:01
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
10.200.0.253    00:00:5a:40:00:01 00:00:00 Sta  to-ser1
=====

A:ALA-A#
    
```

Table 23: Output fields: router static-ARP

Label	Description
IP Address	The IP address of the static ARP entry.
MAC Address	The MAC address of the static ARP entry.
Age	The age of the ARP entry. Static ARPs always have 00:00:00 for the age.
Type	Inv — The ARP entry is an inactive static ARP entry (invalid).
	Sta — The ARP entry is an active static ARP entry.
Interface	The IP interface name associated with the ARP entry.
No. of ARP Entries	The number of ARP entries displayed in the list.

static-route

Syntax

static-route [**family**] [*ip-prefix lmask*] | [**preference** *preference*] | [**next-hop** *ip-address* | **tag** *tag*] [**detail**]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the static entries in the routing table. If no options are present, all static routes are displayed sorted by prefix.

Parameters

family

Specifies the type of routing information to be distributed by this peer group.

Values **ipv4** — Displays only those BGP peers that have the IPv4 family enabled and not those capable of exchanging IP-VPN routes.

ipv6 — Displays the BGP peers that are IPv6 capable.

ip-prefix/mask

Displays static routes only matching the specified *ip-prefix* and *mask*.

Values

ipv4-prefix:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 to 32

Values

ipv6-prefix:	x:x:x:x:x:x:x (eight 16-bit pieces)
	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

detail

Displays detail information.

preference preference

Displays only static routes with the specified route preference.

Values 0 to 65535

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)

Values

ipv6-address: x:x:x:x:x:x:x (eight 16-bit pieces)
 x:x:x:x:x:d.d.d.d
 x: [0 to FFFF]H
 d: [0 to 255]D

tag tag

Displays the tag used to add a 32-bit integer tag to the static route. The tag is used in route policies to control distribution of the route into other protocols.

Values 1 to 4294967295

Output

The following output is an example of static route information, and [Table 24: Output fields: static route](#) describes the output fields.

Sample output

```
A:ALA-A# show router static-route
=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.250.0/24  5    1    ID  10.200.10.1    to-ser1        Y
192.168.252.0/24  5    1    NH  10.10.0.254    n/a            N
192.168.253.0/24  5    1    NH  to-ser1        n/a            N
192.168.253.0/24  5    1    NH  10.10.0.254    n/a            N
192.168.254.0/24  4    1    BH  black-hole     n/a            Y
=====
A:ALA-A#

A:ALA-A# show router static-route 192.168.250.0/24
=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.250.0/24  5    1    ID  10.200.10.1    to-ser1        Y
=====
A:ALA-A#

A:ALA-A# show router static-route preference 4
=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.254.0/24  4    1    BH  black-hole     n/a            Y
=====
A:ALA-A#

A:ALA-A# show router static-route next-hop 10.10.0.254
```

```

=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop          Interface      Active
-----
192.168.253.0/24  5    1    NH   10.10.0.254      n/a            N
=====
A:ALA-A#
  
```

Table 24: Output fields: static route

Label	Description
IP Addr/mask	The static route destination address and mask
Pref	The route preference value for the static route
Metric	The route metric value for the static route
Type	BH — The static route is a blackhole route The next hop for this type of route is black hole ID — The static route is an indirect route, where the next hop for this type of route is the non-directly connected next hop NH — The route is a static route with a directly connected next hop. The next hop for this type of route is either the next-hop IP address or an egress IP interface name
Next Hop	The next hop for the static route destination
Protocol	The protocol through which the route was learned
Interface	The egress IP interface name for the static route n/a — Indicates there is no current egress interface because the static route is inactive or a blackhole route
Active	N — The static route is inactive; for example, the static route is disabled or the next hop IP interface is down Y — The static route is active
No. of Routes	The number of routes displayed in the list

status

Syntax

status

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the router status.

Output

The following output is an example of router status information, and [Table 25: Output fields: router status](#) describes the output fields.

Sample output



Note:

There are multiple instances of OSPF. OSPF-0 is persistent. OSPF-1 through OSPF-31 are present when that particular OSPF instance is configured.

```
*A:7210>show>router# status

=====
Router Status (Router: Base)
=====
-----
Admin State                               Oper State
-----
Router                                    Up
OSPFv2-0                                  Up
ISIS                                       Not configured
MPLS                                       Not configured
RSVP                                       Not configured
LDP                                        Not configured
BGP                                        Not configured
IGMP                                       Not configured
MLD                                        Not configured
OSPFv3                                     Down
MSDP                                       Not configured

Max IPv4 Routes                           No Limit
Max IPv6 Routes                           No Limit
Total IPv4 Routes                         27231
Total IPv4 Destinations                   13614
Total IPv6 Routes                         187
ECMP Max Routes                          2
Mcast Info Policy                        default
Triggered Policies                       No
LDP Shortcut                             Disabled
Single SFM Overload                      Disabled
IP Fast Reroute                          Disabled
=====
*A:7210>show>router#
```

Table 25: Output fields: router status

Label	Description
Router	The administrative and operational states for the router

Label	Description
OSPF	The administrative and operational states for the OSPF protocol
ISIS	The administrative and operational states for the IS-IS protocol
MPLS	The administrative and operational states for the MPLS protocol
LDP	The administrative and operational states for the LDP protocol
BGP	The administrative and operational states for the BGP protocol
Max Routes	The maximum number of routes configured for the system
Total Routes	The total number of routes in the route table
ECMP Max Routes	The number of ECMP routes configured for path sharing
Triggered Policies	No — Triggered route policy reevaluation is disabled
	Yes — Triggered route policy reevaluation is enabled

tunnel-table

Syntax

tunnel-table [*ip-address[/mask]*] [**protocol** *protocol* | **sdp** *sdp-id*] [**summary**]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays tunnel table information.

Parameters

ip-address[/mask]

Displays the specified tunnel table destination IP address and mask.

protocol protocol

Displays LDP protocol information.

sdp sdp-id

Displays information pertaining to the specified SDP.

Values 1 to 17407

summary

Displays summary tunnel table information.

Output

The following output is an example of tunnel table information, and [Table 26: Output fields: router tunnel table](#) describes the output fields.

Sample output

```
A:ALA-A>config>service# show router tunnel-table summary
=====
Tunnel Table Summary (Router: Base)
=====
-----
Active                Available
-----
LDP                   1                1
SDP                   1                1
=====
A:ALA-A>config>service#
```

Table 26: Output fields: router tunnel table

Label	Description
Destination	The route destination address and mask
Owner	Specifies the tunnel owner
Encap	Specifies the tunnel encapsulation type
Tunnel ID	Specifies the tunnel (SDP) identifier
Pref	Specifies the route preference for routes learned from the configured peer
Nexthop	The next hop for the route destination
Metric	The route metric value for the route

2.5.2.3 Clear commands

router

Syntax

router [*router-instance*]

Context

clear

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears for a the router instance in which they are entered.

Parameters

router-instance

Specifies the router name or service ID.

Values Base, management

Default Base

```
arp
```

Syntax

```
arp {all | ip-addr | interface {ip-int-name | ip-addr}}
```

Context

clear>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears all or specific ARP entries.

The scope of ARP cache entries cleared depends on the command line options specified.

Parameters

all

Clears all ARP cache entries.

ip-addr

Clears the ARP cache entry for the specified IP address.

interface ip-int-name

Clears all ARP cache entries for the IP interface with the specified name.

interface ip-addr

Clears all ARP cache entries for the specified IP interface with the specified IP address.

icmp6

Syntax

icmp6 all

icmp6 global

icmp6 interface *interface-name*

Context

clear>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears ICMP statistics.

Parameters

all

Clears all statistics.

global

Clears global statistics.

interface-name

Clears ICMP6 statistics for the specified interface.

bfd

Syntax

bfd src-ip *ip-address* **dst-ip** *ip-address*

bfd all

Context

clear>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears bidirectional forwarding (BFD) sessions and statistics.

Parameters

src-ip ip-address

Specifies the source IP address, in dotted-decimal notation.

dst-ip ip-address

Specifies the destination IP address, in dotted-decimal notation.

dhcp

Syntax

dhcp

Context

clear>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context clear DHCP related information.

statistics

Syntax

statistics [ip-address | ip-int-name]

Context

clear>router>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clear statistics for DHCP relay and snooping statistics.

If no IP address or interface name is specified, then statistics are cleared for all configured interfaces.

If an IP address or interface name is specified, then only data regarding the specified interface is cleared.

Parameters

ip-int-name | ip-address

Clears statistics for the specified IP interface.

session

Syntax

session *src-ip ip-address* **dst-ip** *ip-address*

Context

clear>router>bfd

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears BFD sessions.

Parameters

src-ip *ip-address*

Specifies the address of the local endpoint of this BFD session.

dst-ip *ip-address*

Specifies the address of the remote endpoint of this BFD session.

statistics

Syntax

statistics *src-ip ip-address* **dst-ip** *ip-address*

statistics all

Context

clear>router>bfd

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears BFD statistics.

Parameters

src-ip *ip-address*

Specifies the address of the local endpoint of this BFD session.

dst-ip *ip-address*

Specifies the address of the remote endpoint of this BFD session.

all

Clears statistics for all BFD sessions.

neighbor

Syntax

neighbor {**all** | *ip-address* [**interface** *interface-name*]}

neighbor [**interface** *ip-int-name* | *ipv6-address*]

Context

clear>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears neighbor information.

Parameters

all

Clears IPv6 neighbors.

ip-int-name

Specifies an IPv6 neighbor interface name, up to 32 characters.

ip-address

Specifies an IP neighbor address.

Values a.b.c.d

ipv6-address

Specifies an IPv6 neighbor address.

Values ipv6-address

x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:d.d.d

x - 0 to FFFF (hexadecimal)

d - 0 to 255 (decimal)

router-advertisement

Syntax

```
router-advertisement all
router-advertisement [interface interface-name]
```

Context

clear>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears all router advertisement counters.

Parameters

all

Clears all router advertisement counters for all interfaces.

interface *interface-name*

Clears router advertisement counters for the specified interface.

2.5.2.4 Debug commands

router

Syntax

```
router
```

Context

debug

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures debugging for a router instance.

Parameters

router-instance

Specifies the router name or service ID.

Values	<i>router-name:</i>	Base
	<i>service-id:</i>	1 to 2147483647
Default	Base	

ip

Syntax

ip

Context

debug>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures debugging for IP.

arp

Syntax

arp

Context

debug>router>ip

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures route table debugging.

icmp

Syntax

[no] icmp

Context

debug>router>ip

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables ICMP debugging.

icmp6

Syntax

icmp6 [*ip-int-name*]

no icmp6

Context

debug>router>ip

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables ICMP6 debugging.

interface

Syntax

[no] **interface** [*ip-int-name* | *ip-address*]

Context

debug>router>ip

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the router IP interface table sorted by interface index.

Parameters

ip-address

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 (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x - 0 to FFFF (hexadecimal) d - 0 to 255 (decimal)

ip-int-name

Displays the interface information associated with the specified IP interface name, up to 32 characters.

packet

Syntax

packet [*ip-int-name* | *ip-address*] [**headers**] [*protocol-id*]

no packet [*ip-int-name* | *ip-address*]

Context

debug>router>ip

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for IP packets.

Parameters

ip-int-name

Displays the interface information associated with the specified IP interface name, up to 32 characters.

ip-address

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 (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x - 0 to FFFF (hexadecimal) d - 0 to 255 (decimal)

headers

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)
keywords: none | crtp | crudp | egp | eigrp | encap | ether-ip | icmp | idrp
| igmp | igp | ip | isis | iso-ip | l2tp | ospf-igp | pim | pnni | ptp | rdp | rsvp |
stp | tcp | udp | vrrp
* — udp/tcp wildcard

route-table

Syntax

route-table [*ip-prefix/prefix-length*]

route-table *ip-prefix/prefix-length* **longer**

no route-table

Context

debug>router>ip

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures route table debugging.

Parameters

ip-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

Values	ipv6-prefix	x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:d.d.d
		x - 0 to FFFF (hexadecimal)

d - 0 to 255 (decimal)

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*.

3 VRRP

This chapter provides information about configuring Virtual Router Redundancy Protocol (VRRP) parameters.

3.1 VRRP overview



Note:

- VRRP for IPv4 is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.
- VRRP for IPv6 is supported only for VPRN interfaces on 7210 SAS-Mxp.
- VRRP for IPv6 does not support authentication. See IETF RFC 5798 for more information.

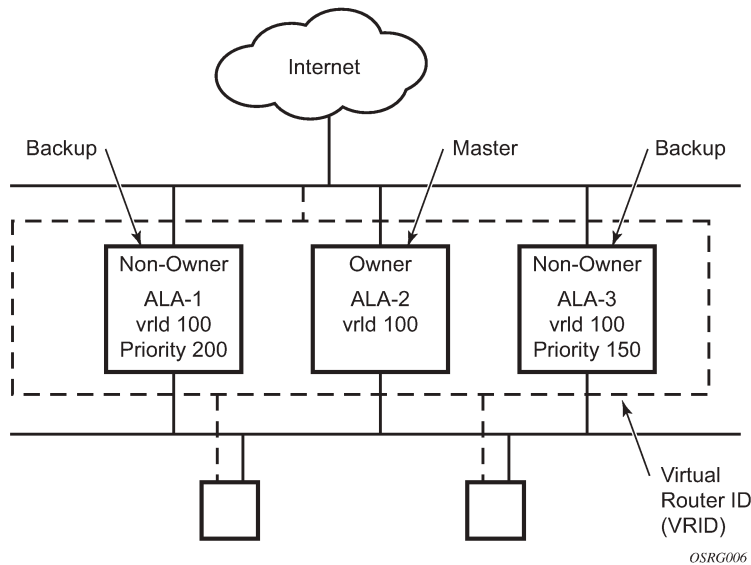
VRRP describes a method of implementing a redundant IP interface shared between two or more routers on a common LAN segment, allowing a group of routers to function as one virtual router. When this IP interface is specified as a default gateway on hosts directly attached to this LAN, the routers sharing the IP interface prevent a single point of failure by limiting access to this gateway address. VRRP can be implemented on IES service interfaces, VPRN interfaces, and on core network IP interfaces.

The VRRP standards RFC 3768 use the term "master" state to denote the virtual router that is currently acting as the active forwarding router for the VRRP instance.

If the virtual router in the master state fails, the backup router configured with the highest acceptable priority becomes the active virtual router. The new active router assumes the normal packet forwarding for the local hosts.

The following figure shows an example of a VRRP configuration.

Figure 7: VRRP configuration



3.1.1 VRRP components

VRRP consists of the following components.

3.1.1.1 Virtual router

A virtual router is a logical entity managed by VRRP that acts as a default router for hosts on a shared LAN. It consists of a Virtual Router Identifier (VRID) and a set of associated IP addresses (or an address) across a common LAN. A VRRP router can backup one or more virtual routers.

The purpose of supporting multiple IP addresses within a single virtual router is for multi-netting. This is a common mechanism that allows multiple local subnet attachment on a single routing interface. Up to four virtual routers are possible on a single IP interface. The virtual routers must be in the same subnet. Each virtual router has its own VRID, state machine, and messaging instance.

3.1.1.2 IP address owner

VRRP can be configured in either the owner or non-owner mode. The owner is the VRRP router whose virtual router IP address is the same as the real interface IP address. The VRRP router configured as the owner responds to packets that are addressed to one of the IP addresses for ICMP pings, TCP connections, and others. Other virtual router instances participating in this message domain must have the same VRID configuration and cannot be configured as owner.

The 7210 SAS allows the virtual routers to be configured as non-owners of the IP address. VRRP on a router can be configured to allow non-owners to respond to ICMP echo requests when they become the virtual router in master state for the VRRP instance. Telnet and other connection-oriented protocols can also be configured for master. However, the individual application conversations (connections) will not

survive a VRRP failover. A non-owner VRRP router operating as a backup will not respond to any packets addressed to any of the virtual router IP addresses.

3.1.1.3 Primary and secondary IP addresses

A primary address is an IP address selected from the set of real interface address. VRRP advertisements are always sent using the primary IP address as the source of the IP packet.

To be active, an IP interface must always have a primary IP address assigned for VRRP. Nokia routers support both primary and secondary IP addresses (multi-netting) on the IP interface. The virtual router VRID primary IP address is always the primary address on the IP interface. VRRP places this primary IP address in the source IP address field of the IP header for all VRRP messages sent on that interface.

3.1.1.4 Virtual router master state

The VRRP router that controls the IP addresses associated with a virtual router is considered to be in the master state, is the active router for the VRRP instance, and is responsible for forwarding packets sent to the VRRP IP address. An election process provides dynamic failover of the forwarding responsibility if the master becomes unavailable. In such an event, any of the virtual router IP addresses on the LAN can be used as the default first hop router by end-hosts. This capability enables a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every end host.

If the active router is unavailable, each backup virtual router for the VRID compares the configured priority values to determine the master role. In case of a tie, the virtual router with the highest primary IP address becomes master.

Setting the **preempt** parameter to **false** prevents a backup virtual router configured with a better priority value from becoming master when an existing non-owner virtual router is the current master. This is determined on a first-come, first-served basis.

While master, a virtual router originates all IP packets and routes them into the LAN using the physical MAC address for the IP interface as the Layer 2 source MAC address, not the VRID MAC address. ARP packets also use the parent IP interface MAC address as the Layer 2 source MAC address and insert the virtual router MAC address in the appropriate hardware address field. VRRP messages are the only packets transmitted using the virtual router MAC address as the Layer 2 source MAC.

3.1.1.5 Virtual router backup

A new virtual router master is selected from the set of VRRP routers available to assume forwarding responsibility for a virtual router in case the current master fails.

3.1.1.6 Owner and non-owner VRRP

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. Only one virtual router in the domain can be configured as owner. All other virtual router instances participating in this message domain must have the same VRID configured.

Priority is the most important parameter to be defined on a non-owner virtual router instance and defines the virtual router selection order in the master election process. The priority value and the preempt mode are used to determine the virtual router with the highest priority that can 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.

See [VRRP non-owner accessibility](#) for information about non-owner access parameters.

3.1.2 Configurable parameters

In addition to backup IP addresses, to facilitate configuration of a virtual router on routers, the following parameters can be defined in owner configurations:

- [Virtual router ID \(VRID\)](#)
- [Message interval and master inheritance](#)
- [VRRP message authentication](#)
- [Authentication data](#)
- [Virtual MAC address](#)

The following parameters can be defined in non-owner configurations:

- [Virtual router ID \(VRID\)](#)
- [Priority](#)
- [IP addresses](#)
- [Message interval and master inheritance](#)
- [Skew time](#)
- [Master down interval](#)
- [Preempt mode](#)
- [VRRP message authentication](#)
- [Authentication data](#)
- [Virtual MAC address](#)
- [VRRP advertisement message IP address list verification](#)
- [IPv6 virtual router instance operationally up](#)
- [Policies](#)

3.1.2.1 Virtual router ID (VRID)

The VRID must be configured with the same value on each virtual router associated with the redundant IP address (IP addresses). It is placed in all VRRP advertisement messages sent by each virtual router.

3.1.2.2 Priority

The priority value affects the interaction between this VRID and the same VRID of other virtual routers participating on the same LAN. A higher priority value defines a greater priority in becoming the virtual router master for the VRID. The priority value can only be configured when the defined IP address on the IP interface is different from the virtual router IP address (non-owner mode).

When the IP address on the IP interface matches the virtual router IP address (owner mode), the priority value is fixed at 255, the highest value possible. This virtual router member is considered the owner of the virtual router IP address. There can only be one owner of the virtual router IP address for all virtual router members.

The priority value 0 is reserved for VRRP advertisement message purposes. It is used to tell other virtual routers in the same VRID that this virtual router is no longer acting as master, triggering a new election process. When this happens, each backup virtual router sets its master down timer equal to the skew time value. This shortens the time until one of the backup virtual routers becomes master.

The current master virtual router must transmit a VRRP advertisement message immediately upon receipt of a VRRP message with priority set to 0. This prevents another backup from becoming master for a short period of time.

Non-owner virtual routers may be configured with a priority of 254 through 1. The default value is 100. Multiple non-owners can share the same priority value. When multiple non-owner backup virtual routers are tied (transmit VRRP advertisement messages simultaneously) in the election process, both become master simultaneously, the one with the best priority will win the election. If the priority value in the message is equal to the master local priority value, then the primary IP address of the local master and the message is evaluated as the tie breaker. The higher IP address becomes master. (The primary IP address is the source IP address of the VRRP advertisement message.)

The priority is also used to determine when to preempt the existing master. If the preempt mode value is true, VRRP advertisement messages from inferior (lower priority) masters are discarded, causing the master down timer to expire and the transition to master state.

The priority value also dictates the skew time added to the master timeout period.

3.1.2.3 IP addresses

Each virtual router participating in the same VRID should be defined with the same set of IP addresses. These are the IP addresses being used by hosts on the LAN as gateway addresses.

Multi-netting supports a total of 64 primary and secondary IP addresses on the IP interface. Up to 64 addresses can be assigned to a specific a virtual router instance.

3.1.2.4 Message interval and master inheritance

Each virtual router is configured with a message interval per VRID within which it participates. This parameter must be the same for every virtual router on the VRID.

For IPv4, the default advertisement interval is 1 s and can be configured between 1 s and 255 s 900 ms. For IPv6, the default advertisement interval is 1 s and can be configured between 1 s and 40 s 950 ms.



Note:

7210 SAS supports a minimum message interval of 1 second. It does not support use of sub-second message intervals.

As specified in the RFCs, the advertisement interval field in every received VRRP advertisement message must match the locally configured advertisement interval. If a mismatch occurs, depending on the inherit configuration, the current master's advertisement interval setting can be used to operationally override the locally configured advertisement interval setting. If the current master changes, the new master setting is used. If the local virtual router becomes master, the locally configured advertisement interval is enforced.

If a VRRP advertisement message is received with an advertisement interval set to a value different from the local value and the inherit parameter is disabled, the message is discarded without processing.

The master virtual router on a VRID uses the advertisement interval to load the advertisement timer, specifying when to send the next VRRP advertisement message. Each backup virtual router on a VRID uses the advertisement interval (with the configured local priority) to derive the master down timer value.

VRRP advertisements messages that are fragmented, or contain IP options (IPv4) or extension headers (IPv6) require a longer message interval.

3.1.2.5 Skew time

The skew time is used to add a time period to the master down interval. This is not a configurable parameter. It is derived from the current local priority of the virtual router VRID. To calculate the skew time, the virtual router evaluates the following formula:

- for IPv4 - Skew Time = $((256 - \text{priority}) / 256)$ seconds
- for IPv6 - Skew Time = $((256 - \text{priority}) * \text{Master_Adver_Interval} / 256)$ centiseconds

The higher the priority value, the smaller the skew time will be. This means that virtual routers with a lower priority will transition to master slower than virtual routers with higher priorities.

3.1.2.6 Master down interval

The master down interval is a calculated value used to load the master down timer. When the master down timer expires, the virtual router enters the master state. To calculate the master down interval, the virtual router evaluates the following formula:

Master Down Interval = $(3 \times \text{Operational Advertisement Interval}) + \text{Skew Time}$

The operational advertisement interval is dependent upon the state of the inherit parameter. When the inherit parameter is enabled, the operational advertisement interval is derived from the current master advertisement interval field in the VRRP advertisement message. When inherit is disabled, the operational advertisement interval must be equal to the locally configured advertisement interval.

The master down timer is only operational when the local virtual router is operating in backup mode.

3.1.2.7 Preempt mode

Preempt mode is a true or false configured value which controls whether a specific backup virtual router preempts a lower priority master. The IP address owner will always become master when available. Preempt mode cannot be set to false on the owner virtual router. The default value for preempt mode is true.

When preempt mode is true, the advertised priority from the incoming VRRP advertisement message from the current master is compared to the local configured priority. If the local priority is higher, the received VRRP advertisement message is discarded. This will result in the eventual expiration of the master down timer causing a transition to the master state. If the received priority is equal to the local priority, the message is not discarded and the current master will not be discarded. Note that when in the backup state, the received primary IP address is not part of the decision to preempt and is not used as a tie breaker when the received and local priorities are equal.

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.

3.1.2.8 VRRP message authentication

The authentication type parameter defines the type of authentication used by the virtual router in VRRP advertisement message authentication. VRRP message authentication is applicable to IPv4 only. The current master uses the configured authentication type to indicate any egress message manipulation that must be performed in conjunction with any supporting authentication parameters before transmitting a VRRP advertisement message. The configured authentication type value is transmitted in the message authentication type field with the appropriate authentication data field filled in. Backup routers use the authentication type message field value in interpreting the contained authentication data field within received VRRP advertisement messages.

VRRP supports two message authentication methods that provide different degrees of security. The supported authentication types are:

- 0 - no authentication
- 1 - simple text password

3.1.2.8.1 Authentication type 0 — no authentication

The use of type 0 indicates that VRRP advertisement messages are not authenticated (provides no authentication). The master transmitting VRRP advertisement messages will transmit the value 0 in the egress messages authentication type field and the authentication data field. Backup virtual routers receiving VRRP advertisement messages with the authentication type field equal to 0 will ignore the authentication data field in the message.

All compliant VRRP advertisement messages are accepted. The following fields within the received VRRP advertisement message are checked for compliance (the VRRP specification may require additional checks):

- IP header checks specific to VRRP:
 - IP header destination IP address - must be 224.0.0.18
 - IP header TTL field - must be equal to 255, the packet must not have traversed any IP routed hops
 - IP header protocol field - must be 112 (decimal)
- VRRP message checks:
 - **Version field**
Must be set to the value 2.
 - **Type field**
Must be set to the value of 1 (advertisement).
 - **Virtual router ID field**
Must match one of the configured VRID on the ingress IP interface (All other fields are dependent on matching the virtual router ID field to one of the interfaces configured VRID parameters).
 - **Priority field**

Must be equal to or greater than the VRID in-use priority or be equal to 0 (Note, equal to the VRID in-use priority and 0 requires further processing regarding master/backup and senders IP address to determine the validity of the message).

– **Authentication type field**

Must be equal to 0.

– **Advertisement interval field**

Must be equal to the VRID configured advertisement interval.

– **Checksum field**

Must be valid

– **Authentication data fields**

Must be ignored.

VRRP messages not meeting the criteria are silently dropped.

3.1.2.8.2 Authentication type 1 — simple text password

The use of type 1 indicates that VRRP advertisement messages are authenticated with a clear (simple) text password. All virtual routers participating in the virtual router instance must be configured with the same 8 octet password. Transmitting virtual routers place a value of 1 in the VRRP advertisement message authentication type field and put the configured simple text password into the message authentication data field. Receiving virtual routers compare the message authentication data field with the local configured simple text password based on the message authentication type field value of 1.

The same checks are performed for type 0 with the following exceptions (the VRRP specification may require additional checks):

- VRRP message checks:
 - authentication type field - must be equal to 1
 - authentication data fields - must be equal to the VRID configured simple text password

Any VRRP message not meeting the type 0 verification checks with the preceding exceptions are silently discarded.

3.1.2.8.3 Authentication failure

Any received VRRP advertisement message that fails authentication must be silently discarded with an invalid authentication counter incremented for the ingress virtual router instance.

3.1.2.9 Authentication data

This feature is different from the VRRP advertisement message field with the same name. This is any required authentication information that is pertinent to the configured authentication type. The type of authentication data used for each authentication type is listed in the following table.

Table 27: Authentication data type

Authentication type	Authentication data
0	None; authentication is not performed
1	Simple text password consisting of 8 octets

3.1.2.10 Virtual MAC address

On the 7210 SAS, the MAC address is not configurable. The 7210 SAS derives the MAC address to use from the VRID assigned as defined in the standard.

3.1.2.11 VRRP advertisement message IP address list verification

VRRP advertisement messages contain an IP address count field that indicates the number of IP addresses listed in the sequential IP address fields at the end of the message.

The implementation always logs mismatching events. The decision on where and whether to forward the generated messages depends on the configuration of the event manager.

To facilitate the sending of mismatch log messages, each virtual router instance keeps the mismatch state associated with each source IP address in the VRRP master table. Whenever the state changes, a mismatch log message is generated indicating the source IP address within the message, the mismatch or match event and the time of the event.

With secondary IP address support, multiple IP addresses may be found in the list and it should match the IP address on the virtual router instance. Owner and non-owner virtual router instances have the supported IP addresses explicitly defined, making mismatched supported IP address within the interconnected virtual router instances a provisioning issue.

3.1.2.12 IPv6 virtual router instance operationally up

If the IPv6 virtual router is configured with a minimum of one link-local backup address, the router advertisement of the parent interface must be configured to use the virtual MAC address for the virtual router to be considered operationally up.

3.1.2.13 Policies

Policies can be configured to control VRRP priority with the virtual router instance. 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 override or diminish the base priority 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.

Policies can only be configured in the non-owner VRRP context. For non-owner virtual router instances, if policies are not configured, then the base priority is used as the in-use priority.

3.2 VRRP priority control policies

This implementation of VRRP supports control policies to manipulate virtual router participation in the VRRP master election process and master self-deprecation. The local priority value for the virtual router instance is used to control the election process and master state.

3.2.1 VRRP virtual router policy constraints

Priority control policies can only be applied to non-owner VRRP virtual router instances. Owner VRRP virtual routers cannot be controlled by a priority control policy because they are required to have a priority value of 255 that cannot be diminished. Only one VRRP priority control policy can be applied to a non-owner virtual router instance.

Multiple VRRP virtual router instances may be associated with the same IP interface, allowing multiple priority control policies to be associated with the IP interface.

An applied VRRP priority control policy only affects the in-use priority on the virtual router instance when the preempt mode has been enabled. A virtual router instance with preempt mode disabled will always use the base priority as the in-use priority, ignoring any configured priority control policy.

3.2.2 VRRP virtual router instance base priority

Non-owner virtual router instances must have a base priority value between 1 and 254. The value 0 is reserved for master termination. The value 255 is reserved for owners. The default base priority for non-owner virtual router instances is the value 100.

The base priority is the starting priority for the VRRP instance. The actual in-use priority for the VRRP instance is derived from the base priority and an optional VRRP priority control policy.

3.2.3 VRRP priority control policy delta in-use priority limit

A VRRP priority control policy enforces an overall minimum value that the policy can inflict on the VRRP virtual router instance base priority. This value provides a lower limit to the delta priority events manipulation of the base priority.

A delta priority event is a conditional event defined in the priority control policy that subtracts a specific amount from the current, in-use priority for all VRRP virtual router instances to which the policy is applied. Multiple delta priority events can apply simultaneously, creating a dynamic priority value. The base priority for the instance, less the sum of the delta values derives the actual priority value in-use.

An explicit priority event is a conditional event defined in the priority control policy that explicitly defines the in-use priority for the virtual router instance. The explicitly defined values are not affected by the delta in-use priority limit. When multiple explicit priority events happen simultaneously, the lowest value is used for the in-use priority. The configured base priority is not a factor in explicit priority overrides of the in-use priority.

The allowed range of the Delta In-Use Priority Limit is 1 to 254. The default is 1, which prevents the delta priority events from operationally disabling the virtual router instance.

3.2.4 VRRP priority control policy priority events

The main function of a VRRP priority control policy is to define conditions or events that impact the system ability to communicate with outside hosts or portions of the network. When one or multiple of these events are true, the base priority on the virtual router instance is either overwritten with an explicit value, or a sum of delta priorities is subtracted from the base priority. The result is the in-use priority for the virtual router instance. Any priority event may be configured as an explicit event or a delta event.

Explicit events override all delta events. When multiple explicit events occur, the event with the lowest priority value is assigned to the in-use priority. As events clear, the in-use priority is reevaluated accordingly and adjusted dynamically.

Delta priority events also have priority values. When no explicit events have occurred within the policy, the sum of the occurring delta events priorities is subtracted from the base priority of each virtual router instance. If the result is lower than the delta in-use priority limit, the delta in-use priority limit is used as the in-use priority for the virtual router instance. Otherwise, the in-use priority is set to the base priority less the sum of the delta events.

Each event generates a VRRP priority event message indicating the policy-id, the event type, the priority type (delta or explicit) and the event priority value. Another log message is generated when the event is no longer true, indicating that it has been cleared.

3.2.4.1 Priority event hold-set timers

Hold-set timers are used to dampen the effect of a flapping event. A flapping event is where the event continually transitions between clear and set. 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 to count down to zero. If the timer reaches zero, the event will be allowed to enter the cleared state again. Entering the cleared state is always dependent on the object controlling the event conforming to the requirements defined in the event. It is possible, on some event types, to have a further set action reload the hold-set timer. This extends the amount of time that must expire before entering the cleared state.

See [LAG degrade priority event](#) for an example of a hold-set timer setting.

3.2.4.2 Port down priority event

The port down priority event is tied to either a physical port or a SONET/SDH channel. The port or channel operational state is evaluated to determine a port down priority event or event clear.

When the port or channel operational state is up, the port down priority event is considered false or cleared. When the port or channel operational state is down, the port down priority event is considered true or set.

3.2.4.3 LAG degrade priority event

The LAG degrade priority event is tied to an existing Link Aggregation Group (LAG). The LAG degrade priority event is conditional to percentage of available port bandwidth on the LAG. Multiple bandwidth percentage thresholds may be defined, each with its own priority value.

If the LAG transitions from one threshold to the next, the previous threshold priority value is subtracted from the total delta sum while the new threshold priority value is added to the sum. The new sum is then subtracted from the base priority and compared to the delta in-use priority limit to derive the new in-use priority on the virtual router instance.

The following example illustrates a LAG priority event and interaction with the hold-set timer in changing the in-use priority.

The following state and timer settings are used for the LAG events listed in the following table:

- user-defined thresholds - 2 ports down 3 ports down
- LAG configured ports - 4 ports
- hold-set timer (hold-set) - 5 seconds

Table 28: LAG events

Time	LAG port state	Parameter	State	Comments
0	All ports down	Event State	Set - 4 ports down	
		Event Threshold	3 ports down	
		Hold-Set Timer	5 seconds	Set to hold-set parameter
1	One port up	Event State	Set - 4 ports down	Cannot change until Hold-Set Timer expires
		Event Threshold	3 ports down	
		Hold-Set Timer	5 seconds	Event does not affect timer
2	All ports up	Event State	Set - 4 ports down	Still waiting for Hold-Set Timer expires
		Event Threshold	3 ports down	
		Hold-Set Timer	3 seconds	
5	All ports up	Event State	Cleared - All ports up	
		Event Threshold	None	Event cleared
		Hold-Set Timer	Expired	
100	Three ports down	Event State	Set - 3 ports down	
		Event Threshold	3 ports down	
		Hold-Set Timer	Expired	Set to hold-set parameter
102	Two ports down	Event State	Set - 3 ports down	
		Event Threshold	3 ports down	
		Hold-Set Timer	3 seconds	
103	All ports up	Event State	Set - 3 ports down	

Time	LAG port state	Parameter	State	Comments
		Event Threshold	3 ports down	
		Hold-Set Timer	2 second	
104	One ports down	Event State	Set - 3 ports down	
		Event Threshold	3 ports down	
		Hold-Set Timer	1 second	Current threshold is 2, so 1 down has no effect
105	One ports down	Event State	Set - 1 port down	
		Event Threshold	2 ports down	
		Hold-Set Timer	Expired	

3.2.4.4 Host unreachable priority event

The host unreachable priority event creates a continuous ping task that is used to test connectivity to a remote host. The path to the remote host and the remote host must be capable and configured to accept ICMP echo request and replies for the ping to be successful.

The ping task is controlled by interval and size parameters that define how often the ICMP request messages are transmitted and the size of each message. A historical missing reply parameter defines when the ping destination is considered unreachable.

When the host is unreachable, the host unreachable priority event is considered true or set. When the host is reachable, the host unreachable priority event is considered false or cleared.

3.2.4.5 Route unknown priority event

The route unknown priority event defines a task that monitors the existence of a specific route prefix in the system routing table.

The route monitoring task can be constrained by a condition that allows a prefix that is less specific than the defined prefix to be considered as a match. The source protocol can be defined to indicate the protocol the installed route must be populated from. To further define match criteria when multiple instances of the route prefix exist, an optional next hop parameter can be defined.

When a route prefix exists within the active route table that matches the defined match criteria, the route unknown priority event is considered false or cleared. When a route prefix does not exist within the active route table matching the defined criteria, the route unknown priority event is considered true or set.

3.3 VRRP non-owner accessibility

Although the RFC states that only VRRP owners can respond to ping and other management-oriented protocols directed to the VRID IP addresses, allows an override of this restraint on a per VRRP virtual router instance basis.

3.3.1 Non-owner access ping reply

When non-owner access ping reply is enabled on a virtual router instance, ICMP echo request messages destined for the non-owner virtual router instance IP addresses are not discarded at the IP interface when operating in master mode. ICMP echo request messages are always discarded in backup mode.

When non-owner access ping reply is disabled on a virtual router instance, ICMP echo request messages destined for the non-owner virtual router instance IP addresses are silently discarded in both the master and backup modes.

3.3.2 Non-owner access Telnet

When non-owner access Telnet is enabled on a virtual router instance, authorized Telnet sessions may be established that are destined to the virtual router instance IP addresses when operating in master mode. Telnet sessions are always discarded at the IP interface when destined to a virtual router IP address operating in backup mode. Enabling non-owner access Telnet does not guarantee Telnet access, proper management and security features must be enabled to allow Telnet on this interface and possibly from the specific source IP address.

When non-owner access Telnet is disabled on a virtual router instance, Telnet sessions destined for the non-owner virtual router instance IP addresses are silently discarded in both master and backup modes.

3.3.3 Non-owner access SSH

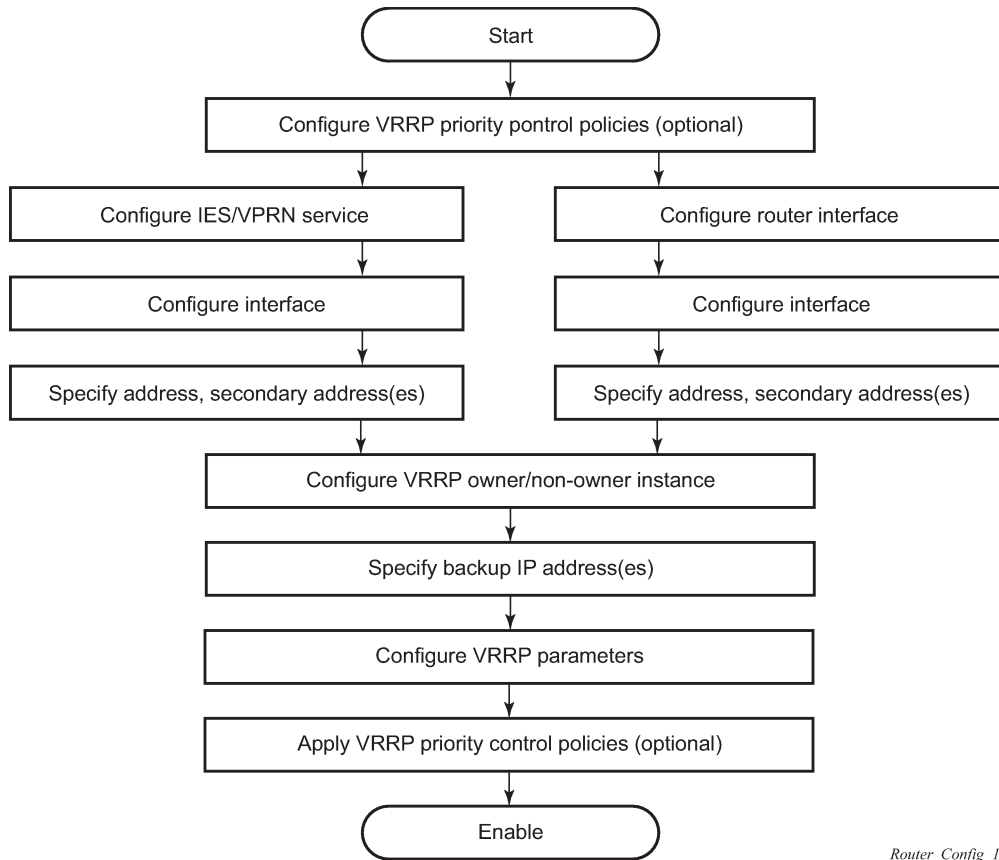
When non-owner access SSH is enabled on a virtual router instance, authorized SSH sessions may be established that are destined to the virtual router instance IP addresses when operating in master mode. SSH sessions are always discarded at the IP interface when destined for a virtual router IP address operating in backup mode. Enabling non-owner access SSH does not guarantee SSH access, proper management and security features must be enabled to allow SSH on this interface and possibly from the specific source IP address. SSH is applicable to IPv4 VRRP only.

When non-owner access SSH is disabled on a virtual router instance, SSH sessions destined for the non-owner virtual router instance IP addresses are silently discarded in both master and backup modes.

3.4 VRRP configuration process overview

The following figure shows the process to provision VRRP parameters.

Figure 8: VRRP configuration and implementation flow



Router_Config_14

3.5 Configuration notes

This section describes VRRP configuration restrictions.

3.5.1 General

- Creating and applying VRRP policies are optional.
- Backup command:
 - The backup IP addresses must be on the same subnet. The backup addresses explicitly define which IP addresses are in the VRRP advertisement message IP address list.
 - In the owner mode, the backup IP address must be identical to one of the interface IP addresses. The backup address explicitly defines which IP addresses are in the VRRP advertisement message IP address list.
 - For IPv6, one of the configured backup addresses must be the link-local address of the owner VRRP instance.

3.6 Configuring VRRP with CLI

This section provides information to configure VRRP using the command line interface.

3.7 VRRP configuration overview

Configuring VRRP policies and configuring VRRP instances on interfaces and router interfaces is optional. The basic owner and non-owner VRRP configurations on an IES or router interface must specify the **backup ip-address** parameter.

VRRP helps eliminate the single point of failure in a routed environment by using virtual router IP address shared between two or more routers connecting the common domain. VRRP provides dynamic fail over of the forwarding responsibility if the master becomes unavailable.

The VRRP implementation allows one master per IP subnet. All other VRRP instances in the same domain must be in backup mode.

3.7.1 Preconfiguration requirements

The following information describes VRRP preconfiguration requirements:

- VRRP policies:
 - VRRP policies must be configured before they can be applied to an interface or IES VRRP instance. VRRP policies are configured in the **config>vrrp** context.
- Configuring VRRP on an IES service interface:
 - The service customer account must be created before configuring an IES VRRP instance.
 - The interface address must be specified in the both the owner and non-owner IES or router interface instances.

3.8 Basic VRRP configurations

This section contains information about basic VRRP configurations.

3.8.1 VRRP policy

Configuring and applying VRRP policies are optional. There are no default VRRP policies. Each policy must be explicitly defined. A VRRP configuration must include the following:

- policy ID
- define at least one of the following priority events:
 - port down
 - LAG port down
 - host unreachable

- route unknown

Output example: Sample VRRP policy configuration output

```
A:SR2>config>vrrp>policy# info
-----
      delta-in-use-limit 50
      priority-event
        port-down /1/2
          hold-set 43200
          priority 100 delta
        exit
      port-down /1/3
        priority 200 explicit
      exit
      lag-port-down 1
        number-down 3
        priority 50 explicit
      exit
      host-unreachable 10.10.24.4
        drop-count 25
      exit
      route-unknown 10.10.0.0/32
priority 50 delta
      exit
      exit
-----
```

3.8.2 VRRP IES service parameters

VRRP parameters are configured within an IES service with two contexts, owner or non-owner. The status is specified when the VRRP configuration is created. When configured as owner, the virtual router instance owns the backup IP addresses. All other virtual router instances participating in this message domain must have the same VRID configured and cannot be configured as owner.

For IPv4, up to 4 virtual routers IDs (vrid) can be configured on an IES service interface.

VRRP parameters configured within an IES service must include the following:

- VRID
- backup IP addresses

Output example

The following is a sample IES service owner and non-owner VRRP configuration output.

```
A:SR2>config>service>ies# info
-----
      interface "tuesday" create
        address 10.10.36.2/24
        sap 7/1/1.2.2 create
        vrrp 19 owner
          backup 10.10.36.2
          authentication-type password
          authentication-key "testabc"
        exit
      exit
-----
```

```
interface "testing" create
  address 10.10.10.16/24
  sap 1/1/55:0 create
  vrrp 12
    backup 10.10.10.15
    policy 1
    authentication-type password
    authentication-key "testabc"
  exit
exit
no shutdown
-----
A:SR2>config>service>ies#
```

3.8.3 VRRP router interface parameters

VRRP parameters are configured on a router interface with two contexts, owner or non-owner. The status is specified when the VRRP configuration is created. When configured as owner, the virtual router instance owns the backed up IP addresses. All other virtual router instances participating in this message domain must have the same VRID configured and cannot be configured as owner.

For IPv4, up to 4 virtual routers IDs (VRIDs) can be configured on a router interface. For IPv6, only one virtual router instance can be configured on a router interface.

VRRP parameters configured on a router interface must include the following:

- VRID
- backup IP addresses

Output example

The following is a sample router interface owner and non-owner VRRP configuration output.

```
A:SR4>config>router# info
#-----
echo "IP Configuration "
#-----
interface "system"
  address 10.10.0.4/32
exit
interface "test1"
  address 10.10.14.1/24

exit
interface "test2"
  address 10.10.10.23/24
  vrrp 1 owner
    backup 10.10.10.23

    authentication-key "testabc"
  exit
exit
#-----
A:SR4>config>router#
```

3.9 Common configuration tasks

This section provides a brief overview of the tasks that must be performed to configure VRRP and provides the CLI commands.

VRRP parameters are defined under a service interface or a router interface context. An IP address must be assigned to each IP interface. Only one primary IP address can be associated with an IP interface, but several secondary IP addresses can also be associated.

Owner and non-owner configurations must include the following parameters:

- All participating routers in a VRRP instance must be configured with the same VRID.
- The owner configuration must include at least one backup IP address.
- For IPv6, all participating routers must be configured with the same link-local backup address (the address configured for the owner instance).

Other owner and non-owner configurations include the following optional commands:

- **authentication-key**
- **message-interval**

In addition to the common parameters, the following non-owner commands can be configured:

- **master-int-inherit**
- **priority**
- **policy**
- **ping-reply**
- **preempt**
- **telnet-reply**
- **ssh-reply** (IPv4 only)
- **[no] shutdown**

3.9.1 Creating interface parameters

If you have multiple subnets configured on an Ethernet interface, you can configure VRRP on each subnet.

Output example

The following is a sample IP interface configuration output.

```
A:SR1>config>router# info
#-----
echo "IP Configuration "
#-----
      interface "system"
        address 10.10.0.1/32
      exit
      interface "testA"
        address 10.123.123.123/24
      exit
      interface "testB"
        address 10.10.14.1/24
```

```
        secondary 10.10.16.1/24
        secondary 10.10.17.1/24
        secondary 10.10.18.1/24
    exit
    router-id 10.10.0.1
#-----
A:SR1>config>router#
```

3.10 Configuring VRRP policy components

Output example

The following is a sample VRRP policy configuration output.

```
A:SR1>config>vrrp# info
-----
    policy 1
        delta-in-use-limit 50
        priority-event
            port-down 1/1/2
                hold-set 43200
                priority 100 delta
        exit
        route-unknown 0.0.0.0/0
            protocol isis
        exit
    exit
    exit
-----
A:SR1>config>vrrp#
```

3.10.1 Configuring service VRRP parameters

VRRP parameters can be configured on an interface in a service to provide virtual default router support which allows traffic to be routed without relying on a single router in case of failure.

3.10.1.1 Non-owner VRRP example

Output example

The following is a sample basic non-owner VRRP configuration output.

```
A:SR2>config>service>ies# info
-----
...
    interface "testing" create
        address 10.10.10.16/24
        sap 1/1/55:0 create
        vrrp 12
            backup 10.10.10.15
            policy 1

            authentication-key "testabc"
        exit
```

```
        exit
        no shutdown
-----
A:SR2>config>service>ies#
```

3.10.1.2 Owner service VRRP example

Output example

The following is a sample owner VRRP configuration output.

```
A:SR4>config>router# info
#-----
echo "IP Configuration "
#-----
...
    interface "test2"
      address 10.10.10.23/24
      vrrp 1 owner
        backup 10.10.10.23

        authentication-key "testabc"
      exit
    exit
#-----
A:SR4>config>router#
```

3.10.2 Configuring router interface VRRP parameters

VRRP parameters can be configured on an interface in an interface to provide virtual default router support which allows traffic to be routed without relying on a single router in case of failure.

3.10.2.1 Router interface VRRP non-owner

Output example

The following is a sample non-owner interface VRRP configuration output.

```
A:SR2>config># info
#-----
interface "if-test"
  address 10.20.30.40/24
  secondary 10.10.50.1/24
  secondary 10.10.60.1/24
  secondary 10.10.70.1/24
  vrrp 1

      backup 10.20.30.41
      ping-reply
      telnet-reply

      authentication-key "testabc"
    exit
  exit
```



```
#-----  
A:SR2>config>#
```

3.10.2.2 Router interface VRRP owner

Output example

The following is a sample router interface owner VRRP configuration output.

```
A:SR2>config>router# info  
#-----  
interface "vrrpowner"  
    address 10.10.10.23/24  
    vrrp 1 owner  
        backup 10.10.10.23  
  
        authentication-key "testabc"  
    exit  
exit  
#-----  
A:SR2>config>router#
```

3.11 VRRP configuration management tasks

This section describes the VRRP configuration management tasks.

3.11.1 Modifying a VRRP policy

To access a specific VRRP policy, you must specify the policy ID. To display a list of VRRP policies, use the **show vrrp policy** command.

Output example

The following is a sample modified VRRP policy configuration output.

```
A:SR2>config>vrrp>policy# info  
-----  
    delta-in-use-limit 50  
    priority-event  
        port-down 1/1/2  
            hold-set 43200  
            priority 100 delta  
    exit  
    port-down 1/1/3  
        priority 200 explicit  
    exit  
    host-unreachable 10.10.24.4  
        drop-count 25  
    exit  
exit  
-----  
A:SR2>config>vrrp>policy#
```

3.11.2 Deleting a VRRP policy

Policies are only applied to non-owner VRRP instances. A VRRP policy cannot be deleted if it is applied to an interface or to an IES service. Each instance in which the policy is applied must be deleted.

The Applied column in the following example displays whether or not the VRRP policies are applied to an entity.

Output example

```
A:SR2#
=====
VRRP Policies
=====
Policy      Current      Current      Current      Delta      Applied
Id          Priority & Effect  Explicit    Delta Sum    Limit
-----
1  200 Explicit      200          100          50          Yes
15 254          None          None          1           No
32 100          None          None          1           No
=====
A:SR2#
```

3.11.3 Modifying service and interface VRRP parameters

3.11.3.1 Modifying non-owner parameters

When a VRRP instance is created as non-owner, it cannot be modified to the owner state. The VRID must be deleted and then recreated with the **owner** keyword to invoke IP address ownership.

3.11.3.2 Modifying owner parameters

When a VRRP instance is created as owner, it cannot be modified to the non-owner state. The VRID must be deleted and then recreated without the **owner** keyword to remove IP address ownership.

Entering the **owner** keyword is optional when entering the VRID for modification purposes.

3.11.3.3 Deleting VRRP on an interface or service

The VRID does not need to be shut down to remove the virtual router instance from an interface or service.

Example

```
config>router#interface
config>router# interface if-test
config>router>if# shutdown
config>router>if# exit
config>router# no interface if-test
config>router#
```

Example

The following shows the command usage to delete a VRRP instance from an interface or IES service.

```
config>service#ies 10
config>service>ies# interface "test"
config>service>ies>if# vrrp 1
config>service>ies>if>vrrp# shutdown
config>service>ies>if>vrrp# exit
config>service>ies>if# no vrrp 1
config>service>ies>if# exit all
```

3.12 VRRP command reference



Note:

VRRP commands are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

3.12.1 Command hierarchies

- [Configuration commands](#)
 - [VRRP network interface commands](#)
 - [VRRP IPv6 interface commands](#)
 - [VRRP priority control event policy commands](#)
- [Show commands](#)
- [Monitor commands](#)
- [Clear commands](#)
- [Debug commands](#)

3.12.1.1 Configuration commands

3.12.1.1.1 VRRP network interface commands

```
config
- router
  - [no] interface interface-name
    - address {ip-address/mask | ip-address netmask} [broadcast all-ones | host-ones]
    - no address
    - arp-timeout seconds
    - no arp-timeout
    - description description-string
    - no description
    - [no] shutdown
    - static-arp ip-address ieee-address
    - [no] static-arp ip-address
    - vrrp virtual-router-id [owner]
```

```

- no vrrp virtual-router-id
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- [no] backup ip-address
- [no] bfd-enable service-id interface interface-name dst-ip ip-address
- [no] bfd-enable interface interface-name dst-ip ip-address
- init-delay seconds
- no init-delay
- [no] master-int-inherit
- message-interval {[seconds] [milliseconds milliseconds]}
- no message-interval
- [no] ping-reply
- policy policy-id
- no policy
- [no] preempt
- priority priority
- no priority
- [no] ssh-reply
- [no] standby-forwarding
- [no] telnet-reply
- [no] shutdown
- [no] traceroute-reply

```

3.12.1.1.2 VRRP IPv6 interface commands



Note:

VRRP IPv6 interface commands are only supported on 7210 SAS-Mxp.

```

config
- router
- [no] interface interface-name
- [no] ipv6
- vrrp virtual-router-id [owner]
- no vrrp virtual-router-id
- [no] backup ip-address
- init-delay seconds
- no init-delay
- [no] master-int-inherit
- message-interval {[seconds] [milliseconds milliseconds]}
- no message-interval
- [no] ping-reply
- policy policy-id
- no policy
- [no] preempt
- priority priority
- no priority
- [no] shutdown
- [no] standby-forwarding
- [no] telnet-reply
- [no] traceroute-reply

```

3.12.1.1.3 VRRP priority control event policy commands

```

config
- vrrp
- [no] policy policy-id [context service-id]
- delta-in-use-limit limit

```

```
- no delta-in-use-limit
- description description string
- no description
- [no] priority-event
  - [no] host-unreachable ip-address
    - drop-count consecutive-failures
    - no drop-count
    - hold-clear seconds
    - no hold-clear
    - hold-set seconds
    - no hold-set
    - interval seconds
    - no interval
    - priority priority-level [{delta | explicit}]
    - no priority
    - timeout seconds
    - no timeout
  - [no] lag-port-down lag-id
    - hold-clear seconds
    - no hold-clear
    - hold-set seconds
    - no hold-set
    - [no] number-down number-of-lag-ports-down
      - priority priority-level [delta | explicit]
      - no priority
  - [no] port-down port-id
    - hold-clear seconds
    - no hold-clear
    - hold-set seconds
    - no hold-set
    - priority priority-level [delta | explicit]
    - no priority
  - [no] route-unknown ip-prefix/mask
    - hold-clear seconds
    - no hold-clear
    - hold-set seconds
    - no hold-set
    - less-specific [allow-default]
    - no less-specific
    - [no] next-hop ip-address
    - priority priority-level [delta | explicit]
    - no priority
    - protocol protocol
    - no protocol[protocol]
    - [no] protocol ospf
    - [no] protocol isis
    - [no] protocol static
```

3.12.1.2 Show commands

```
show
- vrrp
  - policy [policy-id [event event-type specific-qualifier]]
- router
  - vrrp
    - instance
    - instance [interface interface-name [vrid virtual-router-id]]
    - statistics
```

3.12.1.3 Monitor commands

```
monitor
- router
  - vrrp
    - instance interface interface-name vr-id virtual-router-id
    - instance [interval seconds] [repeat repeat] [absolute | rate]
```

3.12.1.4 Clear commands

```
clear
- vrrp
  - statistics
- router
  - vrrp
    - interface ip-int-name [vrid virtual-router-id]
    - statistics interface interface-name [vrid virtual-router-id]
    - statistics
```

3.12.1.5 Debug commands

```
debug
- router
  - vrrp
    - events
    - events interface ip-int-name [vrid virtual-router-id]
    - no events
    - no events interface ip-int-name [vrid virtual-router-id]
    - packets
    - packets interface ip-int-name [vrid virtual-router-id]
    - packets
    - no packets
    - no packets interface ip-int-name [vrid virtual-router-id]
    - no packets
```

3.12.2 Command descriptions

- [Configuration commands](#)
- [Show commands](#)
- [Monitor commands](#)
- [Clear commands](#)
- [Debug commands](#)

3.12.2.1 Configuration commands

- [Interface configuration commands](#)
- [Priority policy commands](#)

- [Priority policy event commands](#)
- [Priority policy port down event commands](#)
- [Priority policy LAG events commands](#)
- [Priority policy host unreachable event commands](#)
- [Priority policy route unknown event commands](#)

3.12.2.1.1 Interface configuration commands

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

config>router>if>vrrp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures 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 any time.

To change the current in-use password key on multiple virtual router instances:

1. Identify the current master.
2. Shutdown the virtual router instance on all backups.
3. Execute the **authentication-key** command on the master to change the password key.
4. Execute the **authentication-key** command and **no shutdown** command on each backup.

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

Default

no authentication-key

Parameters

authentication-key

Specifies 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 (#, \$, spaces, and so on), 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 22 (hash-key1) or 121 (hash-key2) 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 that 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, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Specifies that the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

backup

Syntax

[no] backup *ip-address*

Context

config>router>if>vrrp

config>router>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document

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 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 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 proper list is important. The specified *ip-address* must be equal to the existing parental IP interface IP addresses (primary) or the **backup** command will fail.

For non-owner virtual router instances, the **backup** command creates an IP interface IP address used for routing IP packets and communicating with the system when the access commands are defined (**ping-reply**, **telnet-reply**, and **ssh-reply**). The specified *ip-address* must be an IP address of the parental IP interface local subnets created with the address. 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 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 *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 will not respond to ARP requests to *ip-address*, nor will it route packets received with its *vid* derived source MAC address. A non-master virtual router instance always silently discards packets destined for *ip-address*. A single virtual router instance may only have a single virtual router IP address from a specific 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.

When operating as a (non-owner) master, the default functionality associated with *ip-address* is ARP response to ARP requests to *ip-address*, routing of packets destined for the virtual router instance source MAC address, and silently discarding packets destined for *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 this 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-address* 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

Special Cases

Assigning the Virtual Router ID IP Address

When the *vid* is created on the parent IP interface, IP addresses need to be assigned to the virtual router instance. If the *vid* was created with the keyword **owner**, the virtual router instance IP addresses must have the parent IP interface defined IP addresses (primary). For non-owner virtual router instances, the virtual router IP addresses each must be within one of the parental IP interface IP address defined local subnets. For both **owner** and **non-owner** virtual router instances, the virtual router IP addresses must be explicitly defined using the **backup ip-address** command.

Virtual Router Instance IP Address Assignment Conditions

The RFC does not specify that the assigned IP addresses to the virtual router instance must be in the same subnet as the parent IP interface primary IP address. The only

requirement is that all virtual routers participating in the same virtual router instance have the same virtual router IP addresses assigned. To avoid confusion, the assigned virtual router IP addresses must be in a local subnet of one of the parent IP interfaces IP addresses. For **owner** virtual router instances the assigned virtual router IP address must be the same as the parental IP interface primary.

The following rules apply when adding, changing, or removing parental and virtual router IP addresses:

Owner Virtual Router IP Address Parental Association

When an IP address is assigned to an **owner** virtual router instance, it must be associated with one of the parental IP interface-assigned IP addresses. The virtual router IP address must be equal to the primary oIP address within the parental IP interface.

Example - Owner Virtual Router Instance

Parent IP addresses:	10.10.10.10/24	
Virtual router IP addresses:	10.10.10.11	Invalid (not equal to parent IP address)
	10.10.10.10	Associated (same as parent IP address 10.10.10.10)
	10.10.11.11	Invalid (not equal to parent IP address)

Non-Owner Virtual Router IP Address Parental Association

When an IP address is assigned to a non-owner virtual router instance, it must be associated with one of the parental IP interface assigned IP addresses. The virtual router IP address must be a valid IP address within one of the parental IP interfaces local subnet. Local subnets are created by the primary IP address in conjunction with the IP addresses mask. If the defined virtual router IP address is equal to the associated subnet broadcast address, it is invalid. Virtual router IP addresses for non-owner virtual router instances that are equal to a parental IP interface IP address are also invalid.

The same virtual router IP address may not be assigned to two separate virtual router instances. If the virtual router IP address already exists on another virtual router instance, the virtual router IP address assignment will fail.

Example - Non-Owner Virtual Router Instance

Parent IP addresses:	10.10.10.10/24	
Virtual router IP addresses:	10.10.10.11	Associated with 10.10.10.10 (in subnet)
	10.10.10.10	Invalid (same as parent IP address)
	10.10.11.11	Invalid (outside of all Parent IP subnets)

Virtual Router IP Address Assignment without Parent IP Address

When assigning an IP address to a virtual router instance, an associated IP address (see **Owner Virtual Router IP Address Parental Association** and **Non-Owner Virtual Router IP Address Parental Association**) on the parental IP interface must already exist. If an associated IP address on the parental IP interface is not configured, the virtual router IP address assignment fails.

Parent Primary IP Address Changed

When a virtual router IP address is set and the associated parent IP interface IP address is changed, the new parent IP interface IP address is evaluated to ensure it meets the association rules defined in [backup Owner Virtual Router IP Address Parental Association](#) or [Non-Owner Virtual Router IP Address Parental Association](#). If the association check fails, the parental IP address change is not allowed. If the parental IP address change fails, the previously configured IP address definition remains in effect.

Only the primary parent IP address can be changed. **Parent Primary IP Address Removal** describes IP address removal conditions.

Parent Primary IP Address Removal

When a virtual router IP address is successfully set, but removing the associated parent IP interface IP address is attempted and fails. All virtual router IP addresses associated with the parental IP interface IP address must be deleted before removing the parental IP address. This includes virtual router IP address associations from multiple virtual router instances on the IP interface.

Parameters

ip-address

Specifies the virtual router IP address, 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 the primary IP address for **owner** virtual router instances.

Values 1.0.0.1 to 223.255.255.254

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*

Context

config>router>if>vrrp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This commands assigns a bidirectional forwarding (BFD) session providing heart-beat mechanism for the specific VRRP instance. There can be only one BFD session assigned to any specific VRRP instance, but there can be multiple VRRP sessions using the same BFD session.

By enabling BFD on a specific 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. 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

service-id

Specifies the service ID of the interface running BFD.

Values

<i>service-id</i> :	1 to 2147483647
<i>svc-name</i> :	64 characters maximum

interface *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.

dst-ip *ip-address*

Specifies the destination address to be used for the BFD session.

init-delay

Syntax

init-delay *seconds*

no init-delay

Context

config>router>if>vrrp

config>router>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a VRRP initialization delay timer.

Parameters

seconds

Specifies the initialization delay timer for VRRP, in seconds.

Values 1 to 65535

master-int-inherit

Syntax

[no] **master-int-inherit**

Context

config>router>if>vrrp

config>router>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the virtual router instance to inherit the master VRRP router 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 VRRP advertisement message advertisement interval field value or the message is discarded.

The **no** form of this command reverts to the default operating condition which requires the locally configured **message-interval** to match the received VRRP advertisement message advertisement interval field value.

Default

no master-int-inherit

message-interval

Syntax

message-interval {[seconds] [milliseconds milliseconds]}

no message-interval

Context

config>router>if>vrrp

config>router>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document

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 advertisement messages that are fragmented contain IP options (IPv4) 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:

$(3 \times (\text{in-use message interval}) + \text{skew time})$

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.

In 7210, the least timer values supported is 1 second. Timers less than 1 second cannot be used.

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

Default

1 second

Parameters

seconds

Specifies the number of seconds that will transpire before the advertisement timer expires expressed as a decimal integer.

Values IPv4: 1 to 255

milliseconds *milliseconds*

Specifies the time interval, in milliseconds, between sending advertisement messages.

Values 100 to 900



Note:

The *milliseconds* parameter is only supported on 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T, and 7210 SAS-Mxp.

policy

Syntax

policy *policy-id*

no policy

Context

config>router>if>vrrp

config>router>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document

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 this command removes existing VRRP priority control policy associations from the virtual router instance. All associations must be removed before deleting the policy from the system.

Default

no policy

Parameters

policy-id

Specifies 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

preempt

Syntax

[no] preempt

Context

config>router>if>vrrp

config>router>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command overrides an existing VRRP master if the virtual router in-use priority is higher than the current master.

The priority of the non-owner virtual router instance, the preempt mode allows the best available virtual router to force itself as the master over other available virtual routers.

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.

Enabling **preempt** mode improves the effectiveness of the base **priority** and the VRRP priority control policy mechanisms 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 diminished.

The **preempt** command is only available in the non-owner **vrrp** nodal context. The owner may not be preempted because the priority of non-owners can never be higher than the owner. The owner always preempts all other virtual routers when it is available.

Non-owner virtual router instances 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 only allows 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.

By default, preempt mode is enabled on the virtual router instance.

The **no** form of this command disables preempt mode and prevents the non-owner virtual router instance from preempting another, less desirable virtual router.

priority

Syntax

priority *base-priority*

no priority

Context

config>router>if>vrrp

config>router>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document

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 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 this command reverts to the default value.

Default

100

Parameters

base-priority

Specifies 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

ping-reply

Syntax

[no] ping-reply

Context

config>router>if>vrrp

config>router>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document

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 for 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 command 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 this command configures discarding all ICMP echo request messages destined for the non-owner virtual router instance IP addresses.

Default

no ping-reply

shutdown

Syntax

[no] shutdown

Context

```
config>router>if>vrrp  
config>router>if>ipv6>vrrp (7210 SAS-Mxp only)
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

Special Cases

Non-Owner Virtual Router

Non-owner virtual router instances can be administratively shutdown. This allows the termination of VRRP participation in the virtual router and stops all routing and other access capabilities with regards to the virtual router IP addresses. Shutting down the virtual router instance provides a mechanism to maintain the virtual routers without causing false backup/master state changes.

If the **shutdown** command is executed, no VRRP advertisement messages are generated and all received VRRP advertisement messages are silently discarded with no processing.

By default, virtual router instances are created in the **no shutdown** state.

Whenever the administrative state of a virtual router instance transitions, a log message is generated.

Whenever the operational state of a virtual router instance transitions, a log message is generated.

Owner Virtual Router

An owner virtual router context does not have a **shutdown** command. To administratively disable an owner virtual router instance, use the **shutdown** command within the parent IP interface node which administratively downs the IP interface.

VRRP Protocol Handling

On all 7210 SAS platforms, VRRP is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure>router>if>vrrp** command instantiates the protocol in the **no shutdown** state and resources are allocated to enable the node to process the protocol.

To deallocate resources, you must issue the **configure>router>if>vrrp>shutdown** and **configure>router>if>no vrrp** commands to allow the node to boot up correctly after the reboot. It is not sufficient to only issue a **configure>router>if>vrrp>shutdown** command.



Note:

The resources for VRRP are allocated when the VRRP context is enabled either in the base routing instance or the VPRN service instance. Resources

are deallocated when the configuration of the last VRRP context under either base routing instances or VPRN service is removed.

VRRPv3 Protocol Handling

On all 7210 SAS platforms, VRRPv3 is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure>router>if>ipv6>vrrp** command instantiates the protocol in the **no shutdown** state and resources are allocated to enable the node to process the protocol.

To deallocate resources, you must issue the **configure>router>if>ipv6>vrrp>shutdown** and **configure>router>if>ipv6>no vrrp** commands to allow the node to boot up correctly after the reboot. It is not sufficient to only issue a **configure>router>if>ipv6>vrrp>shutdown** command.



Note:

The resources for VRRPv3 are allocated when the VRRPv3 context is enabled either in the base routing instance, or in the VPRN service instance. Resources are deallocated when the configuration of the last VRRPv3 context, under either base routing instances or VPRN service, is removed.

ssh-reply

Syntax

[no] ssh-reply

Context

config>router>if>vrrp

Platforms

Supported on all 7210 SAS platforms as described in this document

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). Proper 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 this command discards all SSH request messages destined for the non-owner virtual router instance IP addresses.

Default

no ssh-reply

standby-forwarding

Syntax

[no] **standby-forwarding**

Context

config>router>if>vrrp

config>router>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document

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 real MAC address. When enabled, a standby router should forward all traffic.

telnet-reply

Syntax

[no] **telnet-reply**

Context

config>router>if>vrrp

config>router>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document

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 for 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). Proper 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.

The **no** form of this command configures discarding all Telnet request messages destined to the non-owner virtual router instance IP addresses.

Default

no telnet-reply

traceroute-reply

Syntax

[no] **traceroute-reply**

Context

config>router>if>vrrp

config>router>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document

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

vrrp

Syntax

vrrp *vrid* [**owner**]

no vrrp *vrid*

Context

config>router>interface

config>router>if>ipv6 (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures 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**. When 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. For IPv6, only one **vrrp** *vrid* node can be configured on a router interface. Each virtual router instance can manage up to 16 backup IP addresses.

The **no** form of this 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 shutdown to remove the virtual router instance.

Default

no vrrp

Special Cases

Virtual Router Instance Owner IP Address Conditions

It is possible for the virtual router instance **owner** to be created before assigning the parent IP interface primary IP address. When this is the case, the virtual router instance is not associated with an IP address. The operational state of the virtual router instance is down.

VRRP Owner Command Exclusions

By specifying the VRRP *vrid* as **owner**, The following commands are no longer available:

- **vrrp priority** — The virtual router instance **owner** is hard-coded with a **priority** value of 255 and cannot be changed.
- **vrrp master-int-inherit** — Owner virtual router instances do not accept VRRP advertisement messages; the advertisement interval field is not evaluated and cannot be inherited.
- **ping-reply**, **telnet-reply** and **ssh-reply** — The **owner** virtual router instance always allows Ping, Telnet and SSH if the management and security parameters are configured to accept them on the parent IP interface.
- **vrrp shutdown** — The **owner** virtual router instance cannot be shutdown in the **vrrp** node. If this was allowed, VRRP messages would not be sent, but the parent IP interface address would continue to respond to ARPs and forward IP packets. Another virtual router instance may detect the missing master because of the termination of VRRP advertisement messages and become master. This would cause two routers responding to ARP requests for the same IP addresses.

To **shutdown** the **owner** virtual router instance, use the **shutdown** command in the parent IP interface context. This will prevent VRRP participation, IP ARP reply and IP forwarding. To continue parent IP interface ARP reply and forwarding without VRRP participation, remove the **vrrpvrid** instance.

- **traceroute-reply**

Parameters

vrid

Specifies the virtual router ID for the IP interface, expressed as a decimal integer.

Values 1 to 255

owner

Specifies 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.

3.12.2.1.2 Priority policy commands

delta-in-use-limit

Syntax

delta-in-use-limit *in-use-priority-limit*

no delta-in-use-limit

Context

config>vrrp>policy

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures 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.

When 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 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 this command reverts to the default value.

Default

1

Parameters

in-use-priority-limit

Specifies the lower limit of the in-use priority base, as modified by priority control policies. The 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

description

Syntax

description *string*

no description

Context

config>vrrp>policy

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a text description for 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

string

Specifies the description character string. Allowed values are any string up to 80 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.

policy

Syntax

policy *policy-id* [**context** *service-id*]

no policy *policy-id*

Context

config>vrrp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures 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 this 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

vrrp-policy-id

Specifies 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 service-id

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

priority-event

Syntax

[no] **priority-event**

Context

config>vrrp>policy

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures 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 this command clears any configured priority events.

3.12.2.1.3 Priority policy event commands

hold-clear

Syntax

hold-clear *seconds*

no hold-clear

Context

config>vrrp>policy>priority-event>port-down

config>vrrp>policy>priority-event>lag-port-down

config>vrrp>policy>priority-event>route-unknown

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the hold clear time for the event.

The hold-clear time is used to prevent blackhole 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

hold-set

Syntax

hold-set *seconds*

no hold-set

Context

```
config>vrrp>policy>priority-event>host-unreachable
config>vrrp>policy>priority-event>lag-port-down
config>vrrp>policy>priority-event>port-down
config>vrrp>policy>priority-event>route-unknown
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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. 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.

When 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 anytime. 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 this command reverts the default value.

Default

0

Parameters

seconds

Specifies 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

priority

Syntax

```
priority priority-level [{delta | explicit}]
```

no priority

Context

```
config>vrrp>policy>priority-event>host-unreachable
```

```
config>vrrp>policy>priority-event>lag-port-down>number-down
```

```
config>vrrp>policy>priority-event>port-down
```

```
config>vrrp>policy>priority-event>route-unknown
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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 this command reverts to the default values.

Default

0

Parameters

priority-level

Specifies the priority level adjustment value, expressed as a decimal integer.

Values 0 to 254

delta | **explicit**

Specifies what effect the *priority-level* will have on the base priority value.

When **delta** is specified, the *priority-level* value is subtracted from the associated virtual router instance 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.

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.

Values delta, explicit

Default delta

3.12.2.1.4 Priority policy port down event commands

port-down

Syntax

[no] port-down port-id

Context

config>vrrp>policy>priority-event

Platforms

Supported on all 7210 SAS platforms as described in this document

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, 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 will indicate:

- 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. When 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 this 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

Parameters

port-id

Specifies 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 port-id *slot/mda/port[.channel]*

Specifies the POS channel on the port monitored by the VRRP priority control event. The *port-id.channel-id* can only be monitored by a single event in this policy. The channel 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.

If the port is provisioned, but the *channel* does not exist or the port has not been populated, the appropriate event operational state is Set – non-populated.

If the port is not provisioned, the event operational state is Set – non-provisioned.

If the POS interface is configured as a clear-channel, the *channel-id* is 1 and the channel bandwidth is the full bandwidth of the port.

3.12.2.1.5 Priority policy LAG events commands

lag-port-down

Syntax

```
[no] lag-port-down lag-id
```

Context

```
config>vrrp>policy>priority-event
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures 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
- 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 enter 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 than 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 this 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

Parameters

lag-id

Specifies 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 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.

number-down

Syntax

[no] **number-down** *number-of-lag-ports-down*

Context

config>vrrp>policy>priority-event>lag-port-down

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures 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 specific threshold, that is the active threshold.

The **no** form of this 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

Parameters

number-of-lag-ports-down

Specifies 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 4

3.12.2.1.6 Priority policy host unreachable event commands

drop-count

Syntax

drop-count *consecutive-failures*

no drop-count

Context

config>vrrp>priority-event>host-unreachable

Platforms

Supported on all 7210 SAS platforms as described in this document

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 consecutive message drop counter. With each successful attempt, the event consecutive message drop counter resets to zero.

If the event consecutive message drop counter reaches the **drop-count** value, the **host-unreachable** priority event enters the set state.

The event **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 this command reverts to the default value.

Default

3

Parameters

consecutive-failures

Specifies 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

host-unreachable

Syntax

[no] host-unreachable *ip-address*

Context

config>vrrp>policy>priority-event

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures 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 following table lists the possible operational states of the **host-unreachable** event.

Table 29: host-unreachable operational states

Host unreachable operational state	Description
Set – no ARP	No ARP address found for <i>ip-address</i> for drop-count consecutive attempts. Only applies when IP address is considered local.
Set – no route	No route exists for <i>ip-address</i> for drop-count consecutive attempts. Only applies 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 <i>ip-address</i> - not enough failed attempts to set the event.
Cleared – no route	No route exists for <i>ip-address</i> - 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 result of the last attempt. 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 must be expired and the historical success rate must be met before the event operational state becoming cleared.

The **no** form of this command deletes the specific IP host monitoring event. The event may be deleted at anytime. When the event is deleted, the in-use priority of all associated virtual router instances must be reevaluated. The event **hold-set** timer has no effect on the removal procedure.

Default

no host-unreachable

Parameters

ip-address

Specifies the IP address of the host for which the specific event will monitor connectivity. The *ip-address* 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[-*interface*]

x - 0 to FFFF (hexadecimal)

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.

interval

Syntax

interval seconds

no interval

Context

config>vrrp>priority-event>host-unreachable

Platforms

Supported on all 7210 SAS platforms as described in this document

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 this command reverts to the default value.

Default

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

timeout

Syntax

timeout *seconds*

no timeout

Context

config>vrrp>priority-event>host-unreachable

Platforms

Supported on all 7210 SAS platforms as described in this document

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 one of the following occurs.

- 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).



Note:

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 before the **timeout** period for a specific 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 before 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 this command reverts to the default value.

Default

1

Parameters

seconds

Specifies the number of seconds before an ICMP echo request message is timed out. When a message is timed out, a reply with the same identifier and sequence number is discarded.

Values 1 to 60

3.12.2.1.7 Priority policy route unknown event commands

less-specific

Syntax

[no] **less-specific** [allow-default]

Context

config>vrrp>policy>priority-event>route-unknown

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables 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 this command prevents RTM lookup results that are less specific than the route prefix from matching.

Default

no less-specific

Parameters

allow-default

Specifies that 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).

next-hop

Syntax

[no] next-hop *ip-address*

Context

config>vrrp>policy>priority-event>route-unknown

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds an enabled 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 this 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

Parameters

ip-address

Specifies 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

protocol

Syntax

protocol {ospf | is-is | static}

no protocol

Context

config>vrrp>policy>priority-event>route-unknown

Platforms

Supported on all 7210 SAS platforms as described in this document

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 this 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

Parameters

ospf

Specifies 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

Specifies 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.

static

Specifies 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.

route-unknown

Syntax

```
[no] route-unknown prefix/mask-length
```

Context

```
config>vrrp>policy>priority-event
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables 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 proper 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 described in the following table.

Table 30: 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 that matches all the preceding criteria (if defined in the event control commands) is not found, 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 this 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

Parameters

prefix

Specifies the IP prefix address to be monitored by the route unknown priority control event, in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255

mask-length

Specifies 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 0 to 32

ip-address

Specifies the IP address of the host for which the specific event will monitor connectivity. The *ip-address* 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-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	<i>ip-prefix/mask:</i>	ipv4-prefix -	a.b.c.d (host bits must be 0)
		mask-length -	0 to 32
Values	<i>ipv6-address/ prefix:</i>	ipv6-address -	x:x:x:x:x:x (eight 16-bit pieces)
			x:x:x:x:x:d.d.d.d
			x - 0 to FFFF (hexadecimal)
		prefix-length -	1 to 128

3.12.2.2 Show commands

instance

Syntax

instance interface *interface-name* [**vrid** *virtual-router-id*]

Context

show>vrrp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information for VRRP instances.

If no command line options are specified, summary information for all VRRP instances displays.

Parameters

interface *interface-name*

Displays detailed information for the VRRP instances on the specified IP interface including status and statistics.

vrid *virtual-router-id*

Displays detailed information for the specified VRRP instance on the IP interface.

Values 1 to 255

Output

The following output is an example of VRRP instance information, and [Table 31: Output fields: VRRP instance](#) describes the output fields.

Sample output

```
*A:ALA-A# show router vrrp instance
=====
VRRP Instances
=====
Interface Name          VR Id Own Adm State      Base Pri  Msg Int
                       IP      Opr  Pol Id   InUse Pri  Inh Int
-----
n2                      1    No  Up  Master   100      1
                       IPv4    Up  n/a     100      No
Backup Addr: 5.1.1.10

-----
Instances : 2
=====
*A:ALA-A#

*A:ALA-A# show router vrrp instance interface n2 vrid 1
=====
VRRP Instance 1 for interface "n2"
=====
Owner                   : No                VRRP State           : Master
Primary IP of Master: 5.1.1.2 (Self)
Primary IP              : 5.1.1.2          Standby-Forwarding: Disabled
VRRP Backup Addr       : 5.1.1.10
Admin State             : Up                Oper State            : Up
Up Time                 : 09/23/2004 06:53:45 Virt MAC Addr        : 00:00:5e:00:01:01
Auth Type               : None
Config Mesg Intvl      : 1                In-Use Mesg Intvl    : 1
Master Inherit Intvl   : No
```

```

Base Priority      : 100          In-Use Priority   : 100
Policy ID         : n/a          Preempt Mode     : Yes
Ping Reply        : No           Telnet Reply     : No
SSH Reply         : No           Traceroute Reply: No
Init Delay        : 0            Init Timer Expires: 0.000 sec
Creation State    : Active
    
```

 Master Information

```

Primary IP of Master: 5.1.1.2 (Self)
Addr List Mismatch  : No          Master Priority   : 100
Master Since        : 09/23/2004 06:53:49
    
```

 Masters Seen (Last 32)

Primary IP of Master	Last Seen	Addr List Mismatch	Msg Count
5.1.1.2	09/23/2004 06:53:49	No	0

 Statistics

```

Become Master      : 1            Master Changes   : 1
Adv Sent           : 103          Adv Received     : 0
Pri Zero Pkts Sent: 0            Pri Zero Pkts Rcvd: 0
Preempt Events     : 0            Preempted Events: 0
Mesg Intvl Discards: 0          Mesg Intvl Errors: 0
Addr List Discards : 0            Addr List Errors : 0
Auth Type Mismatch: 0            Auth Failures    : 0
Invalid Auth Type  : 0            Invalid Pkt Type  : 0
IP TTL Errors      : 0            Pkt Length Errors: 0
Total Discards     : 0
    
```

=====
 *A:ALA-A#

7210SAS>show>router# vrrp instance interface "n1" vrid 1

=====
 VRRP Instance 1 for interface "n1"
 =====

```

Owner              : No          VRRP State       : Init
Primary IP of Master: 0.0.0.0 (Self)
Primary IP         : 0.0.0.0     Standby-Forwarding: Disabled
VRRP Backup Addr   : None
Admin State        : Up           Oper State        : Down
Up Time            : 02/16/2000 02:01:54 Virt MAC Addr     : 00:00:5e:00:01:01
Auth Type          : None
Config Mesg Intvl  : 1            In-Use Mesg Intvl : 1
Master Inherit Intvl: No
Base Priority       : 100          In-Use Priority   : 100
Policy ID          : n/a          Preempt Mode     : Yes
Ping Reply         : No           Telnet Reply     : No
SSH Reply          : No           Traceroute Reply: No
Init Delay         : 0            Init Timer Expires: 0.000 sec
Creation State     : Init
    
```

 Master Information

```

Primary IP of Master: 0.0.0.0 (Self)
Addr List Mismatch  : Unknown     Master Priority   : 0
Master Since        : 02/16/2000 02:01:54
    
```

 Masters Seen (Last 32)

```

-----
Primary IP of Master   Last Seen           Addr List Mismatch  Msg Count
-----
None
-----
Statistics
-----
Become Master        : 0                Master Changes      : 0
Adv Sent             : 0                Adv Received        : 0
Pri Zero Pkts Sent  : 0                Pri Zero Pkts Rcvd : 0
Preempt Events       : 0                Preempted Events    : 0
Mesg Intvl Discards : 0                Mesg Intvl Errors   : 0
Addr List Discards  : 0                Addr List Errors     : 0
Auth Type Mismatch  : 0                Auth Failures       : 0
Invalid Auth Type   : 0                Invalid Pkt Type    : 0
IP TTL Errors        : 0                Pkt Length Errors   : 0
Total Discards       : 0
=====
7210SAS>show>router#
    
```

Table 31: Output fields: VRRP instance

Label	Description
Interface name	The name of the IP interface
VR ID	The virtual router ID for the IP interface
Own Owner	Yes — Specifies that the virtual router instance as owning the virtual router IP addresses
	No — Indicates that the virtual router instance is operating as a non-owner
Adm	Up — Indicates that the administrative state of the VRRP instance is up
	Down — Indicates that the administrative state of the VRRP instance is down
Opr	Up — Indicates that the operational state of the VRRP instance is up
	Down — Indicates that the operational state of the VRRP instance is down
State	When owner, backup defines the IP addresses that are advertised within VRRP advertisement messages. When non-owner, backup actually creates an IP interface IP address used for routing IP packets and communicating with the system when the access commands are defined (ping-reply, telnet-reply, and ssh-reply).
Pol ID	The value that uniquely identifies a Priority Control Policy

Label	Description
Base Priority	The <i>base-priority</i> value is used to derive the in-use priority of the virtual router instance as modified by any optional VRRP priority control policy.
InUse Priority	The current in-use priority associated with the VRRP virtual router instance
Msg Int	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
Inh Int	Yes — When the VRRP instance is a non-owner and is operating as a backup and the master-int-inherit command is enabled, the master down timer is indirectly derived from the value in the advertisement interval field of the VRRP message received from the current master.
	No — When the VRRP instance is operating as a backup and the master-int-inherit command is not enabled, the configured advertisement interval is matched against the value in the advertisement interval field of the VRRP message received from the current master. If the two values do not match, then the VRRP advertisement is discarded. If the VRRP instance is operating as a master, this value has no effect.
Backup Addr	The backup virtual router IP address
BFD	Indicates BFD is enabled.
VRRP State	Specifies whether the VRRP instance is operating in a master or backup state
Policy ID	The VRRP priority control policy associated with the VRRP virtual router instance A value of 0 indicates that no control policy is associated with the virtual router instance
Preempt Mode	Yes — The preempt mode is enabled on the virtual router instance where it will preempt a VRRP master with a lower priority
	No — The preempt mode is disabled and prevents the non-owner virtual router instance from preempting another, less desirable virtual router
Ping Reply	Yes — A non-owner master is enabled to reply to ICMP Echo requests directed to the virtual router instance IP addresses Ping Reply is valid only if the VRRP virtual router instance associated with this entry is a non-owner

Label	Description
	A non-owner backup virtual router never responds to such ICMP echo requests irrespective if Ping Reply is enabled
	No — ICMP echo requests to the virtual router instance IP addresses are discarded
Telnet Reply	Yes — Non-owner masters can to reply to TCP port 23 Telnet requests directed at the virtual router instances IP addresses
	No — Telnet requests to the virtual router instance IP addresses are discarded
SSH Reply	Yes — Non-owner masters can to reply to SSH requests directed at the virtual router instances IP addresses
	No — All SSH request messages destined for the non-owner virtual router instance IP addresses are discarded
Primary IP of Master	The IP address of the VRRP master
Primary IP	The IP address of the VRRP owner
Up Time	The date and time when the operational state of the event last changed
Virt MAC Addr	The virtual MAC address used in ARP responses when the VRRP virtual router instance is operating as a master
Auth Type	Specifies the VRRP authentication Type 0 (no authentication), Type 1 (simple password), or Type 2 (MD5) for the virtual router
Addr List Mismatch	Specifies whether a trap was generated when the IP address list received in the advertisement messages received from the current master did not match the configured IP address list. This is an edge triggered notification. A second trap will not be generated for a packet from the same master until this event has been cleared.
Master Priority	The priority of the virtual router instance which is the current master
Master Since	The date and time when operational state of the virtual router changed to master For a backup virtual router, this value specifies the date and time when it received the first VRRP advertisement message from the virtual router which is the current master.

policy

Syntax

policy [*vrpp-policy-id* [**event** *event-type specific-qualifier*]]

Context

show>vrpp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays VRRP priority control policy information.

If no command line options are specified, a summary of the VRRP priority control event policies displays.

Parameters

vrpp-policy-id

Displays information about the specified priority control policy ID.

Values 1 — 9999

Default all VRRP policies IDs

event event-type

Displays information about the specified VRRP priority control event within the policy ID.

Values **port-down** *port-id*
lag-port-down *lag-id*
host-unreachable *host-ip-addr*
route-unknown *route-prefix/mask*

Default all event types and qualifiers

specific-qualifier

Display information about the specified qualifier.

Values port-id, lag-id, host-ip-addr, route-prefix/mask

Output

The following outputs are examples of VRRP policy information. The associated tables describe the output fields.

- [Sample output, Table 32: Output fields: VRRP policy](#)
- [Sample output for VRRP policy event, Table 33: Output fields: VRRP policy event](#)

Sample output

```

A:ALA-A# show vrrp policy
=====
VRRP Policies
=====
Policy      Current      Current      Current      Delta      Applied
Id          Priority & Effect  Explicit    Delta Sum   Limit
-----
1           None         None        None        1          Yes
2           None         None        None        1          No
=====

A:ALA-A#

A:ALA-A# show vrrp policy 1
=====
VRRP Policy 1
=====
Description      : 10.10.200.253 reachability
Current Priority: None          Applied           : No
Current Explicit: None          Current Delta Sum : None
Delta Limit      : 1

-----
Applied To      VR   Opr   Base   In-use  Master  Is
Interface Name  Id        Pri   Pri   Pri     Pri     Master
-----
None

-----
Priority Control Events
-----
Event Type & ID      Event Oper State      Hold Set  Priority In
Remaining &Effect  Use
-----
Host Unreach 10.10.200.252      n/a          Expired   20 Del  No
Host Unreach 10.10.200.253      n/a          Expired   10 Del  No
Route Unknown 10.10.100.0/24      n/a          Expired   1 Exp  No
=====

A:ALA-A#
    
```

Table 32: Output fields: VRRP policy

Label	Description
Policy Id	The VRRP priority control policy associated with the VRRP virtual router instance A value of 0 indicates that no control policy is associated with the virtual router instance
Current Priority & Effects	
Current Explicit	When multiple explicitly defined events associated with the priority control policy happen simultaneously, the lowest value of all the current explicit priorities will be used as the in-use priority for the virtual router.

Label	Description
Current Delta Sum	The sum of the priorities of all the delta events when multiple delta events associated with the priority control policy happen simultaneously. This sum is subtracted from the base priority of the virtual router to give the in-use priority.
Delta Limit	The delta-in-use-limit for a VRRP policy. When the total sum of all delta events has been calculated and subtracted from the base-priority of the virtual router, the result is compared to the delta-in-use-limit value. If the result is less than this value, 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. If the delta-in-use-limit is 0, the sum of the delta priority control events to reduce the virtual router's in-use-priority to 0 can prevent it from becoming or staying master.
Current Priority	The configured delta-in-use-limit priority for a VRRP priority control policy or the configured delta or explicit priority for a priority control event
Applied	The number of virtual router instances to which the policy has been applied The policy cannot be deleted unless this value is 0
Description	A text string which describes the VRRP policy
Event Type & ID	A delta priority event is a conditional event defined in a priority control policy that subtracts a specific amount from the base priority to give the current in-use priority for the VRRP virtual router instances to which the policy is applied.
	An explicit priority event is a conditional event defined in a priority control policy that explicitly defines the in-use priority for the VRRP virtual router instances to which the policy is applied. Explicit events override all delta Events. When multiple explicit events occur simultaneously, the event with the lowest priority value defines the in-use priority.
Event Oper State	The operational state of the event
Hold Set Remaining	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.
Priority & Effect	Delta — The <i>priority-level</i> value is subtracted from the associated virtual router instance base priority when the event is set and no explicit events are set. The sum of the priority event <i>priority-level</i> 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.

Label	Description
	If the delta priority event is cleared, the <i>priority-level</i> is no longer used in the in-use priority calculation.
	Explicit — The <i>priority-level</i> 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 <i>priority-level</i> . The set explicit priority value with the lowest <i>priority-level</i> determines the actual in-use protocol value for all virtual router instances associated with the policy.
In Use	Specifies whether the event is currently affecting the in-use priority of some virtual router

Sample output for VRRP policy event

```
A:ALA-A#show vrrp policy 1 event port-down
=====
VRRP Policy 1, Event Port Down 1/1/1
=====
Description      :
Current Priority: None           Applied           : Yes
Current Explicit: None         Current Delta Sum : None
Delta Limit      : 1

-----
Applied To      VR   Opr   Base   In-use   Master   Is
Interface Name  Id   Pri   Pri    Pri     Pri     Master
-----
ies301backup    1   Down  100    100     0       No

-----
Priority Control Event Port Down 1/1/1
-----
Priority          : 30           Priority Effect   : Delta
Hold Set Config  : 0 sec         Hold Set Remaining: Expired
Value In Use     : No           Current State    : Cleared
# trans to Set   : 6           Previous State   : Set-down
Last Transition  : 04/13/2007 04:54:35
=====
A:ALA-A#
A:ALA-A# show vrrp policy 1 event host-unreachable
=====
VRRP Policy 1, Event Host Unreachable 10.10.200.252
=====
Description      : 10.10.200.253 reachability
Current Priority: None           Applied           : No
Current Explicit: None         Current Delta Sum : None
Delta Limit      : 1

-----
Applied To      VR   Opr   Base   In-use   Master   Is
Interface Name  Id   Pri   Pri    Pri     Pri     Master
-----
None

-----
Priority Control Event Host Unreachable 10.10.200.252
```

```

-----
Priority      : 20                Priority Effect   : Delta
Interval     : 1 sec            Timeout          : 1 sec
Drop Count   : 3
Hold Set Config : 0 sec          Hold Set Remaining: Expired
Value In Use  : No              Current State    : n/a
# trans to Set : 0              Previous State   : n/a
Last Transition : 04/13/2007 23:10:24
=====
A:ALA-A#

A:ALA-A# show vrrp policy 1 event route-unknown
=====
VRRP Policy 1, Event Route Unknown 10.10.100.0/24
=====
Description      : 10.10.200.253 reachability
Current Priority: None                Applied          : No
Current Explicit: None              Current Delta Sum : None
Delta Limit      : 1
-----

Applied To      VR      Opr      Base      In-use      Master      Is
Interface Name  Id       Pri      Pri      Pri      Pri      Master
-----
None
-----

Priority Control Event Route Unknown 10.10.100.0/24
-----
Priority      : 1                Priority Effect   : Explicit
Less Specific : No              Default Allowed  : No
Next Hop(s)   : None
Protocol(s)   : None
Hold Set Config : 0 sec          Hold Set Remaining: Expired
Value In Use  : No              Current State    : n/a
# trans to Set : 0              Previous State   : n/a
Last Transition : 04/13/2007 23:10:24
=====
A:ALA-A#
    
```

Table 33: Output fields: VRRP policy event

Label	Description
Description	A text string which describes the VRRP policy
Policy Id	The VRRP priority control policy associated with the VRRP virtual router instance A value of 0 indicates that no control policy is associated with the virtual router instance
Current Priority	The base router priority for the virtual router instance used in the master election process
Current Explicit	When multiple explicitly defined events associated with the priority control policy happen simultaneously, the lowest value of all the current explicit priorities will be used as the in-use priority for the virtual router.

Label	Description
Applied	The number of virtual router instances to which the policy has been applied. The policy cannot be deleted unless this value is 0.
Current Delta Sum	The sum of the priorities of all the delta events when multiple delta events associated with the priority control policy happen simultaneously. This sum is subtracted from the base priority of the virtual router to give the in-use priority.
Delta Limit	<p>The delta-in-use-limit for a VRRP policy. When the total sum of all delta events has been calculated and subtracted from the base-priority of the virtual router, the result is compared to the delta-in-use-limit value. If the result is less than this value, 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.</p> <p>If the delta-in-use-limit is 0, the sum of the delta priority control events to reduce the virtual router's in-use-priority to 0 can prevent it from becoming or staying master.</p>
Applied to Interface Name	The interface name where the VRRP policy is applied
VR ID	The virtual router ID for the IP interface
Opr	Up — Indicates that the operational state of the VRRP instance is up
	Down — Indicates that the operational state of the VRRP instance is down
Base Pri	The base priority used by the virtual router instance
InUse Priority	The current in-use priority associated with the VRRP virtual router instance
Master Priority	The priority of the virtual router instance which is the current master
Priority	The base priority used by the virtual router instance
Priority Effect	Delta — A delta priority event is a conditional event defined in a priority control policy that subtracts a specific amount from the base priority to give the current in-use priority for the VRRP virtual router instances to which the policy is applied.
	Explicit — A conditional event defined in a priority control policy that explicitly defines the in-use priority for the VRRP virtual router instances to which the policy is applied.

Label	Description
	Explicit events override all delta events. When multiple explicit events occur simultaneously, the event with the lowest priority value defines the in-use priority.
Current Priority	The configured delta-in-use-limit priority for a VRRP priority control policy, or the configured delta, or explicit priority for a priority control event
Event Oper State	The operational state of the event
Hold Set Remaining	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
Priority	The base priority used by the virtual router instance
Priority Effect	Delta — The <i>priority-level</i> value is subtracted from the associated virtual router instance base priority when the event is set and no explicit events are set. The sum of the priority event <i>priority-level</i> 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 <i>priority-level</i> is no longer used in the in-use priority calculation.
	Explicit — The <i>priority-level</i> 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 <i>priority-level</i> . The set explicit priority value with the lowest <i>priority-level</i> determines the actual in-use protocol value for all virtual router instances associated with the policy.
Hold Set Config	The configured 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.
Value In Use	Yes — The event is currently affecting the in-use priority of some virtual router
	No — The event is not affecting the in-use priority of some virtual router
# trans to Set	The number of times the event has transitioned to one of the 'set' states
Last Transition	The time and date when the operational state of the event last changed

statistics

Syntax

statistics

Context

show>router>vrrp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays statistics for VRRP instance.

Output

The following output is an example of VRRP statistics information, and [Table 34: Output fields: VRRP statistics](#) describes the output fields.

Sample output

```
A:ALA-48# show router vrrp statistics
=====
VRRP Global Statistics
=====
VR Id Errors      : 0                Version Errors    : 0
Checksum Errors  : 0
=====
A:ALA-48#
```

Table 34: Output fields: VRRP statistics

Label	Description
VR Id Errors	Displays the number of virtual router ID errors
Version Errors	Displays the number of version errors
Checksum Errors	Displays the number of checksum errors

3.12.2.3 Monitor commands

instance

Syntax

instance interface interface-name vr-id virtual-router-id [interval seconds] [repeat repeat] [absolute | rate]

Context

monitor>router>vrrp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command monitors statistics for a VRRP instance.

Parameters

interface interface-name

Specifies the name of the existing IP interface on which VRRP is configured.

vr-id virtual-router-id

Specifies the virtual router ID for the existing IP interface, expressed as a decimal integer.

interval seconds

Specifies the interval for each display, in seconds.

Values 3 to 60

Default 5

repeat repeat

Specifies how many times the command is repeated.

Values 1 to 999

Default 10

absolute

Specifies that the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

rate

Specifies that the rate-per-second for each statistic is displayed instead of the delta.

Output

The following output is an example of VRRP instance information.

Sample output

```
*A:ALA-A# monitor router vrrp instance interface n2 vr-id 1
=====
Monitor statistics for VRRP Instance 1 on interface "n2"
=====
-----
At time t = 0 sec (Base Statistics)
-----
Become Master      : 1                Master Changes    : 1
Adv Sent           : 1439             Adv Received      : 0
Pri Zero Pkts Sent : 0                Pri Zero Pkts Rcvd: 0
Preempt Events     : 0                Preempted Events  : 0
```

```
Mesg Intvl Discards : 0      Mesg Intvl Errors : 0
Addr List Discards  : 0      Addr List Errors  : 0
Auth Type Mismatch  : 0      Auth Failures     : 0
Invalid Auth Type   : 0      Invalid Pkt Type  : 0
IP TTL Errors       : 0      Pkt Length Errors : 0
Total Discards      : 0
```

```
=====
*A:ALA-A#
```

3.12.2.4 Clear commands

interface

Syntax

```
interface ip-int-name [vrid virtual-router-id]
```

Context

```
clear>router>vrrp
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command resets VRRP protocol instances on an IP interface.

Parameters

ip-int-name

Specifies the IP interface to reset the VRRP protocol instances.

vrid *vrid*

Resets the VRRP protocol instance for the specified VRID on the IP interface.

Values 1 to 255

Default all VRIDs on the IP interface

statistics

Syntax

```
statistics [policy policy-id]
```

Context

```
clear>router>vrrp
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears and resets VRRP entities.

Parameters

policy *policy-id*

Clears statistics for the specified policy.

Values 1 to 9999

statistics

Syntax

statistics interface *interface-name* [**vrid** *virtual-router-id*]

statistics policy [*vrrp-policy-id*]

Context

clear>router>vrrp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears statistics for VRRP instances on an IP interface or VRRP priority control policies.

Parameters

interface *interface-name*

Clears the VRRP statistics for all VRRP instances on the specified IP interface.

vrid *virtual-router-id*

Clears the VRRP statistics for the specified VRRP instance on the IP interface.

Values 1 to 255

Default all VRRP instances on the IP interface

policy [*vrrp-policy-id*]

Clears VRRP statistics for all or the specified VRRP priority control policy.

Values 1 to 9999

Default all VRRP policies

3.12.2.5 Debug commands

events

Syntax

events

events interface *ip-int-name* [**vrid** *virtual-router-id*]

no events

no events interface *ip-int-name* [**vrid** *virtual-router-id*]

Context

debug>router>vrrp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for VRRP events.

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

Parameters

ip-int-name

Displays the specified interface name.

vrid virtual-router-id

Displays the specified VRID.

packets

Syntax

packets interface *ip-int-name* [**vrid** *virtual-router-id*]

packets

no packets interface *ip-int-name* [**vrid** *virtual-router-id*]

no packets

Context

debug>router>vrrp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for VRRP packets.

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

Parameters

ip-int-name

Displays the specified interface name.

vrid virtual-router-id

Displays the specified VRID.

4 Filter policies

This chapter provides information about filter policies and management.

4.1 Filter policy configuration overview

Filter policies, also referred to as Access Control Lists (ACLs), are templates applied to services or network IP interfaces to control network traffic into (ingress) or out of (egress) a service access port (SAP) or network IP interface based on IP and MAC matching criteria. Filters are applied to services to look at packets entering or leaving a SAP. Filters can be used on several interfaces. The same filter can be applied to ingress traffic, egress traffic, or both. Ingress filters affect only inbound traffic destined for the routing complex, and egress filters affect only outbound traffic sent from the routing complex.

Configuring an entity with a filter policy is optional. If an entity such as a service or network IP interface is not configured with filter policies, then all traffic is allowed on the ingress and egress interfaces. By default, there are no filters associated with services or interfaces. They must be explicitly created and associated. When you create a new filter, default values are provided although you must specify a unique filter ID value to each new filter policy as well as each new filter entry and associated actions. The filter entries specify the filter matching criteria and also an action to be taken upon a match.

Available ingress and egress CAM hardware resources can be allocated as per user needs for use with different filter criteria. By default, the system allocates resources to maintain backward compatibility with release 4.0. Users can modify the resource allocation based on their need to scale the number of entries or number of associations (that is, number of SAP/IP interfaces using a filter policy that defines particular match criteria). If no CAM resources are allocated to particular match criteria defined in a filter policy, then the association of that filter policy to a SAP will fail. This is true for both ingress and egress filter policy. Please read the configuration notes section below for more information.

Only one ingress IP or MAC filter policy and one egress IP or MAC filter policy can be applied to a Layer 2 SAP. Both IPv4 and IPv6 ingress and egress filter policy can be used simultaneously with a Layer 2 SAP. Only one ingress IP filter policy and one egress IP filter policy can be applied to a network IP interface. Both IPv4 and IPv6 ingress and egress filter policy can be used simultaneously with an IP interface (For example: network Port IP interface in network mode and IES IP interface in access-uplink mode) for which IPv6 addressing is supported. Network filter policies control the forwarding and dropping of packets based on IP match criteria. Note that non-IP packets are not hitting the IP filter policy, so the default action in the filter policy will not apply to these packets.

Note that non-IP packets are not hitting the IP filter policy, so the default action in the filter policy will not apply to these packets.

4.1.1 Service and network IP interface-based filtering

IP and MAC filter policies specify either a forward or a drop action for packets based on information specified in the match criteria.

Filter entry matching criteria can be as general or specific as you require, but all conditions in the entry must be met in order for the packet to be considered a match and the specified entry action performed.

The process stops when the first complete match is found and executes the action defined in the entry, either to drop or forward packets that match the criteria.

4.1.2 Filter policy entities

A filter policy compares the match criteria specified within a filter entry to packets coming through the system, in the order the entries are numbered in the policy. When a packet matches all the parameters specified in the entry, the system takes the specified action to either drop or forward the packet. If a packet does not match the entry parameters, the packet continues through the filter process and is compared to the next filter entry, and so on. If the packet does not match any of the entries, then system executes the default action specified in the filter policy. Each filter policy is assigned a unique filter ID. Each filter policy is defined with:

- scope
- default action
- description

Each filter entry contains:

- match criteria
- an action

4.1.2.1 Applying filter policies

The following tables describe support of filter policies on different 7210 platforms.

Table 35: Applying filter policies for 7210 SAS-T devices configured in network mode

Service	IP filter	IPv6 filter	MAC filter
Network port IP interface	Network port IP interface (ingress and egress)	Network port IP interface (ingress and egress)	Not available
Epipe	Epipe SAP (ingress and egress)	Epipe SAP (ingress and egress)	Epipe SAP (ingress and egress)
VPLS	VPLS SAP (ingress and egress)	VPLS SAP (ingress and egress)	VPLS SAP (ingress and egress)
IES	IES interface SAP (ingress and egress)	IES interface SAP (ingress and egress)	Not available
VPRN	VPRN interface SAP (ingress and egress)	VPRN interface SAP (ingress and egress)	Not available
PBB	Ingress and egress of Epipe I-SAP and I-VPLS I-SAP	Ingress and egress of Epipe I-SAP, and I-VPLS I-SAP	Ingress and egress of Epipe I-SAP, I-VPLS I-SAP and B-VPLS B-SAP

Service	IP filter	IPv6 filter	MAC filter
RVPLS (RVPLS SAPs) ⁹	VPLS access (ingress and egress) and network SAPs (ingress and egress)	Not available	Not available
RVPLS (RVPLS IES IP Interface) ⁹	Ingress override filters (ingress)	Not available	Not available

Table 36: Applying filter policies for 7210 SAS-T (access-uplink mode)

Service	IP filter	IPv6 filter	MAC filter
Epipe	Epipe access SAP (egress and ingress), Epipe access-uplink SAP (egress and ingress)	Epipe (egress and ingress), Epipe access-uplink SAP (egress and ingress)	Epipe (egress and ingress), Epipe access-uplink SAP (egress and ingress)
VPLS	VPLS access SAP (ingress and egress), VPLS access-uplink SAP (ingress and egress)	VPLS access SAP (ingress and egress), VPLS access-uplink SAP (ingress and egress)	VPLS access SAP (ingress and egress), VPLS access-uplink SAP (ingress and egress)
RVPLS (VPLS SAPs) ⁹	VPLS access (ingress and egress) and access-uplink SAPs (ingress and egress)	Not available	Not available
RVPLS (RVPLS IES IP Interface) ⁹	Ingress override filters (ingress)	Not available	Not available
IES	IES access SAP, IES access-uplink SAP	IES access-uplink SAP	Not available

Table 37: Applying filter policies for 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Service	IP filter	IPv6 filter	MAC filter
Network port IP interface	Network port IP interface (ingress and egress)	Network port IP interface (ingress and egress)	Not available
Epipe	Epipe SAP (ingress and egress)	Epipe SAP (ingress and egress)	Epipe SAP (ingress and egress)
VPLS	VPLS SAP (ingress and egress)	VPLS SAP (ingress and egress)	VPLS SAP (ingress and egress)
IES	IES interface SAP (ingress and egress)	IES interface SAP (ingress and egress)	Not available

⁹ See the "Routed VPLS" section in the 7210 SAS-Mxp, S, Sx, T Services Guide for more information.

Service	IP filter	IPv6 filter	MAC filter
VPRN	VPRN interface SAP (ingress and egress)	VPRN interface SAP (ingress and egress)	Not available
PBB	Not supported	Not supported	Not supported
RVPLS (VPLS SAPs) ⁹	VPLS access (ingress and egress) and access-uplink SAPs (ingress and egress)	Available only for 7210 SAS-Mxp Not available for 7210 SAS-R6 and 7210 SAS-R12	Available only for 7210 SAS-Mxp Not available for 7210 SAS-R6 and 7210 SAS-R12
RVPLS (RVPLS IES and VPRN IP interface) ⁹	Ingress override filters (ingress)	Available only for 7210 SAS-Mxp Not available for 7210 SAS-R6 and 7210 SAS-R12	Not available

Table 38: Applying filter policies for 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE devices

Service	IP filter	IPv6 filter	MAC filter
Network port IP interface	Network port IP interface (ingress and egress)	Network port IP interface (ingress and egress)	Not available
Epipe	Epipe SAP (ingress and egress)	Epipe SAP (ingress and egress)	Epipe SAP (ingress and egress)
VPLS	VPLS SAP (ingress and egress)	VPLS SAP (ingress and egress)	VPLS SAP (ingress and egress)
IES	IES interface SAP (ingress and egress)	IES interface SAP (ingress and egress)	Not available
VPRN	VPRN interface SAP (ingress and egress)	VPRN interface SAP (ingress and egress)	Not available
PBB	Not available	Not available	Not available
RVPLS (VPLS SAPs) ⁹	VPLS access (ingress and egress) and access-uplink SAPs (ingress and egress)	Not available	Not available
RVPLS (RVPLS IES and VPRN IP interface) ⁹	Ingress override filters (ingress)	Not available	Not available

4.1.2.2 ACL on range SAPs

The ACLs on VLAN range SAPs are supported only on ingress (for Epipe and VPLS services). The following tables list the support.

Table 39: ACLs support in Epipe services on 7210 SAS-T network and access-uplink modes variants when using range SAPs

Platforms/types of filters	7210 SAS-T (network mode)	7210 SAS-T (access-uplink mode)
Ingress IP or IPv6	Yes	Yes
Ingress MAC	Yes	Yes
Egress IP	No	No
Egress MAC	No	No

Table 40: ACLs support in Epipe services on 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp variants when using range SAPs

Platforms/types of filters	7210 SAS-R6 and 7210 SAS-R12	7210 SAS-Mxp	7210 SAS-Sx/S 1/10GE	7210 SAS-Sx 10/100GE
Ingress IP or IPv6	Yes	Yes	Yes	Yes
Ingress MAC	Yes	Yes	Yes	Yes
Egress IP	No	No	No	No
Egress MAC	No	No	No	No

Table 41: ACLs support in VPLS services on 7210 SAS-T network and access-uplink mode variants when using range SAPs

Platforms/types of filters	7210 SAS-T (access-uplink mode)	7210 SAS-T (network mode)
Ingress IP or IPv6	Yes	No
Ingress MAC	Yes	No
Egress IP	No	No
Egress MAC	No	No

Filter policies are applied to the following service entities:

- **SAP ingress**

IP and MAC filter policies applied on the SAP ingress define the Service Level Agreement (SLA) enforcement of service packets as they ingress a SAP according to the filter policy match criteria.

- **SAP egress**

Filter policies applied on the SAP egress define the Service Level Agreement (SLA) enforcement for service packets as they egress on the SAP according to the filter policy match criteria.

- **network ingress**

IP filter policies are applied to network ingress IP interfaces.

- **network egress**

IP filter policies are applied to network egress IP interfaces.

4.1.2.2.1 Configuration guidelines for Routed VPLS and ACLs

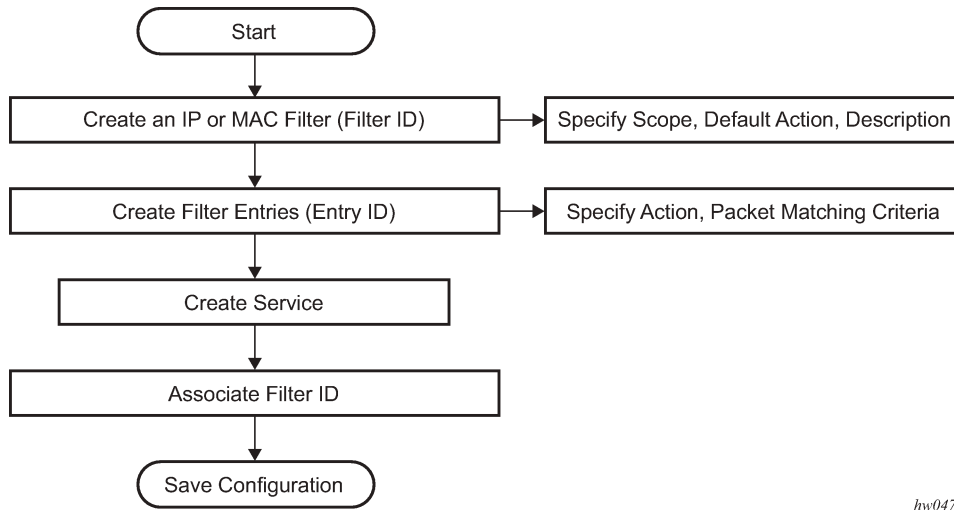
The following information describes Routed VPLS and ACLs configuration guidelines:

- MAC filters are supported on R-VPLS SAPs only on the 7210 SAS-Mxp. See the "Routed VPLS" section in the *7210 SAS-Mxp, S, Sx, T Services Guide* for more information.
- IPv6 filters on RVPLS SAP are supported only on the 7210 SAS-Mxp.
- IP filters using resources from the pool allocated to IPv4 criteria are supported as override filters on all platforms as described in this document.
- IP filters using resources from the pool allocated to IPv4 criteria, and IPv6 filters using resources allocated to IPv6 criteria (both 64-bit and 128-bit resource pools), are supported as override filters only on 7210 SAS-Mxp.
- IP filters using IPv6 resources are not supported on RVPLS SAP.
- Egress override filters are not supported (under the IES and VPRN interface associated with RVPLS).

4.2 Creating and applying policies

The following figure shows the process to create filter policies and apply them to a service network IP interface.

Figure 9: Creating and applying filter policies



hw0473

4.2.1 Packet matching criteria

As few or as many match parameters can be specified as required, but all conditions must be met for the packet to be considered a match and the specified action performed. The process stops when the first complete match is found and then executes the action defined in the entry, either to drop or forward packets that match the criteria.

IP filter policies match criteria that associate traffic with an ingress or egress SAP. Matching criteria to drop or forward IP traffic include:

- **source IP address and mask**

Source IP address and mask values can be entered as search criteria. The IP Version 4 addressing scheme consists of 32 bits expressed in dotted-decimal notation (X.X.X.X).

Address ranges are configured by specifying mask values, the 32-bit combination used to describe the address portion which refers to the subnet and which portion refers to the host. The mask length is expressed as an integer (range 1 to 32).

- **destination IP address and mask**

Destination IP address and mask values can be entered as search criteria. Similar choice as available for source IPv6 addresses is available for destination IPv6 addresses (see above).

- **protocol**

Entering a protocol ID (such as TCP, UDP, and so on) allows the filter to search for the protocol specified in this field.

- **protocol**

For IPv6: entering a next header allows the filter to match the first next header following the IPv6 header.

- **source port**

Entering the source port number allows the filter to search for matching TCP or UDP port values.

- **destination port**
Entering the destination port number allows the filter to search for matching TCP or UDP port.
- **DSCP marking**
Entering a DSCP marking enables the filter to search for the DSCP marking specified in this field. See [Table 42: DSCP name to DSCP value table](#) .
- **ICMP code**
Entering an ICMP code allows the filter to search for matching ICMP code in the ICMP header.
- **ICMP type**
Entering an ICMP type allows the filter to search for matching ICMP types in the ICMP header.
- IPv4 filter created in the mode to use IPv6 resource cannot be applied at egress SAP. Similarly IPv4 filter created in the mode to use IPv6 resource fails to match fragment option.
- **fragmentation**
IPv4 only: Enable fragmentation matching. A match occurs if packets have either the MF (more fragment) bit set or have the Fragment Offset field of the IP header set to a non-zero value.
- **option present**
Enabling the option presence allows the filter to search for presence or absence of IP options in the packet. Padding and EOOL are also considered as IP options.
- **TCP-ACK/SYN flags**
Entering a TCP-SYN/TCP-ACK flag allows the filter to search for the TCP flags specified in these fields.

MAC filter policies match criteria that associate traffic with an ingress or egress SAP. Matching criteria to drop or forward MAC traffic include:

- **source MAC address and mask**
Entering the source MAC address range allows the filter to search for matching a source MAC address or range. Enter the source MAC address and mask in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx; for example, 00:dc:98:1d:00:00.
- **destination MAC address and mask**
Entering the destination MAC address range allows the filter to search for matching a destination MAC address or range. Enter the destination MAC address and mask in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx; for example, 02:dc:98:1d:00:01.
- **dot1p and mask**
Entering an IEEE 802.1p value or range allows the filter to search for matching 802.1p frame. The Dot1p and mask accepts decimal, hex, or binary in the range of 0 to 7.
- **Ethertype**
Entering an Ethernet type II Ethertype value to be used as a filter match criterion. The Ethernet type field is a two-byte field used to identify the protocol carried by the Ethernet frame. The Ethertype accepts decimal, hex, or binary in the range of 1536 to 65535.

4.2.1.1 DSCP values

The following table lists DSCP values.

Table 42: DSCP name to DSCP value table

DSCP name	Decimal DSCP value	Hexadecimal DSCP value	Binary DSCP value
default	0	*	
cp1	1		
cp2	2		
cp3	3		
cp4	4		
cp5	5		
cp6	6		
cp7	7	*	
cs1	8		
cp9	9		
af11	11	*	
af12	12	*	
cp13	13		
cp15	15		
cs2	16	*	
cp17	17		
af21	18	*	
cp19	19		
af22	20	*	
cp21	21		
af23	22	*	
cp23	23		
cs3	24	*	
cp25	25		
af31	26	*	

DSCP name	Decimal DSCP value	Hexadecimal DSCP value	Binary DSCP value
cp27	27		
af32	28	*	
cp29	29		
af33	30	*	
cp21	31		
cs4	32	*	
cp33	33		
af41	34	*	
cp35	35		
af42	36	*	
cp37	37		
af43	38	*	
cp39	39		
cs5	40	*	
cp41	41		
cp42	42		
cp43	43		
cp44	44		
cp45	45		
ef	46	*	
cp47	47		
nc1	48	*	(cs6)
cp49	49		
cp50	50		
cp51	51		
cp52	52		
cp53	53		

DSCP name	Decimal DSCP value	Hexadecimal DSCP value	Binary DSCP value
cp54	54		
cp55	55		
cp56	56		
cp57	57		
nc2	58	*	(cs7)
cp60	60		
cp61	61		
cp62	62		

4.2.2 Ordering filter entries

When entries are created, they should be arranged sequentially from the most explicit entry to the least explicit. Filter matching ceases when a packet matches an entry. The entry action is performed on the packet. 7210 SAS supports either drop or forward action. To be considered a match, the packet must meet all the conditions defined in the entry.

Packets are compared to entries in a filter policy in an ascending entry ID order. To reorder entries in a filter policy, edit the entry ID value; for example, to reposition entry ID "6" to a more explicit location, change the entry ID "6" value to entry ID "2".

When a filter consists of a single entry, the filter executes actions as follows:


- If a packet matches all the entry criteria, the entry specified action is performed (drop or forward).
- If a packet does not match all of the entry criteria, the policy default action is performed.

If a filter policy contains two or more entries, packets are compared in ascending entry ID order (1, 2, 3 or 10, 20, 30, and so on):

- Packets are compared with the criteria in the first entry ID.
- If a packet matches all the properties defined in the entry, the entry specified action is executed.
- If a packet does not completely match, the packet continues to the next entry, and then subsequent entries.
- If a packet does not completely match any subsequent entries, then the default action is performed.

The following figure shows an example of several packets forwarded upon matching the filter criteria and several packets traversing through the filter entries and then dropped.

Figure 10: Filtering process example



hw0474

4.2.3 Applying filters

This section provides information about applying filters to entities.

4.2.3.1 Applying a filter to a SAP

During the SAP creation process, ingress and egress filters are selected from a list of qualifying IP and MAC filters. When ingress filters are applied to a SAP, packets received at the SAP are checked against the matching criteria in the filter entries. If the packet completely matches all criteria in an entry, the checking stops and an entry action is performed. If permitted, the traffic is forwarded according to the specification of the action. If the packets do not match, the default filter action is applied. If permitted, the traffic is forwarded.

When egress filters are applied to a SAP, packets received at the egress SAP are checked against the matching criteria in the filter entries. If the packet completely matches all criteria in an entry, the checking stops. If permitted, the traffic is transmitted. If denied, the traffic is dropped. If the packets do not match, the default filter action is applied.

Filters can be added or changed to an existing SAP configuration by modifying the SAP parameters. Filter policies are not operational until they are applied to a SAP and the service enabled.

4.2.3.2 Applying a filter to a network IP interface

An IP filter can be applied to a network port IP interface. Packets received on the interface are checked against the matching criteria in the filter entries. If the packet completely matches all criteria in an entry, the checking stops. If permitted, the traffic is forwarded. If the packets do not match, they are discarded or forwarded based on the default action specified in the policy.

4.3 Configuration notes



Note:

See the *7210 SAS-Mxp, S, Sx, T Services Guide* and the *7210 SAS-R6, R12 Services Guide* for service-specific ACL support and restrictions.

The following information describes filter implementation guidelines and restrictions:

- Creating a filter policy is optional.
- Associating a service with a filter policy is optional.
- When a filter policy is configured, it should be defined as having either an **exclusive** scope for one-time use, or a **template** scope meaning that the filter can be applied to multiple SAPs.
- A specific filter must be explicitly associated with a specific service in order for packets to be matched.
- A filter policy can consist of zero or more filter entry. Each entry represents a collection of filter match criteria. When packets enter the ingress or egress ports, packets are compared to the criteria specified within the entry or entries.
- When a large (complex) filter is configured, it may take a few seconds to load the filter policy configuration and be instantiated.
- The action keyword must be entered for the entry to be active. Any filter entry without the action keyword will be considered incomplete and be inactive.
- When a filter policy is created with the option `ipv6-64bit-address`, the entries can only use only the IPv6 `src-ip` and `dst-ip` fields in the match criteria.
- When a filter policy is created with the option `ipv6-128bit-address`, the entries can use the IPv6 `src-ip`, `dst-ip`, IPv6 DSCP, TCP/UDP port numbers (source and destination port), ICMP code and type, and TCP flags fields in the match criteria.
- The resources must be allocated for use by ingress IPv6 filters, before associating an IPv6 filter policy to a SAP. By default, the software does not enable the use of IPv6 resources. Until resources are allocated for use by IPv6 filters, software fails all attempts to associate a IPv6 filter policy with a SAP.
- The available ingress CAM hardware resources can be allocated as per user needs for use with different filter criteria using the commands under `config> system>resource-profile>ingress-internal-tcam>acl-sap-ingress`. By default, the system allocates resources to maintain backward compatibility

with Release 4.0. Users can modify the resource allocation based on their need to scale the number of entries or number of associations (that is, number of SAP/IP interfaces using a filter policy that defines a particular match criterion).

- The available egress CAM hardware resources can be allocated as per user needs for use with different filter criteria using the commands under **config> system>resource-profile>egress-internal-tcam>acl-sap-egress**. By default, the system allocates resources to maintain backward compatibility with Release 4.0. Users can modify the resource allocation based on their needs to scale the number of entries or the number of associations (that is, number of SAP/IP interfaces using a filter policy that defines a particular match criterion).
- IPv6 ACLs and MAC QoS policies cannot co-exist on the SAP.
- If no CAM resources are allocated to a particular match criterion defined in a filter policy, then the association of that filter policy to a SAP will fail. This is true for both ingress and egress filter policy.
- IPv6 ACLs and MAC QoS policies cannot co-exist on the SAP.
- For traffic ingressing a B-VPLS SAP and destined to a B-VPLS SAP, the MAC filter matches the B-domain, MAC header fields (that is, B-DA, B-SA, and others). The MAC filter can be used to match customer payload MAC header fields for traffic ingressing a B-VPLS SAP and destined to an I-VPLS SAP.

4.3.1 MAC filters

The following information describes MAC filters:

- If a MAC filter policy is created with an entry and entry action specified but the packet matching criteria is not defined, then all packets processed through this filter policy entry will pass and take the action specified. There are no default parameters defined for matching criteria.
- MAC filters cannot be applied to network interfaces. On the 7210 SAS-Mxp, they can be applied on an R-VPLS service with IES or VPRN. See [Applying filter policies](#) and the *7210 SAS-Mxp, S, Sx, T Services Guide* for more information.
- Some of the MAC match criteria fields are exclusive to each other, based on the type of Ethernet frame. Use the following table to determine the exclusivity of fields. On the 7210 SAS, the default frame-format is "EthernetII".

Table 43: MAC match criteria exclusivity rules

Frame format	Etype
Ethernet – II	Yes
802.3	No
802.3 – snap	No

4.3.2 IP filters

The following information describes IP filterSs:

- **define filter entry packet matching criteria**

If a filter policy is created with an entry and entry action specified but the packet matching criteria is not defined, then all packets processed through this filter policy entry will pass and take the action specified. There are no default parameters defined for matching criteria.

- **action**

An action parameter must be specified for the entry to be active. Any filter entry without an action parameter specified will be considered incomplete and be inactive.

4.3.3 IPv6 filters

The following information describes IPv6 filters:

- **define filter entry packet matching criteria**

If a filter policy is created with an entry and entry action specified, but the packet matching criteria is not defined, then all packets processed through this filter policy entry passes and takes the action specified. There are no default parameters defined for matching criteria.

- **action**

An action parameter must be specified for the entry to be active. Any filter entry without an action parameter specified is considered incomplete and inactive.

4.3.3.1 Resource usage for ingress filter policies



Note:

- The number of entries per slice/chunk is different for both ingress-internal-tcam resource pool and egress-internal-tcam resource pool for different platforms.
- The example below assumes number of entries to be 256 per slice/chunk for ingress-internal-tcam resource pool (for example: 7210 SAS-T). It is valid for other platforms with suitable modification of number of entries per slice.

When the user allocates resources from the ingress CAM resource pool for use by filter policies using the **configure> system> resource-profile** CLI commands, the system allocates resources in chunks of fixed-size entries (example - 256 entries per chunk on 7210 SAS-T). The usage of these entries by different type of match criteria is described below:

- **mac-criteria** - User needs to allocate resources for mac-criteria from the filter resource pool by using the command **config>system>resource-profile>ingress-internal-tcam>acl-sap-ingress>mac-match-enable** before using ingress ACLs with mac-criteria. Every entry configured in the filter policy using the mac-criteria uses one (1) entry from the chunks allocated for use by mac-criteria in the hardware. For example, assume a filter policy is configured with 50 entries and uses **config>system>resource-profile>ingress-internal-tcam>acl-sap-ingress>mac-match-enable 1**, the user configures one chunk for use by mac-criteria (allowing a total of 256 entries. one reserved for internal use entries for use by SAPs using filter policies that use mac-criteria). In this case, the user can have 5 SAPs using mac-criteria filter policy and consumes 250 entries.
- **ipv4-criteria** - User needs to allocate resources for IPv4 criteria from the filter resource pool by using the command **config>system>resource-profile>ingress-internal-tcam>acl-sap-ingress>ipv4-match-enable** before using ingress ACLs with IPv4 criteria. The resource usage per IPv4 match entry is same as the mac-criteria. Please check the preceding example. When created with **use-ipv6-resource**, the resource usage is the same as IPv6 filters using ipv6-128-bit-addresses.

- **ipv4-criteria with port range values** - Users must allocate resources for the IPv4 criteria using the port-range values from the filter resource pool by using the **config>system>resource-profile>ingress-internal-tcam>acl-sap-ingress>ipv4-port-range-match-enable** command before using ingress ACLs with an IPv4 criteria port range. Every entry configured in the filter policy using the IPv4 criteria port range uses two entries from the chunks allocated for use by the IPv4 criteria port range in the hardware. For example, if a filter policy is configured with 50 entries and includes the **config>system>resource-profile>ingress-internal-tcam>acl-sap-ingress>ipv4-port-range-match-enable 2** configuration, the user configures one chunk for use by the IPv4 criteria port range values (allowing a total of 256 entries; one entry is reserved for internal use entries for use by SAPs using filter policies that use an IPv4 criteria port range). In this case, the user can have two filter policies using the IPv4 criteria filter policy and consumes 200 entries (each policy requires 100 entries, as each of the 50 entries must have two entries each).
- **ipv6-criteria using ipv6-64-bit addresses** - User needs to allocate resources for IPv6 criteria with 64-bit address match from the filter resource pool by using the command **config>system>resource-profile>ingress-internal-tcam>acl-sap-ingress>ipv6-64only-match-enable** before using ingress ACLs with ipv6-criteria that use only IPv6 64-bit address for source and destination IPv6 addresses. The IPv6 headers fields available for match is limited. Please see the CLI description for filter below for more information. The usage is same as the ipv4 and mac-criteria.
- **ipv6-criteria using ipv6-128-bit addresses** - User needs to allocate resources for IPv6 criteria with 128-bit address match from the filter resource pool by using the command **config>system>resource-profile>ingress-internal-tcam>acl-sap-ingress>ipv4-ipv6-128-match-enable** before using ingress ACLs with ipv6-criteria that use only IPv6 128-bit address for source and destination IPv6 addresses. These resources can be shared by a policy that uses only IPv4 criteria entries. Every entry configured in the filter policy using the ipv6-criteria with 128-bit addresses uses two entries from the chunks allocated for use by IPv6 criteria (128-bit) in the hardware. For example: Assume a filter policy is configured with 50 entries and using **config>system>resource-profile>ingress-internal-tcam>acl-sap-ingress>ipv4-ipv6-128-match-enable 1**, the user configures one chunk for use by IPv6 criteria with 128-bit addresses (allowing for a total of 128 entries for use by SAPs using filter policies that use this criteria). In this case, user can have five (5) SAPs using this filter policy and consumes 125 entries. Note when a chunk is allocated to IPv6 criteria, software automatically adjusts the number of available entries in that chunk to 128, instead of 256, since two entries are needed to match IPv6 fields.

The users can use the **tools>dump>system-resources** command to know the current usage and availability. For example: Though chunks are allocated in 256 entries, only 128 entries show up against filters using those of IPv6 128-bit addresses. One or more entries are reserved for system use and is not available for user.

4.3.3.2 Resource usage for egress filter policies

When the user allocates resources for use by filter policies using the **config>system>resource-profile>egress-internal-tcam** CLI commands, the system allocates resources in chunks of fixed-size entries (example - 256 entries per slice on 7210 SAS-Mxp) from the egress internal tcam pool in hardware. The usage of these entries by different type of match criteria is described below:

- **mac-criteria** - The user needs to allocate resources for using mac-criteria using the command **config>system>resource-profile>egress-internal-tcam>acl-sap-egress>mac-match-enable 2** or **config>system>resource-profile>egress-internal-tcam>acl-sap-egress>mac-ipv4-match-enable 2** or **config>system>resource-profile>egress-internal-tcam>acl-sap-egress>mac-ipv6-64bit-match-enable 2**. In the last two cases, the resources can be shared with SAPs that use IPv4 or IPv6 64-bit filter policies. The first case allocates resources for exclusive use by MAC filter policies. The resource usage varies based how resources have been allocated:

- If resources are allocated for use by mac-criteria only (using mac-match-enable), then every entry configured in the filter policy uses one (1) entry from the chunks allocated for use by mac-criteria in the hardware. For example: Assume a filter policy is configured with 25 mac-criteria entries and uses **config>system>resource-profile>egress-internal-tcam>acl-sap-egress>mac-match-enable 2**, the user configures two chunks (each chunk having 256 entries each) for use by mac-criteria, allowing a total of 512 entries for use by SAPs using filter policies that use mac-criteria. Therefore, the user can have about 10 SAPs using mac-criteria filter policy and consumes 500 entries. With this, SAPs using ipv4 criteria or ipv6 criteria cannot share the resources along with SAPs using mac-criteria.
- If the resources are allocated for sharing between mac-criteria and ipv4-criteria, then every entry configured in the filter policy uses 2 (two) entries from the chunks allocated in hardware. For example, assume a filter policy is configured with 25 mac-criteria entries and another filter policy configured with 25 IPv4 criteria entries and, with mac-ipv4-match-enable set to 2, that is, user configures two chunks (each chunk having 256 entries each) for sharing between MAC and IPv4, allowing for a total of 128 entries for use by SAPs that use filter policies using ipv4-criteria or mac-criteria. Therefore, the user can have about 5 SAPs using filter policies, such that 3 SAPs uses mac-criteria and the other 2 SAPs use ipv4-criteria or any combination thereof.
- If the resources are allocated for sharing between mac-criteria and ipv6-64bit-criteria, then every entry configured in the filter policy uses 2 (two) entries from the chunks allocated in hardware.

For example, assume a filter policy is configured with 50 mac-criteria entries and another filter policy configured with 50 IPv6 64-bit criteria entries and, with mac-ipv6-64bit-match-enable set to 2, that is, user configures two chunks (with 256 entries each) for sharing between MAC and IPv6-64bit, allowing for a total of 128 entries for use by SAPs that use filter policies using ipv6-64bit-criteria or mac-criteria. Therefore, the user can have about 2 SAPs using filter policies, such that one SAP uses mac-criteria and the other one SAP uses ipv6-64bit-criteria or any combination thereof.

- **ipv4-criteria** - The user need to allocate resources using the command **config>system>resource-profile>egress-internal-tcam>acl-sap-egress>mac-ipv4-match-enable**. The resource usage is as described previously.
- **ipv6-criteria using ipv6-64-bit addresses** - The user need to allocate resources using the command **config>system>resource-profile>egress-internal-tcam>acl-sap-egress>mac-ipv6-64bit-match-enable**. The resource usage is as described previously.
- **ipv6-criteria using ipv6-128-bit addresses** - The user need to allocate resources using the command **config>system>resource-profile>egress-internal-tcam>acl-sap-egress>ipv6-128bit-match-enable**. This command allocates resources for exclusive by IPv6-128bit criteria filter policies and cannot be shared by SAPs using any another criteria. If resources are allocated for use by ipv6-128bit-criteria only, then every entry configured in the filter policy uses two (2) entries from the chunks allocated for use in hardware. For example: Assume a filter policy is configured with 50 ipv6-128bit-criteria entries and user uses **config>system>resource-profile>egress-internal-tcam>acl-sap-egress>ipv6-128bit-match-enable 2**, to configure two chunks (each chunk having 256 entries each) for use by ipv6-128bit-criteria. This allows for a total of 128 entries for use by SAPs using filter policies that use ipv6-128bit-criteria. Therefore the user can have about 2 SAPs using ipv6-128bit-criteria filter policy and consumes 100 entries.

The user can use the **tools>dump>system-resources** command to know the current usage and availability.

4.4 Configuring filter policies with CLI

This section provides information to configure filter policies using the command line interface.

4.5 Basic configuration

The most basic IP and MAC filter policies must have the following:

- a filter ID
- template scope, either **exclusive** or **template**
- default action, either drop or forward
- at least one filter entry
 - specified action, either drop or forward
 - specified matching criteria
- allocates the required amount of resources for ingress and egress filter policies

Output example

The following is a sample configuration output of allocation of ingress internal CAM resources for ingress policy for 7210 SAS.

```
*A:7210SAS>config>system>res-prof>ing-internal-tcam>acl-sap-ing# info detail
-----
                ipv4-match-enable max
                ipv6-64-only-match-enable 1
                no ipv4-ipv6-128-match-enable
-----
*A:7210SAS>config>system>res-prof>ing-internal-tcam>acl-sap-ing# back
```

Output example

The following is a sample configuration output of allocation of egress internal CAM resources for egress policy for 7210 SAS-Sx/S 1/10GE.

```
A:7210SAS>config>system>res-prof>egr-internal-tcam# info detail
-----
                acl-sap-egress 2
                mac-ipv4-match-enable 2
                ipv6-128bit-match-enable 0
                mac-ipv6-64bit-match-enable 0
                mac-match-enable 0
                exit
-----
*A:7210SAS>config>system>res-prof>egr-internal-tcam# acl-sap-egress
```

Output example

The following is a sample configuration output of allocation of egress internal CAM resources for egress policy for 7210 SAS-Sx 10/100GE.

```
*A:7210SAS>config>system>res-prof>egr-internal-tcam# info detail
-----
      acl-sap-egress 2
        mac-ipv4-match-enable 2
        ipv6-128bit-match-enable 0
        ipv6-64bit-match-enable 0
        mac-match-enable 0
      exit
      no egress-sap-aggregate-meter
-----
*A:7210SAS>config>system>res-prof>egr-internal-tcam# acl-sap-egress
```

Output example

The following is a sample configuration output of allocation of egress internal CAM resources for egress policy for 7210 SAS-Mxp.

```
*A:sim_dutc>config>system>res-prof>egr-internal-tcam>acl-sap-egr# info detail
-----
      mac-ipv4-match-enable 2
      ipv6-128bit-match-enable 0
      mac-ipv6-64bit-match-enable 0
      mac-match-enable 0
-----
*A:sim_dutc>config>system>res-prof>egr-internal-tcam>acl-sap-egr#
```


Output example

The following is a sample configuration output of an IP filter policy. The configuration blocks all incoming TCP session except Telnet and allows all outgoing TCP sessions from IP net 10.67.132.0/24. CAM resources must be allocated to IPv4 criteria before associating the filter with a SAP. [Figure 11: Applying an IP filter to an ingress interface](#) shows the interface to apply the filter.

```
A:ALA-1>config>filter# info
-----
      ip-filter 3 create
        entry 10 create
          match protocol 6
            dst-port eq 23
            src-ip 10.67.132.0/24
          exit
          action
            forward
        exit
        entry 20 create
          match protocol 6
            tcp-syn true
            tcp-ack false
          exit
          action
            drop
        exit
      exit
-----
A:ALA-1>config>filter#
```

The following figure shows the IP filter applied to an ingress interface.

Figure 11: Applying an IP filter to an ingress interface



4.6 Common configuration tasks

This section provides a brief overview of the tasks that must be performed for both IP and MAC filter configurations and provides the CLI commands.

4.6.1 Creating an IP filter policy

Configuring and applying filter policies is optional. Each filter policy must have the following:

- the filter type specified (IP)
- a filter policy ID
- a default action, either **drop** or **forward**
- filter policy scope specified, either **exclusive** or **template**
- at least one filter entry with matching criteria specified
- configure CAM hardware resource for use by the filter policy match-criteria

4.6.1.1 IP filter policy

Output example

The following is a sample exclusive filter policy configuration output.

```
A:ALA-7>config>filter# info
-----
...
    ip-filter 12 create
        description "IP-filter"
        scope exclusive
    exit
...
-----
A:ALA-7>config>filter#
```

4.6.1.2 IP filter entry

Within a filter policy, configure filter entries which contain criteria against which ingress, egress, or network traffic is matched. The action specified in the entry determine how the packets are handled, either dropped or forwarded:

- Enter a filter entry ID. The system does not dynamically assign a value.
- Assign an action, either drop or forward.
- Specify matching criteria.

Use the following syntax to create an IP filter entry.

```
config>filter# ip-filter filter-id [create]
      entry entry-id [time-range time-range-name][create]
      description description-string
```

Output example

The following is a sample IP filter entry configuration output.

```
A:ALA-7>config>filter>ip-filter# info
-----
      description "filter-main"
      scope exclusive
      entry 10 create
          description "no-91"
          match
          exit
          no action
      exit
  exit
-----
A:ALA-7>config>filter>ip-filter#
```

4.6.1.3 IP entry matching criteria

Use the following syntax to configure IP filter matching criteria.

Output example

The following is a sample IP filter matching configuration output.

```
*A:ALA-48>config>filter>ip-filter# info
-----
      description "filter-mail"
      scope exclusive
      entry 10 create
          description "no-91"

          match
            dst-ip 10.10.10.91/24
            src-ip 10.10.10.103/24
          exit
          action
            forward
      exit
```

```
-----  
*A:ALA-48>config>filter>ip-filter#
```

4.6.2 Creating an IPv6 filter policy

Configuring and applying IPv6 filter policies is optional. Each filter policy must have the following:

- the IPv6 filter type specified
- an IPv6 filter policy ID
- a default action, either **drop** or **forward**
- template scope specified, either **exclusive** or **template**
- at least one filter entry with matching criteria specified

4.6.2.1 IPv6 filter policy

Use the following syntax to create an IPv6 filter policy.

To create an IPv6 filter using 64-bit-address, the user can use the command **config> filter>ipv6-filter filter-id ipv6-64bit-address create**.

By default, the IPv6 filters are configured using 128-bit-address, the output is as follows.

```
*A:7210SAS>config>filter>ipv6-filter# info detail  
-----  
          default-action drop  
          no description  
          scope template  
          exit  
*A:7210SAS>config>filter>ipv6-filter#
```

4.6.2.2 IPv6 filter entry

Within an IPv6 filter policy, configure filter entries which contain criteria against which ingress, egress, or network traffic is matched. The action specified in the entry determine how the packets are handled, either dropped or forwarded:

- Enter an IPv6 filter entry ID. The system does not dynamically assign a value.
- Assign an action, either drop or forward.
- Specify matching criteria.

Output example

The following is a sample IPv6 filter entry configuration output.

```
*A:7210SAS>config>filter>ipv6-filter# info detail  
-----  
          default-action drop  
          no description  
          scope template  
          entry 1 create  
          no description
```

```
        match next-header none
          no dscp
          no dst-ip
          no dst-port
          src-ip 2001:db8::1/128
          no src-port
          no tcp-syn
          no tcp-ack
          no icmp-type
          no icmp-code
        exit
      action
        forward
    exit
*A:7210SAS>config>filter>ipv6-filter#
```

4.6.3 Creating a MAC filter policy

Configuring and applying filter policies is optional. Each filter policy must have the following:

- The filter type specified (MAC).
- A filter policy ID.
- A default action, either **drop** or **forward**
- Filter policy scope, either **exclusive** or **template**.
- At least one filter entry.
- Matching criteria specified.

4.6.3.1 MAC filter policy

Output example

The following is a sample MAC filter policy configuration output.

```
A:ALA-7>config>filter# info
-----
...
    mac-filter 90 create
      description "filter-west"
      scope exclusive
    exit
-----
A:ALA-7>config>filter#
```

4.6.3.2 MAC filter entry

Within a filter policy, configure filter entries which contain criteria against which ingress, egress, or network traffic is matched. The action specified in the entry determine how the packets are handled, either dropped or forwarded:

- Enter a filter entry ID. The system does not dynamically assign a value.
- Assign an action, either drop or forward.

- Specify matching criteria.

Output example

The following is a sample MAC filter entry configuration output.

```
A:sim1>config>filter# info
-----
      mac-filter 90 create
        entry 1 create
          description "allow-104"
          match
          exit
          action
            drop
        exit
      exit
-----
A:sim1>config>filter#
```

4.6.3.3 MAC entry matching criteria

The following is a sample filter matching configuration output.

Output example: Sample filter matching configuration output

```
A:ALA-7>config>filter>mac-filter# info
-----
      description "filter-west"
      scope exclusive
      entry 1 create
        description "allow-104"
        match
          src-mac 00:dc:98:1d:00:00 ff:ff:ff:ff:ff:ff
          dst-mac 02:dc:98:1d:00:01 ff:ff:ff:ff:ff:ff
        exit
        action
          drop
      exit
-----
```

4.6.3.4 Apply IP and MAC filter policies

Use the following syntax to apply an IP and a MAC filter policy to an Epipe service.

```
config>service# epipe service-id
  sap sap-id
  egress
    filter {ip ip-filter-id | mac mac-filter-id}
  ingress
    filter {ip ip-filter-id | mac mac-filter-id}
```

Output example

The following is a sample of IP and MAC filters assigned to an ingress and egress SAP output.

```
A:ALA-48>config>service>epipe# info
```

```
-----  
      sap 1/1/1.1.1 create  
      ingress  
          filter ip 10  
      exit  
      egress  
          filter mac 92  
      exit  
      exit  
      no shutdown  
-----  
A:ALA-48>config>service>epipe#
```

4.6.3.5 Apply an IPv6 filter policy to VPLS

Output example

The following is a sample of IPv6 filters assigned to VPLS service interface output.

```
*A:7210SAS>config>service#vpls#sap info detail  
-----  
.....  
      ingress  
          counter-mode in-out-profile-count  
          no drop-count-extra-vlan-tag-pkts  
      exit  
      exit  
      ingress  
          qos 1  
          no aggregate-meter-rate  
          filter ipv6 1  
      exit  
      egress  
          no filter  
      exit  
      no collect-stats  
      no accounting-policy  
      no shutdown  
      exit  
*A:7210SAS>config>service#vpls#sap info detail
```

4.6.4 Applying filter policies to a network IP interface

IP filter policies can be applied to network IP interfaces. MAC filters cannot be applied to network IP interfaces or to routable IES services.

4.6.4.1 Applying a filter policy to an IP interface

```
config>router# interface ip-int-name
```


Output example

The following is a sample IP filter applied to an interface at ingress output.

```
A:ALA-48>config>router# info
#-----
# IP Configuration
#-----
...
    interface "to-104"
      address 10.0.0.103/24
      port 1/1/1
      ingress
        filter ip 10
      exit
      egress
        filter ip 10
      exit
    exit
...
#-----
A:ALA-48>config>router#
```

4.7 Filter management tasks

This section describes the filter policy management tasks.

4.7.1 Renumbering filter policy entries

The system exits the matching process when the first match is found and then executes the actions in accordance with the specified action. Because the ordering of entries is important, the numbering sequence can be rearranged. Entries should be numbered from the most explicit to the least explicit.

Use the following syntax to renumber existing MAC or IP filter entries to re-sequence filter entries.

```
config>filter
ip-filter filter-id
renum old-entry-numbernew-entry-number
mac-filter filter-id
renum old-entry-numbernew-entry-number
```

Example

```
config>filter>ip-filter# renum 10 15
config>filter>ip-filter# renum 20 10
config>filter>ip-filter# renum 40 1
```

Output example

The following is a sample of the original filter entry order output.

```
A:ALA-7>config>filter# info
-----
...
```

```
ip-filter 11 create
  description "filter-main"
  scope exclusive
  entry 10 create
    description "no-91"
    match
      dst-ip 10.10.10.91/24
      src-ip 10.10.10.103/24
    exit
    action forward
  exit
  entry 20 create
    match
      dst-ip 10.10.10.91/24
      src-ip 10.10.0.100/24
    exit
    action drop
  exit
  entry 30 create
    match
      dst-ip 10.10.10.91/24
      src-ip 10.10.0.200/24
    exit
    action forward
  exit
  entry 40 create
    match
      dst-ip 10.10.10.91/24
      src-ip 10.10.10.106/24
    exit
    action drop
  exit
exit
...
-----
A:ALA-7>config>filter#
```

Example

The following is a sample of the reordered filter entries output.

```
A:ALA-7>config>filter# info
-----
...
ip-filter 11 create
  description "filter-main"
  scope exclusive
  entry 1 create
    match
      dst-ip 10.10.10.91/24
      src-ip 10.10.10.106/24
    exit
    action drop
  exit
  entry 10 create
    match
      dst-ip 10.10.10.91/24
      src-ip 10.10.0.100/24
    exit
    action drop
  exit
  entry 15 create
    description "no-91"
```

```

        match
            dst-ip 10.10.10.91/24
            src-ip 10.10.10.103/24
        exit
        action forward

    exit
    entry 30 create
        match
            dst-ip 10.10.10.91/24
            src-ip 10.10.0.200/24
        exit
        action forward
    exit
exit
...
-----
A:ALA-7>config>filter#

```

4.7.2 Modifying an IP filter policy

To access a specific IP filter, you must specify the filter ID. Use the **no** form of this command to remove the command parameters or return the parameter to the default setting.

Example

```

config>filter>ip-filter# description "New IP filter info"
config>filter>ip-filter# entry 2 create
config>filter>ip-filter>entry$ description "new entry"
config>filter>ip-filter>entry# action drop
config>filter>ip-filter>entry# match dst-ip 10.10.10.104/32
config>filter>ip-filter>entry# exit
config>filter>ip-filter#

```

Output example

The following is a sample of the modified IP filter output.

```

A:ALA-7>config>filter# info
-----
...
    ip-filter 11 create
        description "New IP filter info"
        scope exclusive
        entry 1 create
            match
                dst-ip 10.10.10.91/24
                src-ip 10.10.10.106/24
            exit
            action
                drop
        exit
        entry 2 create
            description "new entry"
            match
                dst-ip 10.10.10.104/32
            exit

```

```
        action
            drop
    exit
    entry 10 create
        match
            dst-ip 10.10.10.91/24
            src-ip 10.10.0.100/24
        exit
        action
            drop
    exit
    entry 15 create
        description "no-91"
        match
            dst-ip 10.10.10.91/24
            src-ip 10.10.10.103/24
        exit
        action
            forward
    exit
    entry 30 create
        match
            dst-ip 10.10.10.91/24
            src-ip 10.10.0.200/24
        exit
        action
            forward
    exit
exit
..
-----
A:ALA-7>config>filter#
```

4.7.3 Modifying an IPv6 filter policy

To access a specific IPv6 filter, you must specify the filter ID. Use the **no** form of this command to remove the command parameters or return the parameter to the default setting.

Example

```
config>filter# ipv6-filter 11
config>filter>ipv6-filter# description "IPv6 filter for Customer 1"
config>filter>ipv6-filter# scope exclusive
config>filter>ipv6-filter# entry 1
config>filter>ipv6-filter>entry# description "Fwds matching packets"
config>filter>ipv6-filter>entry# action forward
config>filter>ipv6-filter>entry# exit
```

Output example

The following is a sample output of the modified IPv6 filter output.

```
A:7210SAS>config>filter>ipv6-filter# info detail
-----
    default-action drop
    no description
    scope template
    entry 1 create
        description "Test"
        match next-header none
```

```
        no dscp
        no dst-ip
        no dst-port
        src-ip 2001:db8::1/128
        no src-port
        no tcp-syn
        no tcp-ack
        no icmp-type
        no icmp-code
    exit
    action
        forward
    exit
...
A:7210SAS>config>filter>ipv6-filter
```

4.7.4 Modifying a MAC filter policy

To access a specific MAC filter, you must specify the filter ID. Use the **no** form of this command to remove the command parameters or return the parameter to the default setting.

Example

```
- config>filter# mac-filter 90
- config>filter>mac-filter# description "New filter info"
- config>filter>mac-filter# entry 1
- config>filter>mac-filter>entry# description "New entry info"
- config>filter>mac-filter>entry# action forward
- config>filter>mac-filter>entry# exit
- config>filter>mac-filter# entry 2 create
- config>filter>mac-filter>entry$ action drop
- config>filter>mac-filter>entry# match
- config>filter>mac-filter>entry>match# dot1p 7 7
```

Output example

The following is a sample of the modified MAC filter output.

```
A:ALA-7>config>filter# info
-----
...
    mac-filter 90 create
    description "New filter info"
    scope exclusive
    entry 1 create
        description "New entry info"
        match
            src-mac 00:dc:98:1d:00:00 ff:ff:ff:ff:ff:ff
            dst-mac 02:dc:98:1d:00:01 ff:ff:ff:ff:ff:ff
        exit
        action
            forward
    exit
    entry 2 create
        match
            dot1p 7 7
        exit
        action
            drop
```

```
        exit
      exit
    ...
-----
A:ALA-7>config>filter#
```

4.7.5 Detaching/deleting a filter policy

Before you can delete a filter, you must remove the filter association from the applied ingress and egress SAPs and network interfaces.

4.7.5.1 From an ingress SAP

The following shows the command usage to remove a filter from an ingress SAP.

```
config>service# [epipe | vpls] service-id
  sap port-id[:encap-val]
  ingress
  no filter
```

Example

```
config>service# epipe 5
config>service>epipe# sap 1/1/2:3
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress# no filter
```

4.7.5.2 From an egress SAP

The following shows the command usage to remove a filter from an egress SAP.

```
config>service# [epipe | vpls] service-id
  sap port-id[:encap-val]
  egress
  no filter
```

Example

```
config>service# epipe 5
config>service>epipe# sap 1/1/2:3
config>service>epipe>sap# egress
config>service>epipe>sap>egress# no filter
```

4.7.5.3 From a network interface

The following shows the command usage to delete a filter from a network interface.

```
config>router# interface ip-int-name
  ingress
```

Example

```
config>router>if>ingress# no filter ip 2
config>router>if>ingress# exit
```

4.7.5.4 From the filter configuration

Use the following syntax to delete the filter after you have removed the filter from the SAP.

```
config>filter# no ip-filter filter-id
```

```
config>filter# no mac-filter filter-id
```

Example

```
config>filter# no ip-filter 11
config>filter# no mac-filter 13
```

4.7.6 Copying filter policies

When changes are made to an existing filter policy, they are applied immediately to all services where the policy is applied. If numerous changes are required, the policy can be copied so you can edit the "work in progress" version without affecting the filtering process. When the changes are completed, you can overwrite the work in progress version with the original version.

New filter policies can also be created by copying an existing policy and renaming the new filter.

```
config>filter# copy filter-type src-filter-id [src-entry src-entry-id] to dst-filter-id [dst-
entry dst-entry-id][overwrite]
```

The following shows the command usage to copy an existing IP filter (11) to create a new filter policy (12).

Example

```
config>filter# copy ip-filter 11 to 12
```

Output example

```
A:ALA-7>config>filter# info
-----
...
    ip-filter 11 create
      description "This is new"
      scope exclusive
      entry 1 create
        match
          dst-ip 10.10.10.91/24
          src-ip 10.10.10.106/24
        exit
      action
        drop
      exit
    entry 2 create
```

```
...
    ip-filter 12 create
      description "This is new"
      scope exclusive
      entry 1 create
        match
          dst-ip 10.10.10.91/24
          src-ip 10.10.10.106/24
        exit
      action
        drop
      exit
    entry 2 create
  ...
-----
A:ALA-7>config>filter#
```

4.8 Filter command reference

4.8.1 Command hierarchies

- [Configuration commands](#)
 - [IP filter policy commands](#)
 - [IPv6 filter policy commands](#)
 - [MAC filter policy commands](#)
 - [Generic filter commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Monitor commands](#)

4.8.1.1 Configuration commands

4.8.1.1.1 IP filter policy commands

```
config
- filter
- [no] ip-filter filter-id [use-ipv6-resource] [create]
- [no] ip-filter filter-id [use-port-range] [create]
- no ip-filter filter-id
  - default-action {drop | forward}
  - description description-string
  - no description
  - filter-name filter-name
  - no filter-name
  - renam old-entry-id new-entry-id
  - scope {exclusive | template}
  - no scope
  - entry entry-id [time-range time-range-name] [create]
```



```

- no entry entry-id
- action
- no action
- drop
- forward
- description description-string
- no description
- match [protocol protocol-id]
- no match
- dscp dscp-name
- no dscp
- dst-ip {ip-address/mask | ip-address ipv4-address-mask}
- no dst-ip
- dst-port {eq} dst-port-number
- dst-port range dst-port-number1 dst-port-number2
- no dst-port
- fragment {true | false}
- no fragment
- icmp-code icmp-code
- no icmp-code
- icmp-type icmp-type
- no icmp-type
- option-present {true | false}
- no option-present
- src-ip {ip-address/mask | ip-address ipv4-address-mask}
- no src-ip
- src-port {eq} src-port-number
- src-port range src-port-number1 src-port-number2
- no src-port
- tcp-ack {true | false}
- no tcp-ack
- tcp-syn {true | false}
- no tcp-syn
    
```

4.8.1.1.2 IPv6 filter policy commands

```

config
- filter
- ipv6-filter ipv6-filter-id [ipv6-128bit-address | ipv6-64bit-address] [create]
- no ipv6-filter ipv6-filter-id
- default-action {drop | forward}
- description description-string
- no description
- filter-name filter-name
- no filter-name
- entry entry-id [time-range time-range-name] [create]
- no entry entry-id
- action
- no action
- drop
- forward
- description description-string
- no description
- match [next-header next-header]
- no match
- dscp dscp-name
- no dscp
- dst-ip {ipv6-address/prefix-length}
- no dst-ip
- dst-port {eq} dst-port-number
- no dst-port
    
```

```

- icmp-code icmp-code
- no icmp-code
- icmp-type icmp-type
- no icmp-type
- dst-ip {ipv6-address/prefix-length}
- no dst-ip
- src-port {eq} src-port-number
- src-port range start end}
- no src-port
- src-ip {ipv6-address/prefix-length}
- no src-ip
- tcp-ack {true | false}
- no tcp-ack
- tcp-syn {true | false}
- no tcp-syn
- renum old-entry-id new-entry-id
- scope {exclusive | template}
- no scope
  
```

4.8.1.1.3 MAC filter policy commands

```

config
- filter
  - mac-filter filter-id [create]
  - no mac-filter filter-id
  - default-action {drop | forward}
  - description description-string
  - no description
  - entry entry-id [time-range time-range-name]
  - no entry entry-id
    - description description-string
    - no description
    - action [drop]
    - action forward
    - no action
    - match
    - no match
      - dot1p dot1p-value [dot1p-mask]
      - no dot1p
      - dst-mac ieee-address [ieee-address-mask]
      - no dst-mac
      - etype 0x0600..0xffff
      - no etype
      - src-mac ieee-address [ieee-address-mask]
      - no src-mac
  - filter-name filter-name
  - no filter-name
  - renum old-entry-id new-entry-id
  - scope {exclusive | template}
  - no scope
  - type filter-type
  
```

4.8.1.1.4 Generic filter commands

```

config
- filter
  
```

```
- copy ip-filter | ipv6-filter| mac-filter src-filter-id [src-entry src-entry-id] to dst-filter-id [dst-entry dst-entry-id] [overwrite]
```

4.8.1.2 Show commands

```
show
- filter
  - download-failed
  - ip [ip-filter-id [entry entry-id] [association | counters]]
  - ipv6 [ipv6-filter-id [entry entry-id] [association | counters]]
  - mac {mac-filter-id [entry entry-id] [association | counters]}
```

4.8.1.3 Clear commands

```
clear
- filter
  - ip filter-id [entry entry-id] [ingress | egress]
  - ipv6 filter-id [entry entry-id] [ingress | egress]
  - mac filter-id [entry entry-id] [ingress | egress]
```

4.8.1.4 Monitor commands

```
monitor
- filter
  - ip ip-filter-id entry entry-id [interval seconds] [repeat repeat] [absolute | rate]
  - ipv6 ipv6-filter-id entry entry-id [interval seconds] [repeat repeat] [absolute | rate]
rate]
  - mac mac-filter-id entry entry-id [interval seconds] [repeat repeat] [absolute | rate]
```

4.8.2 Command descriptions

- [Configuration commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Monitor commands](#)

4.8.2.1 Configuration commands

- [Generic commands](#)
- [Global filter commands](#)
- [Filter policy commands](#)
- [General filter entry commands](#)
- [IP filter entry commands](#)
- [MAC filter entry commands](#)

- [IP filter match criteria](#)
- [MAC filter match criteria](#)
- [Policy and entry maintenance commands](#)

4.8.2.1.1 Generic commands

description

Syntax

description *string*

no description

Context

config>filter>ip-filter

config>filter>ip-filter>entry

config>filter>ipv6-filter

config>filter>ipv6-filter>entry

config>filter>mac-filter

config>filter>mac-filter>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a text description for a configuration context to help identify the content in the configuration file.

The **no** form of this command removes any description string from the context.

Parameters

string

Specifies the description character string. Allowed values are any string up to 80 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.

4.8.2.1.2 Global filter commands

ip-filter

Syntax

[no] ip-filter *filter-id* [use-ipv6-resource] [create]

[no] ip-filter *filter-id* [use-port-range] [create]

Context

config>filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure an IP filter policy.

The IP filter policies specify either a forward or a drop action for packets based on the specified match criteria.

The IP filter policy, sometimes referred to as an access control list (ACL), is a template that can be applied to multiple services or multiple network ports as long as the scope of the policy is template.

Any changes made to the existing policy, using any of the subcommands, will be applied immediately to all services where this policy is applied. For this reason, when many changes are required on an IP filter policy, Nokia recommends that the policy be copied to a work area. That work-in-progress policy can be modified until complete and then written over the original filter policy. Use the **config filter copy** command to maintain policies in this manner.

By default, when an IPv4 filter policy is associated with a service entity (For example: SAP), the software attempts to allocate resources for the filter policy entries from the IPv4 resource pool. If resources unavailable in the pool, then the software fails to associate and display an error. If the user knows that resources are free in the IPv6 resource pool, then the **use-ipv6-resource** parameter is used to allow the user to share the entries in the resource chunks allocated for use by IPv6 128-bit resource pool, if available. If this parameter is specified, the resource for this filter policy is always allocated from the IPv6 128-bit filter resource pool.



Note:

By default, IPv4 filters are created using IPv4 resources, assuming an unspecified use-ipv6-resource. If such filters are to be created using IPv6 resources, the **use-ipv6-resource** option needs to be specified. Ahead of the application of such a filter, the user should ensure the number of policies in the newly created policy is within the limit of available resources in the IPv6 128-bit resource pool, by considering the output of the **tools>dump>system-resources** command.

The **no** form of this command deletes the IP filter policy. A filter policy cannot be deleted until it is removed from all SAPs or network ports where it is applied.

Parameters

filter-id

Specifies the IP filter policy ID number.

Values 1 to 65535

create

Specifies that when the context is created, one can navigate into the context without the **create** keyword. This keyword is required when first creating the configuration context.

use-ipv6-resource

Specifies that the hardware resources for the entries in this filter policy must be allocated from the IPv6 filter resource pool, if available.

use-port-range

Keyword to indicate that the IP criteria match entries can use the port range for matching L4 port values and to allocate resources for this filter policy from the **ipv4-port-range-match-enable** pool. This keyword is only available on the 7210 SAS-Mxp.

There are limited amount of unique port range values that can be used (see the **tools>dump>system-resources** command to know the current usage). The port range entries are shared among QoS and ACLs, along with source and destination port range values.

ipv6-filter

Syntax

```
[no] ipv6-filter ipv6-filter-id [ipv6-128bit-address | ipv6-64bit-address] [create]
```

Context

```
config>filter
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates an IPv6 filter policy. During the 'create', the user must specify if IPv6 addresses, both source and destination IPv6 addresses, specified in the match criteria uses complete 128-bits or uses only the upper 64 bits of the IPv6 addresses.

The **no** form of this command deletes the IPv6 filter policy. A filter policy cannot be deleted until it is removed from all SAPs or network ports where it is applied.

Default

128-bit addresses

Parameters

ipv6-filter-id

Specifies the IPv6 filter policy ID number.

Values 1 to 65535

ipv6-128bit-address

Specifies that if the user intends to use complete 128-bit addresses, then the user requires the *ipv6-128bit-address* CLI parameter with the *create* command. When this policy is associated with a SAP, software allocates resources for the filter entries from the IPv6 128-bit resource pool for the SAP.

ipv6-64bit-address

Specifies that if the user intends to use upper most significant bit (MSB) 64-bit addresses, then the user requires the *ipv6-64bit-address* CLI parameter with the *create* command. When this policy is associated with a SAP, the software allocates resources for the filter entries from the IPv6 64-bit resource pool for the SAP. All the IP packet fields are not available for match are when using 64-bit addresses. For more information, see [Configuration notes](#), to know the packet header fields available for match when using this option.

create

Specifies that when the context is created, one can navigate into the context without the **create** keyword. This keyword is required when first creating the configuration context.

mac-filter

Syntax

```
[no] mac-filter filter-id [create]
```

Context

```
config>filter
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the context for a MAC filter policy.

The *mac-filter* policy specifies either a forward or a drop action for packets based on the specified match criteria.

The *mac-filter* policy, sometimes referred to as an access control list, is a template that can be applied to multiple services as long as the scope of the policy is template.



Note:

It is not possible to apply a MAC filter policy to a network port network IP interface.

Any changes made to the existing policy, using any of the subcommands, will be applied immediately to all services where this policy is applied. For this reason, when many changes are required on a *mac-filter*

policy, Nokia recommends that the policy be copied to a work area. That work-in-progress policy can be modified until complete and then written over the original filter policy. Use the **config filter copy** command to maintain policies in this manner.

The **no** form of this command deletes the mac-filter policy. A filter policy cannot be deleted until it is removed from all SAP where it is applied.

Parameters

filter-id

Specifies the MAC filter policy ID number.

Values 1 to 65535

create

Specifies that when the context is created, one can navigate into the context without the **create** keyword. This keyword is required when first creating the configuration context.

4.8.2.1.3 Filter policy commands

default-action

Syntax

default-action {drop | forward}

Context

config>filter>ip-filter

config>filter>ipv6-filter

config>filter>mac-filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the action to be applied to packets when the packets do not match the specified criteria in all of the IP filter entries of the filter.

When multiple **default-action** commands are entered, the last command will overwrite the previous command.

Default

drop

Parameters

drop

Specifies that all packets will be dropped unless there is a specific filter entry which causes the packet to be forwarded.

forward

Specifies that all packets will be forwarded unless there is a specific filter entry which causes the packet to be dropped.

scope

Syntax

scope {**exclusive** | **template**}

no scope

Context

config>filter>ip-filter

config>filter>ipv6-filter

config>filter>mac-filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the filter policy scope as exclusive or template. If the scope of the policy is template and is applied to one or more services or network interfaces, the scope cannot be changed.

The **no** form of this command reverts the scope of the policy to the default.

Default

template

Parameters

exclusive

Specifies that the policy can only be applied to a single entity (SAP or network IP interface). Attempting to assign the policy to a second entity will result in an error message. If the policy is removed from the entity, it will become available for assignment to another entity.

template

Specifies that the policy can be applied to multiple SAPs or network IP interfaces.

4.8.2.1.4 General filter entry commands

entry

Syntax

entry *entry-id* [**time-range** *time-range-name*] [**create**]

no entry *entry-id*

Context

config>filter>ip-filter

config>filter>ipv6-filter

config>filter>mac-filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates or edits an IP or MAC filter entry. Multiple entries can be created using unique *entry-id* numbers within the filter. The implementation exits the filter on the first match found and executes the actions in accordance with the accompanying 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 the **action** command for it to be considered complete. Entries without the **action** command will be considered incomplete and therefore will be rendered inactive.

The **no** form of this command removes the specified entry from the IP or MAC filter. Entries removed from the IP or MAC filter are immediately removed from all services or network ports where that filter is applied.

Parameters

entry-id

Specifies a match criteria and the corresponding action. Nokia recommends that multiple entries be specified *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 65535

time-range *time-range-name*

Specifies the time range name to be associated with this filter entry up to 32 characters. The time-range name must already exist in the **config>cron** context.

create

Specifies that when the context is created, one can navigate into the context without the **create** keyword. This keyword is required when first creating the configuration context.

4.8.2.1.5 IP filter entry commands

action

Syntax

action

no action

Context

config>filter>ip-filter>entry

config>filter>ipv6-filter>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the action taken when a packet meets filtering criteria.

The **no** form of this command removes the currently configured action.

Default

no action

drop

Syntax

drop

Context

config>filter>ip-filter>entry>action

config>filter>ipv6-filter>entry>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the filter action to drop packets matching the filter entry.

forward

Syntax

forward

Context

config>filter>ip-filter>entry>action

config>filter>ipv6-filter>entry>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the filter action to forward packets matching the filter entry.

match

Syntax

match [**protocol**] [*protocol-id*]

no match

Context

config>filter>ip-filter>entry

config>filter>ipv6-filter>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enters match criteria for the filter entry. 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, then all criteria must be satisfied (AND function) before the action associated with the match is executed.

A **match** context may 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

protocol-id

Specifies the decimal value representing the IP protocol used as the IP filter match criterion. Well known protocol IDs, including the ICMP(1), TCP(6), UDP(17), are listed in the following table. The values can be expressed in decimal, hexadecimal, or binary.

Values 0 to 255
 * — udp/tcp

Table 44: IP protocol IDs and 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
idrp	45	Inter-Domain Routing Protocol
rsvp	46	Reservation Protocol
iso-ip	80	ISO Internet Protocol
eigrp	88	EIGRP
ospf-igp	89	OSPF/IGP
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

Protocol	Protocol ID	Description
isis	124	ISIS over IPv4
crtp	126	Combat Radio Transport Protocol
crudp	127	Combat Radio User Datagram

4.8.2.1.6 MAC filter entry commands

action

Syntax

action drop

action forward

no action

Context

config>filter>mac-filter>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the action for a MAC filter entry. The **action** keyword must be entered for the entry to be active. Any filter entry without the **action** keyword will be considered incomplete and will be inactive.

If neither drop nor forward is specified, this is considered a No-Op filter entry used to explicitly set a filter entry inactive without modifying match criteria or removing the entry.

Multiple action statements entered will overwrite previous actions parameters when defined. To remove a parameter, use the no form of the action command with the specified parameter.

The **no** form of this command removes the specified **action** statement. The filter entry is considered incomplete and therefore rendered inactive without the **action** command.

Parameters

drop

Specifies that packets matching the entry criteria will be dropped.

forward

Specifies that packets matching the entry criteria will be forwarded.

If neither drop nor forward is specified, the filter action is no-op and the filter entry is inactive.

match

Syntax

match [**frame-type** *keyword*]

no match

Context

config>filter>mac-filter>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures match criteria for the filter entry and specifies an Ethernet frame type for the entry. 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 then all criteria must be satisfied (AND function) before the action associated with the match will be executed.

A **match** context may 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 *keyword*

Specifies an Ethernet frame type to be used for the MAC filter match criteria.

Default ethernet_II

4.8.2.1.7 IP filter match criteria

dscp

Syntax

dscp *dscp-name*

no dscp

Context

config>filter>ip-filter>entry>match

config>filter>ipv6-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a DiffServ Code Point (DSCP) name to be used as an IP filter 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 may 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 | 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

dst-ip

Syntax

dst-ip {*ip-address/mask* | *ip-address ipv4-address-mask*}

no dst-ip

Context

config>filter>ip-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a destination IPv4 address range to be used as an IP filter match criterion.

To match on the destination IPv4 address, specify the address and its associated mask, such as 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 may also be used.

The **no** form of this command removes the destination IPv4 address match criterion.

Default

none

Parameters

ip-address

Specifies the IPv4 prefix for the IP match criterion in dotted decimal notation.

Values a.b.c.d

mask

Specifies the subnet mask length expressed as a decimal integer.

Values 0 to 32

ipv4-address-mask

Specifies any mask expressed in dotted quad notation.

Values 0 to 255

dst-ip

Syntax

dst-ip {*ipv6-address*|*prefix-length*}

no dst-ip

Context

config>filter>ipv6-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a destination IPv6 address range to be used as an IP filter match criterion.

To match on the destination IPv6 address, specify the address and its associated mask.

The **no** form of this command removes the destination IPv6 address match criterion.

Default

none

Parameters

ipv6-address

Specifies the IPv6 prefix for the IP match criterion in hex digits.

Values x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:d.d.d.d

x - 0 to FFFF (hexadecimal)

d - 0 to 255 (decimal)

prefix-length

Specifies the IPv6 prefix length for the IPv6 address as a decimal integer.

Values 1 to 128

dst-port

Syntax

dst-port {eq} *dst-port-number*

dst-port range *dst-port-number1 dst-port-number2*

no dst-port

Context

config>filter>ip-filter>entry>match

config>filter>ipv6-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a destination TCP or UDP port number for an IP filter match criterion. An entry containing L4 match criteria will not match non-initial (2nd, 3rd, and so on) fragments of a fragmented packet since only the first fragment contains the L4 information.



Note:

A limited number of unique port range values are available for use (see the **tools>dump>system-resources** command for the current usage). The port range entries are shared among the QoS classification and ACL entries, and among source and destination port range values. The available entries in the hardware port range resource pool is allocated on a first come, first served basis. The following rules apply to the allocation of these port range pool entries:

- One port range hardware entry is required to match a unique combination of port range values configured by the user. Two port range values (for example, port range 1024 to 2048 and port range 3050 to 3055) are considered to be unique if their range 1 and range 2 values do not match. Each unique port range value consumes one entry each in hardware.
- Each unique port range entry requires two entries in the hardware table if it is used for both a source and destination match. In other words, a unique port range value requires one entry for matching a source port and another entry for matching a destination port.
- If one or more QoS or ACL policy use the same unique port range value for either the source port match or destination port match but not both, only a single entry in the hardware table is required. For example, if port range 1000 to 2000 is used in both the access ingress policy and filter policy to match on source port, it requires one entry in the hardware port range table.
- If one or more QoS and or ACL policy use the same unique port range value for both source port match and destination port, two entries in the hardware table are required. For example, if port range 1000 to 2000 is used in both access ingress policy and filter policy to match on

both the source port and destination port, it requires two entries in the hardware port range table.

The **no** form of this command removes the destination port match criterion.

Parameters

eq

Specifies the operator to use relative to *dst-port-number* for specifying the port number match criteria. The **eq** keyword specifies that *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.

Values 1 to 65535

range *dst-port-number1*

Specifies the first destination port number in the range to match a range of ports. This parameter is only available on the 7210 SAS-Mxp.

Values 1 to 65535

range *dst-port-number2*

Specifies the last destination port number in the range to match a range of ports. The second number must be greater than the first number in the range. This parameter is only available on the 7210 SAS-Mxp.

Values 1 to 65535

fragment

Syntax

fragment {**true** | **false**}

no fragment

Context

config>filter>ip-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures fragmented or non-fragmented IP packets as an IP filter match criterion. An entry containing L4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the L4 information.

The **no** form of this command removes the match criterion.

Default

no fragment

Parameters

true

Specifies 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

Specifies 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.

icmp-code

Syntax

icmp-code *icmp-code*

no icmp-code

Context

config>filter>ip-filter>entry>match

config>filter>ipv6-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures matching on ICMP code field in the ICMP header of an IP packet as a filter match criterion. An entry containing L4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the L4 information.

This option is only meaningful if the protocol match criteria specifies ICMP (1).

The **no** form of this command removes the criterion from the match entry.

Default

no icmp-code

Parameters

icmp-code

Specifies the ICMP code values that must be present to match.

Values 0 to 255

icmp-type

Syntax

icmp-type *icmp-type*

no icmp-type

Context

config>filter>ip-filter>entry>match

config>filter>ipv6-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures matching on the ICMP type field in the ICMP header of an IP or packet as a filter match criterion. An entry containing L4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the L4 information.

This option is only meaningful if the protocol match criteria specifies ICMP (1).

The **no** form of this command removes the criterion from the match entry.

Default

no icmp-type

Parameters

icmp-type

Specifies the ICMP type values that must be present to match.

Values 0 to 25A

option-present

Syntax

option-present {true | false}

no option-present

Context

config>filter>ip-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures matching packets that contain the option field or have an option field of zero in the IP header as an IP filter match criterion.

The **no** form of this command removes the checking of the option field in the IP header as a match criterion.

Parameters

true

Specifies matching on all IP packets that contain the option field in the header. A match will occur for all packets that have the option field present. An option field of zero is considered as no option present.

false

Specifies matching on IP packets that do not have any option field present in the IP header. (an option field of zero). An option field of zero is considered as no option present.

src-ip

Syntax

src-ip {*ip-address/mask* | *ip-address ipv4-address-mask*}

no src-ip

Context

config>filter>ip-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a source IPv4 address range to be used as an IP filter match criterion.

To match on the source IPv4 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 may also be used.

The **no** form of this command removes the source IPv4 address match criterion.

Default

no src-ip

Parameters

ip-address

Specifies the IPv4 prefix for the IP match criterion in dotted decimal notation.

Values a.b.c.d

mask

Specifies the subnet mask length, expressed as a decimal integer.

Values 0 to 32

ipv4-address-mask

Specifies any mask, expressed in dotted quad notation.

Values 0 to 255

src-ip

Syntax

src-ip {*ipv6-address/prefix-length*}

no src-ip

Context

config>filter>ipv6-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a source IPv6 address range to be used as an IP filter match criterion.

To match on the source IPv6 address, specify the address and its associated mask.

If the filter is created to match 64-bit address, the IPv6 address specified for the match must contain only the first 64-bits (that is, the first four 16-bit groups of the IPv6 address).

The **no** form of this command removes the source IPv6 address match criterion.

Default

no src-ip

Parameters

ipv6-address

Specifies the IPv6 prefix for the IP match criterion in hex digits.

Values x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:d.d.d.d

x - 0 to FFFF (hexadecimal)

d - 0 to 255 (decimal)

prefix-length

Specifies the IPv6 prefix length for the IPv6 address as a decimal integer.

Values 1 to 128

src-port

Syntax

src-port {eq} *src-port-number*

src-port range *src-port-number1* *src-port-number2*

no src-port

Context

config>filter>ip-filter>entry>match

config>filter>ipv6-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a source TCP or UDP port number for an IP filter match criterion. An entry containing L4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the L4 information.



Note:

A limited number of unique port range values are available for use (see the **tools>dump>system-resources** command for the current usage). The port range entries are shared among the QoS classification and ACL entries, and among source and destination port range values. The available entries in the hardware port range resource pool is allocated on a first come, first served basis. The following rules apply to the allocation of these port range pool entries:

- One port range hardware entry is required to match a unique combination of port range values configured by the user. Two port range values (for example, port range 1024 to 2048 and port range 3050 to 3055) are considered to be unique if their range 1 and range 2 values do not match. Each unique port range value consumes one entry each in hardware.
- Each unique port range entry requires two entries in the hardware table if it is used for both a source and destination match. In other words, a unique port range value requires one entry for matching a source port and another entry for matching a destination port.
- If one or more QoS or ACL policy use the same unique port range value for either the source port match or destination port match but not both, only a single entry in the hardware table is required. For example, if port range 1000 to 2000 is used in both the access ingress policy and filter policy to match on source port, it requires one entry in the hardware port range table.
- If one or more QoS and or ACL policy use the same unique port range value for both source port match and destination port, two entries in the hardware table are required. For example, if port range 1000 to 2000 is used in both access ingress policy and filter policy to match on both the source port and destination port, it requires two entries in the hardware port range table.

The **no** form of this command removes the source port match criterion.

Default

no src-port

Parameters

eq

Specifies the operator to use relative to *src-port-number* for specifying the port number match criteria. The **eq** keyword 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.

Values 0 to 65535

range *src-port-number1*

Specifies the first source port number in the range to match a range of ports. This parameter is only available on the 7210 SAS-Mxp.

Values 1 to 65535

range *src-port-number2*

Specifies the last source port number in the range to match a range of ports. The second number must be greater than the first number in the range. This parameter is only available on the 7210 SAS-Mxp.

Values 1 to 65535

tcp-ack

Syntax

tcp-ack {true | false}

no tcp-ack

Context

config>filter>ip-filter>entry>match

config>filter>ipv6-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures matching on the ACK bit being set or reset in the control bits of the TCP header of an IP packet as an IP filter match criterion. An entry containing L4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the L4 information.

The **no** form of this command removes the criterion from the match entry.

Default

no tcp-ack

Parameters

true

Specifies matching on IP packets that have the ACK bit set in the control bits of the TCP header of an IP packet.

false

Specifies matching on IP packets that do not have the ACK bit set in the control bits of the TCP header of the IP packet.

tcp-syn

Syntax

tcp-syn {true | false}

no tcp-syn

Context

config>filter>ip-filter>entry>match

config>filter>ipv6-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures matching on the SYN bit being set or reset in the control bits of the TCP header of an IP packet as an IP filter match criterion. An entry containing L4 match criteria will not match non-initial (2nd, 3rd, and so on) fragments of a fragmented packet since only the first fragment contains the L4 information.

The SYN bit is normally set when the source of the packet needs to initiate a TCP session with the specified destination IP address.

The **no** form of this command removes the criterion from the match entry.

Default

no tcp-syn

Parameters

true

Specifies matching on IP packets that have the SYN bit set in the control bits of the TCP header.

false

Specifies matching on IP packets that do not have the SYN bit set in the control bits of the TCP header.

4.8.2.1.8 MAC filter match criteria

dot1p

Syntax

dot1p *ip-value* [*mask*]

no dot1p

Context

config>filter>mac-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an IEEE 802.1p value or range to be used as a MAC filter match criterion.

When a frame is missing the 802.1p bits, specifying an dot1p match criterion will fail for the frame and result in a non-match for the MAC filter entry.

The **no** form of this command removes the criterion from the match entry.

The MAC filter applied on the SAP egress can match the details of the packet-on-the-wire. For example, a QinQ packet came in on a null SAP and egressing on a dot1p-encapsulated port, the packet-on-the-wire will have three tags. Now, the etype=0x8100 and dot1p will equal the outer VLAN tag dot1p. This Etype and dot1p can be configured on the egress filter to match this packet.

Default

no dot1p

Parameters

ip-value

Specifies the IEEE 802.1p value, in decimal notation

Values 0 to 7

mask

Specifies a 3-bit mask.

This 3-bit mask can be specified using the following formats:

Table 45: 3-bit mask format

Format style	Format syntax	Example
Decimal	D	4
Hexadecimal	0xH	0x4
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 1 to 7

Default 7

dst-mac

Syntax

dst-mac *ieee-address* [*mask*]

no dst-mac

Context

config>filter>mac-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Configures a destination MAC address or range to be used as a MAC filter match criterion.

The **no** form of this command removes the destination mac address as the match criterion.

Default

no dst-mac

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.

This 48-bit mask can be configured using the following formats:

Table 46: 48-bit mask format

Format style	Format syntax	Example
Decimal	DDDDDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHHH	0xFFFFFFFF000000
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: 0003FA000000 0xFFFFFFFF000000

Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

Default 0xFFFFFFFFFFFF (exact match)

etype

Syntax

etype *ethernet-type*

no etype

Context

config>filter>mac-filter>entry>match

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an Ethernet type II Ethertype value for use 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. [Table 43: MAC match criteria exclusivity rules](#) describes fields that are exclusive based on the frame format.

The data plane processes a maximum of two VLAN tags in a received packet. The Ethertype used in the MAC matching criteria for ACLs is the Ethertype that is found in the packet after processing single-tagged frames, double-tagged frames, and no-tag frames

The packet is considered to have no tags if at least one of the following criteria is true.

- The packet is a null-tagged frame
- The packet is a priority-tagged frame

- The outermost Ethertype does not match the default Ethertype (0x8100)
- The outermost Ethertype does not match the configured dot1q-etype on Dot1q encapsulated ports
- The outermost Ethertype does not match the configured qinq-etype on QinQ encapsulated ports

The packet is considered to have a single tag if at least one of the following criteria is true.

- The outermost Ethertype matches the default Ethertype (0x8100)
- The outermost Ethertype matches the configured dot1q-etype on Dot1q encapsulated ports
- The outermost Ethertype matches the configured qinq-etype on QinQ encapsulated ports

The packet is considered to have double tags if at least one of the following criteria is true.

- The outermost Ethertype matches the default Ethernet type (0x8100)
- The configured dot1q-etype on Dot1q encapsulated ports and the immediately following Ethertype match the default Ethertype (0x8100)
- The configured qinq-etype on QinQ encapsulated ports and the immediately following Ethertype match the default Ethertype (0x8100)

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

src-mac

Syntax

src-mac *ieee-address* [*ieee-address-mask*]

no src-mac

Context

config>filter>mac-filter>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

This 48-bit mask can be configured using the following formats:

Table 47: 48-bit mask format

Format style	Format syntax	Example
Decimal	DDDDDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHHHH	0xFFFFFFFF000000
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, the entry should be specified as: 003FA000000 0xFFFFFFFF000000

Values 0x0000000000000000 to 0xFFFFFFFFFFFFFFF

Default 0xFFFFFFFFFFFFFFF (exact match)

4.8.2.1.9 Policy and entry maintenance commands

copy

Syntax

copy {*ip-filter* | *ipv6-filter* | *mac-filter*} *source-filter-id* *dest-filter-id* *dest-filter-id* [*overwrite*]

Context

config>filter

Platforms

Supported on all 7210 SAS platforms as described in this document

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 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 occurs if the destination policy ID exists.

Parameters

ip-filter

Specifies that the *source-filter-id* and the *dest-filter-id* are IP filter IDs.

ipv6-filter

Specifies that the *source-filter-id* and the *dest-filter-id* are IPv6 filter IDs.

mac-filter

Specifies that the *source-filter-id* and the *dest-filter-id* are MAC filter IDs.

source-filter-id

Specifies the *source-filter-id*, which 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** or **mac-filter**).

dest-filter-id

Specifies the destination filter policy where the copy command attempts to copy. If the **overwrite** keyword does not follow, the filter policy ID cannot already exist within the system for the filter type the copy command is issued for. If the **overwrite** keyword is present, the destination policy ID may or may not exist.

overwrite

Specifies that the destination filter ID may exist. If it does, the existing destination filter ID will be overwritten with the contents of the source filter ID. If the destination filter ID exists, either **overwrite** must be specified or an error message is returned. If **overwrite** is specified, the function of copying from source to destination occurs in a “break before make” manner, so must be handled with care.

filter-name

Syntax

filter-name *filter-name*

Context

config>filter>ip-filter

config>filter>ipv6-filter

config>filter>mac-filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures filter-name attribute of a specific filter. filter-name, when configured, can be used instead of filter ID to reference the specific policy in the CLI.

Default

no filter-name

Parameters

filter-name

Specifies a string of up to 64 characters uniquely identifying this filter policy.

```
renum
```

Syntax

```
renum old-entry-id new-entry-id
```

Context

```
config>filter>ip-filter
```

```
config>filter>ipv6-filter
```

```
config>filter>mac-filter
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command renumbers existing MAC or IP 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.

Values 1 to 65535

new-entry-id

Specifies the new entry-number to be assigned to the old entry.

Values 1 to 65535

type

Syntax

type *filter-type*

Context

config>filter>mac-filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the type of mac-filter as normal, ISID or VID types.

Default

normal

Parameters

filter-type

Specifies which type of entries this MAC filter can contain.

- Values**
- normal — Regular match criteria are allowed; ISID or VID filter match criteria not allowed.
 - isid — Only ISID match criteria are allowed.
 - vid — Only VID match criteria are allowed on ethernet_II frame types.

4.8.2.2 Show commands

download-failed

Syntax

download-failed

Context

show>filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays all filter entries for which the download has failed.

Output

The following output is an example of failed download information, and [Table 48: Output fields: filter download-failed](#) describes the output fields.

Sample output

```
A:ALA-48# show filter download-failed
=====
Filter entries for which download failed
=====
Filter-type  Filter-Id  Filter-Entry
-----
ip           1          10
=====
A:ALA-48#
```

Table 48: Output fields: filter download-failed

Label	Description
Filter-type	Displays the filter type
Filter-ID	Displays the ID of the filter
Filter-Entry	Displays the entry number of the filter

```
ip
```

Syntax

ip *ip-filter-id* [**association** | **counters**]

ip *ip-filter-id* **entry** *entry-id* [**counters**]

Context

show>filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays IP filter information.

Parameters

ip-filter-id

Displays detailed information for the specified filter ID and its filter entries.

Values 1 to 65535

entry entry-id

Displays information about the specified filter entry ID for the specified filter ID only.

Values 1 to 65535

associations

Displays information about where the filter policy ID is applied to the detailed filter policy ID output.

counters

Displays counter information for the specified filter ID. Egress counters count the packets without Layer 2 encapsulation. Ingress counters count the packets with Layer 2 encapsulation.

type entry-type

Displays information about the specified filter ID for the specified *entry-type* only

Output

The following outputs are examples of IP filter information, and the associated tables describe the output fields.

- IP Filter: [Sample Output](#) , [Table 49: Output fields: filter IP](#)
- IP Filter with Filter ID Specified: [Sample Output](#) , [Sample Output on the 7210 SAS-Mxp](#), [Table 50: Output fields: filter with filter ID specified](#)
- IP Filter with Time-range Specified: [Sample Output](#)
- IP Filter Associations: [Sample Output](#) , [Table 51: Output fields: filter IP associations](#)
- IP Filter Counters: [Table 52: Output fields: filter counters](#)

Sample Output

```
A:ALA-49# show filter ip
=====
IP Filters
=====
Filter-Id Scope    Applied Description
-----
1         Template Yes
3         Template Yes
6         Template Yes
10        Template No
11        Template No
-----
Num IP filters: 5
=====
A:ALA-49#
*A:Dut-C>config>filter# show filter ip
=====
IP Filters                                     Total:    2
=====
Filter-Id  Scope    Applied Description
-----
10001     Template Yes
fSpec-1   Template Yes    BGP FlowSpec filter for the Base router
```

```
-----
Num IP filters: 2
=====
*A:Dut-C>config>filter#
```

Table 49: Output fields: filter IP

Label	Description
Filter Id	The IP filter ID.
Scope	Template The filter policy is of type template. Exclusive The filter policy is of type exclusive.
Applied	No The filter policy ID has not been applied. Yes The filter policy ID has been applied.
Description	The IP filter policy description.

Sample Output

```
A:ALA-49>config>filter# show filter ip 3
=====
IP Filter
=====
Filter Id      : 3                      Applied       : Yes
Scope         : Template                Def. Action   : Drop
Entries       : 1
-----
Filter Match Criteria : IP
-----
Entry         : 10
Src. IP       : 10.1.1.1/24             Src. Port     : None
Dest. IP      : 0.0.0.0/0               Dest. Port    : None
Protocol      : 2                       Dscp         : Undefined
ICMP Type     : Undefined                ICMP Code     : Undefined
TCP-syn       : Off                      TCP-ack       : Off
Match action  : Drop
Ing. Matches  : 0                       Egr. Matches  : 0
=====
A:ALA-49>config>filter#

*A:Dut-C>config>filter# show filter ip fSpec-1 associations
=====
IP Filter
=====
Filter Id      : fSpec-1                 Applied       : Yes
Scope         : Template                Def. Action   : Forward
Radius Ins Pt : n/a
CrCtl. Ins Pt : n/a
Entries       : 2 (insert By Bgp)
Description   : BGP FlowSpec filter for the Base router
```

```

-----
Filter Association : IP
-----
Service Id   : 1                               Type       : IES
- SAP       1/1/3:1.1 (merged in ip-fltr 10001)
=====
*A:Dut-C>config>filter#

*A:Dut-C>config>filter# show filter ip 10001
=====
IP Filter
=====
Filter Id    : 10001                          Applied     : Yes
Scope       : Template                        Def. Action  : Drop
Radius Ins Pt: n/a
CrCtl. Ins Pt: n/a
Entries      : 1
BGP Entries  : 2
Description  : (Not Specified)
-----
Filter Match Criteria : IP
-----
Entry       : 1
Description  : (Not Specified)
Log Id      : n/a
Src. IP     : 0.0.0.0/0                        Src. Port   : None
Dest. IP    : 0.0.0.0/0                        Dest. Port  : None
Protocol    : 6                               Dscp        : Undefined
ICMP Type   : Undefined                       ICMP Code   : Undefined
Fragment    : Off                             Option-present : Off
Sampling    : Off                             Int. Sampling : On
IP-Option   : 0/0                             Multiple Option: Off
TCP-syn     : Off                             TCP-ack     : Off
Match action : Forward
Next Hop    : Not Specified
Ing. Matches : 0 pkts
Egr. Matches : 0 pkts

Entry       : fSpec-1-32767 - inserted by BGP FLOWSpec
Description  : (Not Specified)
Log Id      : n/a
Src. IP     : 0.0.0.0/0                        Src. Port   : None
Dest. IP    : 0.0.0.0/0                        Dest. Port  : None
Protocol    : 6                               Dscp        : Undefined
ICMP Type   : Undefined                       ICMP Code   : Undefined
Fragment    : Off                             Option-present : Off
Sampling    : Off                             Int. Sampling : On
IP-Option   : 0/0                             Multiple Option: Off
TCP-syn     : Off                             TCP-ack     : Off
Match action : Drop
Ing. Matches : 0 pkts
Egr. Matches : 0 pkts

Entry       : fSpec-1-49151 - inserted by BGP FLOWSpec
Description  : (Not Specified)
Log Id      : n/a
Src. IP     : 0.0.0.0/0                        Src. Port   : None
Dest. IP    : 0.0.0.0/0                        Dest. Port  : None
Protocol    : 17                              Dscp        : Undefined
ICMP Type   : Undefined                       ICMP Code   : Undefined
Fragment    : Off                             Option-present : Off
Sampling    : Off                             Int. Sampling : On
IP-Option   : 0/0                             Multiple Option: Off
    
```

```
TCP-syn      : Off                TCP-ack      : Off
Match action : Drop
Ing. Matches : 0 pkts
Egr. Matches : 0 pkts
```

```
=====
*A:Dut-C>config>filter#
```

Sample Output on the 7210 SAS-Mxp

```
=====
IP Filter
=====
Filter Id      : 1                Applied      : Yes
Scope         : Template         Def. Action  : Drop
System filter: Unchained
RadSh. Ins Pt: n/a
PccRl. Ins Pt: n/a
Entries       : 16
Description   : (Not Specified)
-----
Filter Match Criteria : IP
-----
Entry         : 1
Description   : Entry1
Src. IP      : 0.0.0.0/0
Src. Port    : 10..15
Dest. IP     : 0.0.0.0/0
Dest. Port   : 10..15
Protocol     : 6                  Dscp        : Undefined
ICMP Type    : Undefined         ICMP Code    : Undefined
Fragment     : Off
Sampling     : Off                Int. Sampling : On
TCP-syn      : Off                TCP-ack      : Off
Option-pres  : Off
Egress PBR   : Disabled
Match action : Forward
Ing. Matches : 0 pkts
Egr. Matches : 0 pkts

Entry        : 2 (Inactive)
Description  : Entry2
Src. IP     : 0.0.0.0/0
Src. Port   : 20..25
Dest. IP    : 0.0.0.0/0
Dest. Port  : 20..25
Protocol    : 6                  Dscp        : Undefined
ICMP Type   : Undefined         ICMP Code    : Undefined
Fragment    : Off
Sampling    : Off                Int. Sampling : On
TCP-syn     : Off                TCP-ack      : Off
Option-pres : Off
Egress PBR  : Disabled
Match action : Drop
Ing. Matches : 0 pkts
Egr. Matches : 0 pkts
```

Table 50: Output fields: filter with filter ID specified

Label	Description
Filter Id	The IP filter policy ID.
Scope	Template — Indicates the filter policy is of type template.
	Exclusive — Indicates the filter policy is of type exclusive.
Entries	The number of entries configured in this filter ID.
Description	The IP filter policy description.
Applied	No — Indicates the filter policy ID has not been applied.
	Yes — Indicates the filter policy ID has been applied.
Def. Action	Forward — The default action for the filter ID for packets that do not match the filter entries is to forward.
	Drop — The default action for the filter ID for packets that do not match the filter entries is to drop.
Filter Match Criteria	IP — Indicates the filter is an IP filter policy.
Entry	The filter ID filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified.
ICMP Type	The ICMP type match criterion. Undefined indicates no ICMP type specified.
Fragment	False — Configures a match on all non-fragmented IP packets.
	True — Configures a match on all fragmented IP packets.
	Off — Fragments are not a matching criteria. All fragments and non-fragments implicitly match.
TCP-syn	False — Configures a match on packets with the SYN flag set to false.
	True — Configured a match on packets with the SYN flag set to true.
	Off — The state of the TCP SYN flag is not considered as part of the match criteria.
Match action	Default — The filter does not have an explicit forward or drop match action specified. If the filter entry ID indicates the entry is Inactive, the filter entry is incomplete, no action was specified.
	Drop — Drop packets matching the filter entry.

Label	Description
	Forward — The explicit action to perform is forwarding of the packet.
Ing. Matches	The number of ingress filter matches or hits for the filter entry.
Src. IP	The source IP address.
Src. Port	The source TCP or UDP port number.
Dest. IP	The destination IP address.
Dest. Port	The destination TCP or UDP port number.
Dscp	The DiffServ Code Point (DSCP) name.
ICMP Code	The ICMP code field in the ICMP header of an IP packet.
Option-present	Off — Does not search for packets that contain the option field or have an option field of zero.
	On — Matches packets that contain the option field or have an option field of zero be used as IP filter match criteria.
TCP-ack	False — Configures a match on packets with the ACK flag set to false.
	True — Configures a match on packets with the ACK flag set to true.
	Off — The state of the TCP ACK flag is not considered as part of the match criteria.
Egr. Matches	The number of egress filter matches or hits for the filter entry.

Sample Output

```
A:ALA-49# show filter ip 10
=====
IP Filter
=====
Filter Id      : 10                               Applied       : No
Scope         : Template                       Def. Action   : Drop
Entries       : 2
-----
Filter Match Criteria : IP
-----
Entry         : 1010
time-range    : day                            Cur. Status   : Inactive
Src. IP       : 0.0.0.0/0                       Src. Port     : None
Dest. IP      : 10.10.100.1/24                  Dest. Port    : None
Protocol      : Undefined                       Dscp         : Undefined
ICMP Type     : Undefined                       ICMP Code     : Undefined
Fragment      : Off                            Option-present : Off
TCP-syn       : Off                            TCP-ack       : Off
Match action  : Forward
Ing. Matches  : 0                               Egr. Matches  : 0
```

```

Entry       : 1020
time-range  : night
Src. IP     : 0.0.0.0/0
Dest. IP    : 10.10.1.1/16
Protocol    : Undefined
ICMP Type   : Undefined
Fragment    : Off
TCP-syn     : Off
Match action : Forward
Ing. Matches : 0
Cur. Status : Active
Src. Port    : None
Dest. Port   : None
Dscp        : Undefined
ICMP Code    : Undefined
Option-present : Off
TCP-ack     : Off
Egr. Matches : 0
=====
A:ALA-49#
    
```

Sample Output

```

A:ALA-49# show filter ip 1 associations
=====
IP Filter
=====
Filter Id    : 1
Scope       : Template
Entries     : 1
-----
Filter Association : IP
-----
Service Id   : 1001
- SAP       1/1/1:1001 (Ingress)
Service Id   : 2000
- SAP       1/1/1:2000 (Ingress)
Type        : VPLS
Type        : Epipe
=====
A:ALA-49#
    
```

```

A:ALA-49# show filter ip 160 associations
=====
IP Filter
=====
Filter Id    : 160
Scope       : Template
Entries     : 0
-----
Filter Association : IP
-----
Tod-suite "english_suite"
- ingress, time-range "day" (priority 5)
=====
A:ALA-49#
    
```

Table 51: Output fields: filter IP associations

Label	Description
Filter Id	The IP filter policy ID.
Scope	Template — The filter policy is of type Template.
	Exclusive — The filter policy is of type Exclusive.
Entries	The number of entries configured in this filter ID.

Label	Description
Applied	Yes — The filter policy ID has been applied.
	No — The filter policy ID has not been applied.
Def. Action	Forward — The default action for the filter ID for packets that do not match the filter entries is to forward.
	Drop — The default action for the filter ID for packets that do not match the filter entries is to drop.
Service Id	The service ID on which the filter policy ID is applied.
SAP	The Service Access Point on which the filter policy ID is applied.
(Ingress)	The filter policy ID is applied as an ingress filter policy on the interface.
(Egress)	The filter policy ID is applied as an egress filter policy on the interface.
Type	The type of service of the service ID.

Show Filter Counters

Table 52: Output fields: filter counters

Label	Description
IP Filter Filter Id	The IP filter policy ID.
Scope	Template — The filter policy is of type Template.
	Exclusive — The filter policy is of type Exclusive.
Applied	No — The filter policy ID has not been applied.
	Yes — The filter policy ID has been applied.
Def. Action	Forward — The default action for the filter ID for packets that do not match the filter entries is to forward.
	Drop — The default action for the filter ID for packets that do not match the filter entries is to drop.
Filter Match Criteria	IP — Indicates the filter is an IP filter policy.
Entry	The filter ID filter entry ID. If the filter entry ID indicates the entry is Inactive, then the filter entry is incomplete as no action has been specified.
Ing. Matches	The number of ingress filter matches or hits for the filter entry. Ingress counters count the packets with Layer 2 encapsulation.

Label	Description
Egr. Matches	The number of egress filter matches or hits for the filter entry. Egress counters count the packets without Layer 2 encapsulation.

ipv6

Syntax

ipv6 {*ipv6-filter-id* [**entry** *entry-id*] [**association** | **counters**]}

Context

show>filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays IPv6 filter information.

Parameters

ipv6-filter-id

Displays detailed information for the specified IPv6 filter ID and filter entries.

Values 1 to 65535

entry entry-id

Displays information about the specified IPv6 filter entry ID for the specified filter ID.

Values 1 to 9999

associations

Displays information about where the IPv6 filter policy ID is applied to the detailed filter policy ID output.

counters

Displays counter information for the specified IPv6 filter ID.

Egress counters count the packets without Layer 2 encapsulation. Ingress counters count the packets with Layer 2 encapsulation.

Output

The following outputs are examples of IP filter information, and the associated tables describe the output fields.

- [Sample Output, Table 53: Output fields: IP filter \(no filter-id specified\)](#)

- [Sample Output for IPv6 Filter with Filter ID Specified](#), Table 54: Output fields: filter (with filter-id specified)
- [Sample Output for IPv6 Filter Associations](#), Table 55: Output fields: filter associations
- [Sample Output for IPv6 Filter Counters](#), Table 56: Output fields: filter counters

Sample Output

```
*A:7210SAS>show>filter# ipv6

=====
IPv6 Filters                                     Total:    1
=====
Filter-Id Scope   Applied Description
-----
1       Template Yes
-----
Num IPv6 filters: 1
=====
*A:7210SAS>show>filter#
```

Table 53: Output fields: IP filter (no filter-id specified)

Label	Description
Filter Id	The IP filter ID.
Scope Template	The filter policy is of type template.
Exclusive	The filter policy is of type exclusive.
Applied	No — The filter policy ID has not been applied.
	Yes — The filter policy ID has been applied.
Description	The IP filter policy description.

Sample Output for IPv6 Filter with Filter ID Specified

```
*A:7210SAS>show>filter# ipv6 1

=====
IPv6 Filter
=====
Filter Id      : 1                      Applied       : Yes
Scope         : Template                Def. Action   : Drop
Entries       : 2
Description    : (Not Specified)
-----
Filter Match Criteria : IPv6
-----
Entry         : 1
Description   : Test
Src. IP       : 2001:db8::1/128          Src. Port     : None
Dest. IP      : 2001:db8::/0             Dest. Port    : None
Next Header   : Undefined                Dscp         : Undefined
ICMP Type     : Undefined                ICMP Code     : Undefined
TCP-syn       : Off                      TCP-ack       : Off
Match action  : Forward
```

```

Ing. Matches : 0 pkts
Egr. Matches : 0 pkts

Entry       : 2
Description : (Not Specified)
Src. IP     : ::/0
Dest. IP    : 1:2::1AFC/128
Next Header : Undefined
ICMP Type   : Undefined
TCP-syn     : Off
Match action : Drop
Ing. Matches : 819 pkts
Egr. Matches : 0 pkts

Src. Port   : None
Dest. Port  : None
Dscp        : Undefined
ICMP Code   : Undefined
TCP-ack     : Off

=====
*A:7210SAS>show>filter#
    
```

Table 54: Output fields: filter (with filter-id specified)

Label	Description
Filter Id	The IP filter policy ID.
Scope	Template — The filter policy is of type template.
	Exclusive — The filter policy is of type exclusive.
Entries	The number of entries configured in this filter ID.
Description	The IP filter policy description.
Applied	No — The filter policy ID has not been applied.
	Yes — The filter policy ID has been applied.
Def. Action	Forward — The default action for the filter ID for packets that do not match the filter entries is to forward.
	Drop — The default action for the filter ID for packets that do not match the filter entries is to drop.
Filter Match Criteria	IP — Indicates the filter is an IP filter policy.
Entry	The filter ID filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified.
Src. IP	The source IP address and mask match criterion. 0.0.0.0/0 indicates no criterion specified for the filter entry.
Dest. IP	The destination IP address and mask match criterion. 0.0.0.0/0 indicates no criterion specified for the filter entry.
ICMP Type	The ICMP type match criterion. Undefined indicates no ICMP type specified.

Label	Description
IP-Option	Specifies matching packets with a specific IP option or a range of IP options in the IP header for IP filter match criteria.
TCP-syn	False — Configures a match on packets with the SYN flag set to false.
	True — Configured a match on packets with the SYN flag set to true.
	Off — The state of the TCP SYN flag is not considered as part of the match criteria.
Match action	Default — The filter does not have an explicit forward or drop match action specified. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified.
	Drop — Drop packets matching the filter entry.
	Forward — The explicit action to perform is forwarding of the packet. If the action is Forward, then, if configured, the next-hop information should be displayed, including Nexthop: <IP address>, Indirect: <IP address> or Interface: <IP interface name>.
Ing. Matches	The number of ingress filter matches or hits for the filter entry.
Src. Port	The source TCP or UDP port number or port range.
Dest. Port	The destination TCP or UDP port number or port range.
Dscp	The DiffServ Code Point (DSCP) name.
ICMP Code	The ICMP code field in the ICMP header of an IP packet.
TCP-ack	False — Configures a match on packets with the ACK flag set to false.
	True — Configured a match on packets with the ACK flag set to true.
	Off — The state of the TCP ACK flag is not considered as part of the match criteria
Ing. Matches	The number of ingress filter matches or hits for the filter entry.
Egr. Matches	The number of egress filter matches or hits for the filter entry.

Sample Output for IPv6 Filter Associations

```
*A:7210SAS>show>filter# ipv6 1 associations
```

=====

```

IPv6 Filter
=====
Filter Id      : 1                               Applied      : Yes
Scope         : Template                       Def. Action  : Drop
Entries       : 2
Description   : (Not Specified)
-----
Filter Association : IPv6
-----
Service Id    : 1                               Type         : Epipe
- SAP        1/1/1:1 (Ingress)
Service Id    : 2                               Type         : VPLS
- SAP        1/1/1:2 (Ingress)
- SAP        1/1/1:3 (Ingress)
=====
*A:7210SAS>show>filter#
    
```

Table 55: Output fields: filter associations

Label	Description
Filter Id	The IPv6 filter policy ID.
Scope	Template — The filter policy is of type Template.
	Exclusive — The filter policy is of type Exclusive.
Entries	The number of entries configured in this filter ID.
Applied	No — The filter policy ID has not been applied.
	Yes — The filter policy ID has been applied.
Def. Action	Forward — The default action for the filter ID for packets that do not match the filter entries is to forward.
	Drop — The default action for the filter ID for packets that do not match the filter entries is to drop.
Description	The IP filter policy description.
Service Id	The service ID on which the filter policy ID has been applied.
SAP	The Service Access Point on which the filter policy ID is applied.
	Ingress — The filter policy ID is applied as an ingress filter policy on the interface.
	Egress — The filter policy ID is applied as an egress filter policy on the interface.
Type	The type of service of the service ID.

Sample Output for IPv6 Filter Counters

```
*A:7210SAS>show>filter# ipv6 1 counters

=====
IPv6 Filter
=====
Filter Id      : 1                               Applied       : Yes
Scope         : Template                       Def. Action   : Drop
Entries       : 2
Description    : (Not Specified)
-----
Filter Match Criteria : IPv6
-----
Entry         : 1
Ing. Matches  : 0 pkts
Egr. Matches  : 0 pkts

Entry        : 2
Ing. Matches : 819 pkts
Egr. Matches : 0 pkts

=====
*A:7210SAS>show>filter#
```

Table 56: Output fields: filter counters

Label	Description
Filter Id	The IPv6 filter policy ID.
Scope	Template — The filter policy is of type Template.
	Exclusive — The filter policy is of type Exclusive.
Entries	The number of entries configured in this filter ID.
Applied	No — The filter policy ID has not been applied.
	Yes — The filter policy ID has been applied.
Def. Action	Forward — The default action for the filter ID for packets that do not match the filter entries is to forward.
	Drop — The default action for the filter ID for packets that do not match the filter entries is to drop.
Description	The IP filter policy description.
Entry	The filter ID filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified.
Ing. Matches	The number of ingress filter matches or hits for the filter entry.
Egr. Matches	The number of egress filter matches or hits for the filter entry.

Label	Description
	Egress counters count the packets without Layer 2 encapsulation. Ingress counters count the packets with Layer 2 encapsulation.

mac

Syntax

mac [*mac-filter-id* [**associations** | **counters**] [**entry** *entry-id*]]

Context

show>filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays MAC filter information.

Parameters

mac-filter-id

Displays detailed information for the specified filter ID and its filter entries.

Values 1 to 65535

associations

Displays information about where the filter policy ID is applied to the detailed filter policy ID output.

counters

Displays counter information for the specified filter ID.

entry *entry-id*

Displays information about the specified filter entry ID for the specified filter ID only.

Values 1 to 65535

Output

The following outputs are examples of MAC filter information, and the associated tables describe the output fields.

- [Sample Output, Sample Detailed Output, Table 57: Output fields: filter MAC](#)
- [Sample Output for MAC Filter Associations, Table 58: Output fields: filter MAC associations](#)
- [Sample Output for MAC Filter Counters, Table 59: Output fields: filter MAC counters](#)

Sample Output

```

=====
Mac Filter
=====
Filter Id      : 8                               Applied       : Yes
Scope         : Template                       Def. Action   : Forward
Entries       : 2
Description    : Description for Mac Filter Policy id # 8
-----
Filter Match Criteria : Mac
-----
Entry         : 8                               FrameType     : Ethernet
Ing. Matches  : 80 pkts
Egr. Matches  : 62 pkts

Entry         : 10                              FrameType     : Ethernet
Ing. Matches  : 80 pkts
Egr. Matches  : 80 pkts
    
```

Sample Detailed Output

```

=====
Mac Filter : 200
=====
Filter Id      : 200                               Applied       : No
Scope         : Exclusive                       D. Action     : Drop
Description    : Forward SERVER sourced packets
-----
Filter Match Criteria : Mac
-----
Entry         : 200                              FrameType     : 802.2SNAP
Description   : Not Available
Src Mac       : 00:00:5a:00:00:00 ff:ff:ff:00:00:00
Dest Mac      : 00:00:00:00:00:00 00:00:00:00:00:00
Dot1p        : Undefined                       Ethertype     : 802.2SNAP
Match action  : Forward
Ing. Matches  : 0                               Egr. Matches  : 0
Entry         : 300 (Inactive)                 FrameType     : Ethernet
Description   : Not Available
Src Mac       : 00:00:00:00:00:00 00:00:00:00:00:00
Dest Mac      : 00:00:00:00:00:00 00:00:00:00:00:00
Dot1p        : Undefined                       Ethertype     : Ethernet
Match action  : Default
Ing. Matches  : 0                               Egr. Matches  : 0
=====
    
```

Table 57: Output fields: filter MAC

Label	Description
MAC Filter Filter Id	The MAC filter policy ID.
Scope	Template — The filter policy is of type Template.
	Exclusive — The filter policy is of type Exclusive.
Description	The IP filter policy description.

Label	Description
Applied	No — The filter policy ID has not been applied.
	Yes — The filter policy ID has been applied.
Def. Action	Forward — The default action for the filter ID for packets that do not match the filter entries is to forward.
	Drop — The default action for the filter ID for packets that do not match the filter entries is to drop.
Filter Match Criteria	MAC — Indicates that the filter is an MAC filter policy.
Entry	The filter ID filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified.
Description	The filter entry description.
FrameType	Ethernet — The entry ID match frame type is Ethernet IEEE 802.3.
	Ethernet II — The entry ID match frame type is Ethernet Type II
Src MAC	The source MAC address and mask match criterion. When both the MAC address and mask are all zeros, no criterion specified for the filter entry.
Dest MAC	The destination MAC address and mask match criterion. When both the MAC address and mask are all zeros, no criterion specified for the filter entry.
Dot1p	The IEEE 802.1p value for the match criteria. Undefined indicates no value is specified.
Ethertype	The Ethertype value match criterion.
Match action	Default — The filter does not have an explicit forward or drop match action specified. If the filter entry ID indicates the entry is Inactive, the filter entry is incomplete, no action was specified.
	Drop — Packets matching the filter entry criteria are dropped.
	Forward — Packets matching the filter entry criteria are forwarded.
Ing. Matches	The number of ingress filter matches or hits for the filter entry.
Egr. Matches	The number of egress filter matches or hits for the filter entry.

Sample Output for MAC Filter Associations

```
A:ALA-49# show filter mac 3 associations
```

```
=====
Mac Filter
```

```

=====
Filter ID   : 3                               Applied    : Yes
Scope      : Template                       Def. Action : Drop
Entries    : 1
-----
Filter Association : Mac
-----
Service Id  : 1001                            Type       : VPLS
- SAP      1/1/1:1001 (Egress)
=====
A:ALA-49#
    
```

Table 58: Output fields: filter MAC associations

Label	Description
Filter Association	Mac — The filter associations displayed are for a MAC filter policy ID.
Service Id	The service ID on which the filter policy ID is applied.
SAP	The Service Access Point on which the filter policy ID is applied.
Type	The type of service of the Service ID.
(Ingress)	The filter policy ID is applied as an ingress filter policy on the interface.
(Egress)	The filter policy ID is applied as an egress filter policy on the interface.

Sample Output for MAC Filter Counters

```
A:ALA-49# show filter mac 8 counters
```

Table 59: Output fields: filter MAC counters

Label	Description
Mac Filter	The MAC filter policy ID.
Filter Id	
Scope	Template — The filter policy is of type Template.
	Exclusive — The filter policy is of type Exclusive.
Description	The MAC filter policy description.
Applied	No — The filter policy ID has not been applied.
	Yes — The filter policy ID has been applied.
Def. Action	Forward — The default action for the filter ID for packets that do not match the filter entries is to forward.

Label	Description
	Drop — The default action for the filter ID for packets that do not match the filter entries is to drop.
Filter Match Criteria	Mac — Indicates that the filter is an MAC filter policy.
Entry	The filter ID filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified.
FrameType	Ethernet II — The entry ID match frame type is Ethernet Type II.
Ing. Matches	The number of ingress filter matches or hits for the filter entry.
Egr. Matches	The number of egress filter matches or hits for the filter entry.

4.8.2.3 Clear commands

ip

Syntax

ip *ip-filter-id* [**entry** *entry-id*] [**ingress** | **egress**]

Context

clear>filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears the counters associated with the IP filter policy.

By default, all counters associated with the filter policy entries are reset. The scope of which counters are cleared can be narrowed using the command line parameters.

Parameters

ip-filter-id

Specifies the IP filter policy ID.

Values 1 to 65535

entry-id

Specifies that only the counters associated with the specified filter policy entry will be cleared.

Values 1 to 65535

ingress

Specifies to only clear the ingress counters.

egress

Specifies to only clear the egress counters.

ipv6

Syntax

ipv6 *ip-filter-id* [**entry** *entry-id*] [**ingress** | **egress**]

Context

clear>filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears the counters associated with the IPv6 filter policy.

By default, all counters associated with the filter policy entries are reset. The scope of which counters are cleared can be narrowed using the command line parameters.

Parameters

ip-filter-id

Specifies the IP filter policy ID.

Values 1 to 65535

entry-id

Specifies that only the counters associated with the specified filter policy entry will be cleared.

Values 1 to 65535

ingress

Specifies to only clear the ingress counters.

egress

Specifies to only clear the egress counters.

mac

Syntax

mac *mac-filter-id* [**entry** *entry-id*] [**ingress** | **egress**]

Context

clear>filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears the counters associated with the MAC filter policy.

By default, all counters associated with the filter policy entries are reset. The scope of which counters are cleared can be narrowed using the command line parameters.

Parameters

mac-filter-id

Specifies the MAC filter policy ID.

Values 1 to 65535

entry-id

Specifies that only the counters associated with the specified filter policy entry will be cleared.

Values 1 to 65535

ingress

Specifies to only clear the ingress counters.

egress

Specifies to only clear the egress counters.

4.8.2.4 Monitor commands

ip

Syntax

ip *ip-filter-id* **entry** *entry-id* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]

Context

monitor>filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command monitors the counters associated with the IP filter policy.

Parameters

ip-filter-id

Specifies the IP filter policy ID.

Values 1 to 65535

entry-id

Specifies that only the counters associated with the specified filter policy entry will be monitored.

Values 1 to 65535

interval

Specifies the interval for each display, in seconds.

Values 3 to 60

Default 10 seconds

repeat repeat

Specifies how many times the command is repeated.

Values 1 to 999

Default 10

absolute

Displays raw statistics without processing. No calculations are performed on the delta or rate statistics.

rate

Displays the rate-per-second for each statistic instead of the delta.

ipv6

Syntax

```
ipv6 ip-filter-id entry entry-id [interval seconds] [repeat repeat] [absolute | rate]
```

Context

```
monitor>filter
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command monitors the counters associated with the IPv6 filter policy.

Parameters

ip-filter-id

Specifies the IP filter policy ID.

Values 1 to 65535

entry-id

Specifies that only the counters associated with the specified filter policy entry will be monitored.

Values 1 to 65535

interval seconds

Specifies the interval for each display, in seconds.

Values 3 to 60

Default 10

repeat repeat

Specifies how many times the command is repeated.

Values 1 to 999

Default 10

absolute

Displays raw statistics without processing. No calculations are performed on the delta or rate statistics.

rate

Displays the rate-per-second for each statistic instead of the delta.

mac

Syntax

mac *mac-filter-id* **entry** *entry-id* [**interval seconds**] [**repeat repeat**] [**absolute | rate**]

Context

monitor>filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command monitors the counters associated with the MAC filter policy.

Parameters

mac mac-filter-id

Specifies the MAC filter policy ID.

Values 1 to 65535

entry entry-id

Specifies that only the counters associated with the specified filter policy entry will be cleared.

Values 1 to 65535

interval seconds

Specifies the interval for each display, in seconds.

Values 3 to 60

Default 5

repeat repeat

Specifies how many times the command is repeated.

Values 1 to 999

Default 10

absolute

Displays raw statistics without processing. No calculations are performed on the delta or rate statistics.

rate

Displays the rate-per-second for each statistic instead of the delta.

5 Cflowd



Note:

Cflowd is supported only on the 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone).

This chapter provides information to configure the cflowd tool.

5.1 Cflowd overview

Cflowd is a tool used to sample IPv4, IPv6, MPLS, and Ethernet traffic data flows through a router. Cflowd enables ISPs and traffic engineers to perform traffic sampling and analysis to support capacity planning, trends analysis, and characterization of workloads in a network service provider environment.

Cflowd is also useful for traffic engineering, network planning and analysis, network monitoring, developing user profiles, data warehousing and mining, and performing security-related investigations. Collected information can be interpreted in several ways such as in port, autonomous system (AS), or network matrices, and pure flow structures. The amount of data stored depends on the cflowd configurations.

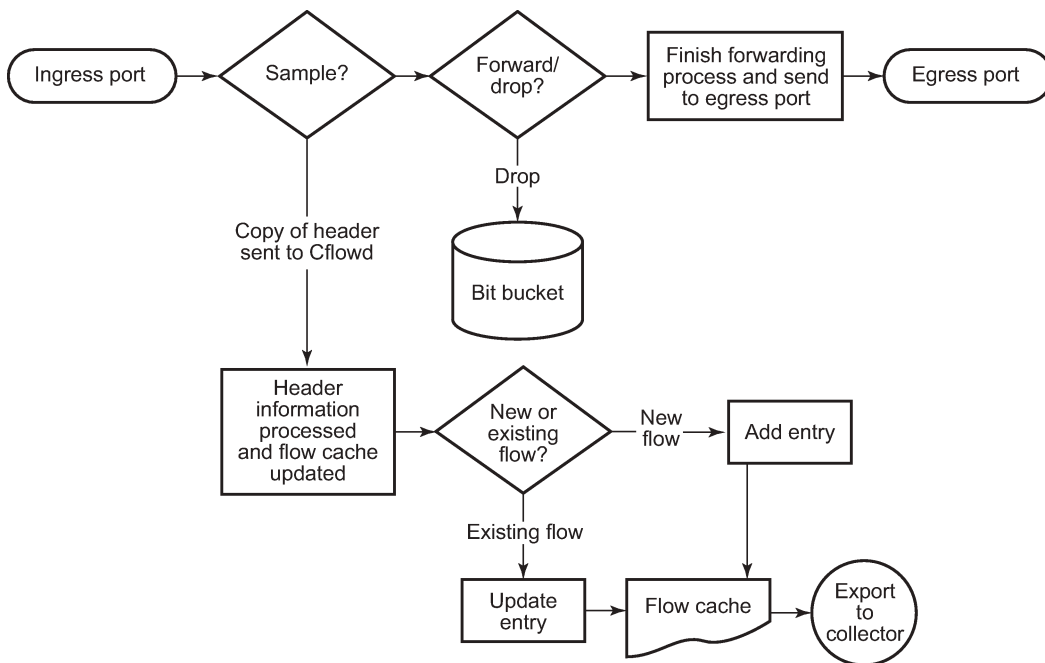
Cflowd maintains a list of router data flows. A flow is a unidirectional traffic stream defined by several characteristics such as source and destination IP addresses, source and destination ports, inbound interface, IP protocol, and Type-of-Service (TOS) bits.

When a router receives a packet for which it currently does not have a flow entry, a flow structure is initialized to maintain state information about that flow, such as the number of bytes exchanged, IP addresses, port numbers, AS numbers, and so on. Each subsequent packet matching the same parameters of the flow contributes to the byte and packet count of the flow until the flow is terminated and exported to a collector for storage.

5.1.1 Operation

The following figure shows the basic operation of the cflowd feature. This sample flow only describes the basic cflowd operation overview and is not intended to specify implementation and support on the 7210 SAS.

Figure 12: Basic cflowd steps



Router_Config_30

The logical sequence of cflowd operation is as follows:

1. The system decides whether to forward or drop packets as the packets ingress a port.
2. The system decides whether to sample the packet for cflowd, then the packet is forwarded or dropped.
3. If a new flow is found, the system adds a new entry to the cache. If the flow already exists in the cache, the system updates the flow statistics.
4. If a new flow is detected and the maximum number of entries are already present in the flow cache, the system removes the entry with the earliest expiry time. The earliest expiry entry/flow is the next flow that will expire based on the active or inactive timer expiration.
5. If a flow has been inactive for a period of time equal to or greater than the inactive timer (default 15 seconds), or has been active for a period of time equal to or greater than the active timer (default 30 minutes), the system removes the entry from the flow cache.

When a flow is exported from the cache, the collected data is sent to an external collector that maintains an accumulation of historical data flows, which network operators can use to analyze traffic patterns.

Data is exported in one of the following formats:

- **Version 5**

This format generates a fixed export record for each individual flow captured.

- **Version 8**

This format aggregates multiple individual flows into a fixed aggregate record.

- **Version 9**

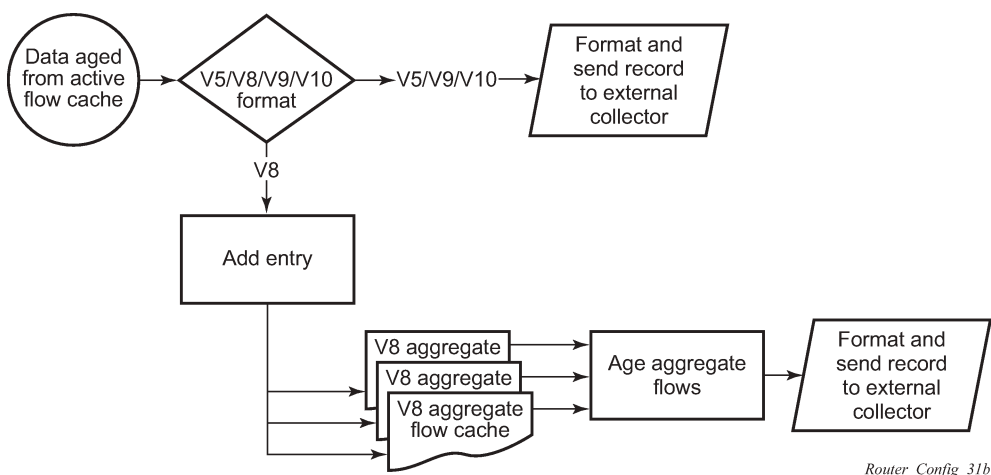
This format generates a variable export record, depending on user configuration and sampled traffic type (IPv4, IPv6, or MPLS), for each individual flow captured.

- **Version 10 (IPFIX)**

This format generates a variable export record, depending on user configuration and sampled traffic type (IPv4, IPv6, or MPLS), for each individual flow captured.

The following figure shows Version 5, Version 8, Version 9, and Version 10 flow processing.

Figure 13: V5, V8, V9, V10, and flow processing



As flows expire and are removed from the active flow cache, the export format is determined (either Version 5, Version 8, Version 9, and Version 10 record format) and one of the following processes occurs:

- If the export format is Version 5, Version 9, or Version 10, no further processing is performed and the flow data is accumulated to be sent to the external collector.
- If the export format is Version 8, the flow entry is added to one or more of the configured aggregation matrices.

As the entries within the aggregate matrices are aged out, they are accumulated to be sent to the external flow collector in Version 8 format.

The sample rate and cache size are configurable values. The cache size is set up with the default number of entries.

A flow terminates when one of the following conditions is met:

- The inactive timeout period expires (default 15 seconds). A flow is considered terminated when no packets are seen for the flow for the configured number of seconds.
- An active timeout expires (default 30 minutes). A flow terminates according to the time duration, regardless of whether packets are coming in for the flow.
- The user executes a **clear cflowd** command.
- Other conditions are met to aggressively age flows as the cache becomes too full, such as **overflow percent**.

5.1.1.1 Version 8

There are several aggregate flow types including:

- AS matrix

- destination prefix matrix
- source prefix matrix
- source-destination prefix matrix
- protocol/port matrix

Version 8 is an aggregated export format. As individual flows are aged out of the raw flow cache, the data is added to the aggregate flow cache for each configured aggregate type. Each of these aggregate flows are also aged in a manner similar to the method the active flow cache entries are aged. When an aggregate flow is aged out, it is sent to the external collector in the Version 8 record format.

5.1.1.2 Version 9

The Version 9 format is a more flexible and allows for different templates or sets of cflowd data to be sent based on the sampled traffic type and the configured template set.

Version 9 is interoperable with RFC 3954, *Cisco Systems NetFlow Services Export Version 9*.

5.1.1.3 Version 10

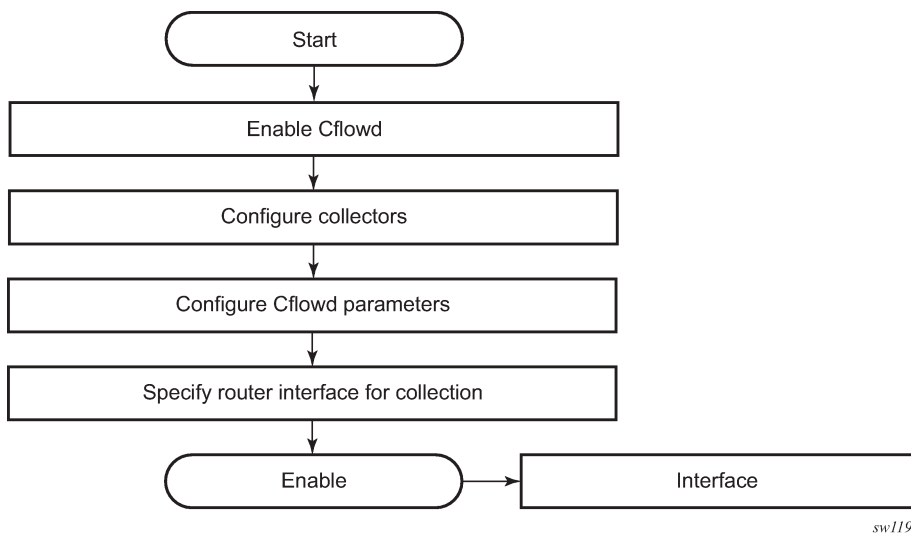
Version 10 is a new format and protocol that interoperates with the IETF specifications described in the IP Flow Information Export (IPFIX) standard. Like Version 9, Version 10 uses templates to export different data elements for a flow and handle different types of data flows, such as IPv4, IPv6, and MPLS.

Version 10 is interoperable with RFC 5150 and RFC 5102.

5.2 Cflowd configuration process overview

The following figure shows the process to configure cflowd parameters.

Figure 14: Cflowd configuration and implementation flow



Cflowd can be enabled to sample traffic on a specific interface in the cflowd interface mode. In this mode, all traffic entering a specific port is subject to sampling as the configured sampling rate.

5.3 Configuration notes

The following cflowd components must be configured for cflowd to be operational:

- Cflowd must be enabled globally.
- At least one collector must be configured and enabled.
- A cflowd option must be specified and enabled on a router interface.
- Sampling must be enabled on the interface (ingress only).
- On the 7210 SAS, when cflowd is enabled on an IP interface, the sampling rate is applied to a port and only the samples that match the IP interface for which cflowd is enabled are processed further to update or create flow records in the flow cache. Samples received that do not match the IP interface for which cflowd is enabled are not processed further, and flow records are not created for them.
- On the 7210 SAS, when cflowd is enabled on a SAP in a Layer 2 service (for example, VPLS and Epipe), the sampling rate is applied to a port. Only samples that match the SAP tag for which cflowd is enabled are processed further to update or create flow records in the flow cache. Samples received that do not match the SAP for which cflowd is enabled are not processed further, and no flow records are created.
- On the 7210 SAS, samples are collected only in the ingress direction. Sampling in the egress direction is not supported.
- On the 7210 SAS-Mxp and 7210 SAS-R6, packets received on RSVP LSPs are eligible for cflowd processing.
- Cflowd is not supported on SAPs in EVPN-VPWS services.

5.4 Configuring cflowd with CLI

This section provides information to configure cflowd using the command line interface.

5.5 Cflowd configuration overview

The cflowd implementation supports traffic flow analysis and the use of traffic and access list (ACL) filters to limit the type of traffic analyzed.

5.5.1 Traffic sampling

Traffic sampling does not examine all packets received by a router. The user can configure command parameters to modify the rate at which traffic is sampled and sent for flow analysis. The default sampling rate is one out of every 1000 packets.



Caution:

Excessive sampling, such as one out of every 100 packets, over an extended period of time can burden router processing resources.

The following data is maintained for each individual flow in the raw flow cache:

- source IP address
- destinations IP address
- source port
- destination port
- forwarding status
- input interface
- output interface
- IP protocol
- TCP flags
- first timestamp (of the first packet in the flow)
- last timestamp (timestamp of last packet in the flow before expiry of the flow)
- source AS number for peer and origin (taken from BGP)
- destination AS number for peer and origin (taken from BGP)
- IP next hop
- BGP next hop
- ICMP type and code
- IP version
- source prefix (from routing)
- destination prefix (from routing)
- MPLS label stack from label 1 to 6

Within the raw flow cache, the following characteristics are used to identify an individual flow:

- ingress interface
- source IP address
- destination IP address
- source transport port number
- destination transport port number
- IP protocol type
- IP TOS byte
- virtual router ID
- ICMP type and code
- direction
- MPLS labels

The user enables cflowd at the interface level. By enabling cflowd at the interface level, all IP packets forwarded by the interface are subject to cflowd analysis.

5.5.2 Collectors

A collector defines how data flows are exported from the flow cache. The user can configure a maximum of five collectors. Each collector is identified by a unique IP address and UDP port value. Each collector can only export traffic in one version type: Version 5, Version 8, Version 9, or Version 10.

The user can modify the parameters of a collector configuration or retain the defaults.

The **autonomous-system-type** command defines whether the autonomous system (AS) information is included in the flow data based on the originating AS or external peer AS of the flow.

5.5.2.1 Aggregation

Version 8 allows the aggregation of flow data into larger, less granular flows. Use aggregation commands to specify the type of data to collect. These aggregation types are only applicable to flows that are exported to a Version 8 collector.

The following aggregation schemes are supported:

- **AS matrix**
Flows are aggregated based on source and destination AS and ingress and egress interfaces.
- **protocol-port**
Flows are aggregated based on the IP protocol, source port number, and destination port number.
- **source prefix**
Flows are aggregated based on source prefix and mask, source AS, and ingress interface.
- **destination prefix**
Flows are aggregated based on destination prefix and mask, destination AS, and egress interface.
- **source-destination prefix**
Flows are aggregated based on source prefix and mask, destination prefix and mask, source and destination AS, ingress and egress interfaces.
- **raw**
Flows are not aggregated and are sent to the collector in a Version 5 record.

5.6 Basic cflowd configuration

This section provides information to configure cflowd and examples of common configuration tasks. To sample traffic, the user must configure the following minimal cflowd parameters:

- Cflowd must be enabled.
- At least one collector must be configured and enabled.
- Sampling must be enabled on the interface (ingress only)

Output example

The following is a sample of cflowd configuration output.

```
A:Dut-D>config>cflowd$ info detail
-----
active-timeout 30
cache-size 65536
inactive-timeout 15
export-mode automatic
overflow 1
rate 1000
template-retransmit 600
no use-vrtr-if-index
collector 10.10.10.103:2055 version 9
  description "V9 collector"
  template-set basic
  no shutdown
exit
no shutdown
```

5.7 Common configuration tasks

This section provides an overview of the cflowd configuration tasks and CLI commands. To begin traffic flow sampling, cflowd and the user must enable at least one collector.

5.7.1 Global cflowd components

The following common (global) attributes apply to all instances of cflowd:

- **active timeout**
This attribute controls the maximum time a flow record can be active before it is automatically exported to defined collectors.
- **inactive timeout**
This attribute controls the minimum time before a flow is declared inactive. If no traffic is sampled for an existing flow for the inactive timeout duration, the flow is declared inactive and marked to be exported to the defined collectors.
- **cache size**
This attribute defines the maximum size of the flow cache.
- **overflow**
This attribute defines the percentage of flow records that are exported to all collectors if the flow cache size is exceeded.
- **rate**
This attribute defines the system wide sampling rate for cflowd.
- **template retransmit**
This attribute defines the interval (in seconds) at which the Version 9 and Version 10 templates are retransmitted to all configured Version 9 or Version 10 collectors.

5.7.2 Configuring cflowd

Use the following CLI syntax to perform cflowd configuration tasks.

```
config>cflowd#
  active-timeout minutes
  cache-size num-entries
  inactive-timeout seconds
  template-retransmit seconds
  overflow percent
  rate sample-rate
  collector ip-address[:port] {version [5 | 8 | 9 |10]}
    aggregation
      as-matrix
      destination-prefix
      protocol-port
      raw
      source-destination-prefix
      source-prefix
    template-set {basic | mpls-ip}
    autonomous-system-type [origin | peer]
    description description-string
  no shutdown
no shutdown
```

5.7.3 Enabling cflowd

Cflowd is disabled by default. Executing the **configure cflowd** command enables Cflowd. By default, cflowd is not shut down but must be configured, including at least one collector, to be active.

Use the following CLI syntax to enable cflowd.

```
config# cflowd
no shutdown
```

Output example

The following is a sample configuration output that shows the default values when cflowd is initially enabled. No collectors or collector options are configured.

```
A:ALA-1>config# info detail
...
#-----
echo "Cflowd Configuration"
#-----
  cflowd
    active-timeout 30
    cache-size 65536
    inactive-timeout 15
    overflow 1
    rate 1000
    template-retransmit 600
    no shutdown
  exit
#-----
A:ALA-1>config#
```

5.7.4 Configuring global cflowd parameters

This section describes the cflowd parameters that apply to all instances where cflowd (traffic sampling) is enabled.

Use the following syntax to configure cflowd parameters.

```
config>cflowd#
  active-timeout minutes
  cache-size num-entries
  inactive-timeout seconds
  overflow percent
  rate sample-rate
  template-retransmit seconds
  no shutdown
```

Output example

The following is an example of a common cflowd component configuration.

```
A:ALA-1>config>cflowd# info
#-----
      active-timeout 20
      inactive-timeout 10
      overflow 10
      rate 100
#-----
A:ALA-1>config>cflowd#
```

5.7.5 Configuring cflowd collectors

Use the following syntax to configure cflowd collector parameters.

```
config>cflowd#
  collector ip-address[:port] [version version]
  aggregation
    as-matrix
    destination-prefix
    protocol-port
    raw
    source-destination-prefix
    source-prefix
  autonomous-system-type [origin | peer]
  description description-string
  no shutdown
  template-set {basic | mpls-ip}
```

Example

The following is a sample configuration output.

```
A:ALA-1>config>cflowd# info
-----
active-timeout 20
  inactive-timeout 10
  overflow 10
  rate 100
```

```

collector 10.10.10.1:2000 version 8
  aggregation
    as-matrix
    raw
  exit
  description "AS info collector"
exit
collector 10.10.10.2:5000 version 8
  aggregation
    protocol-port
    source-destination-prefix
  exit
  autonomous-system-type peer
  description "Neighbor collector"
exit
-----
A:ALA-1>config>cflowd#

```

Example

The following is a sample configuration output for a Version 9 collector.

```

collector 10.10.10.9:2000 version 9
  description "v9collector"
  template-set mpls-ip
  no shutdown
exit

```

5.7.5.1 Version 9 and Version 10 templates

If the collector is configured to use either Version 9 or Version 10 (IPFIX) formats, the flow data is sent to the designated collector using one of the predefined templates. The template used is based on the type of flow for which the data was collected (IPv4, IPv6, MPLS, or Ethernet (Layer 2)), and the configuration of the **template-set** parameter. The following table lists traffic flow types and the corresponding template used to export the flow data.

Table 60: Template-set

Traffic type	Basic	MPLS-IP
IPv4	Basic IPv4	MPLS-IPv4
IPv6	Basic IPv6	MPLS-IPv6
MPLS	Basic MPLS	MPLS-IP
Ethernet	L2-IP	L2-IP

Each flow exported to a collector, configured for either Version 9 or Version 10 formats, is sent using one of the preceding flow template sets. The template is used based on the flow type and how the **template-set** parameter of the collector is configured.

The following tables list the fields present in each template set listed in the preceding table:

- [Table 61: Basic IPv4 template](#)
- [Table 62: MPLS-IPv4 template](#)

- [Table 63: Basic IPv6 template](#)
- [Table 64: MPLS-IPv6 template](#)
- [Table 65: Basic MPLS template](#)
- [Table 66: MPLS-IP template](#)
- [Table 67: Ethernet \(L2-IP\) flow template](#)

Table 61: Basic IPv4 template

Field name	Field ID
IPv4 Src Addr	8
IPv4 Dest Addr	12
IPv4 Nexthop	15
BGP Nexthop	18
Ingress Interface	10
Egress Interface	14
Packet Count	2
Byte Count	1
Start Time	22
End Time	21
Flow Start Milliseconds ¹⁰	152
Flow End Milliseconds1	153
Src Port	7
Dest Port	11
Forwarding Status	89
TCP control Bits (Flags)	6
IPv4 Protocol	4
IPv4 TOS	5
IP version	60
ICMP Type & Code	32
Direction	61

¹⁰ Only sent to collectors configured for the Version 10 format

Field name	Field ID
BGP Source ASN	16
BGP Dest ASN	17
Source IPv4 Prefix Length	9
Dest IPv4 Prefix Length	13

Table 62: MPLS-IPv4 template

Field name	Field ID
IPv4 Src Addr	8
IPv4 Dest Addr	12
IPv4 Nexthop	15
BGP Nexthop	18
Ingress Interface	10
Egress Interface	14
Packet Count	2
Byte Count	1
Start Time	22
End Time	21
Flow Start Milliseconds 11	152
Flow End Milliseconds	153
Src Port	7
Dest Port	11
Forwarding Status	89
TCP control Bits (Flags)	6
IPv4 Protocol	4
IPv4 TOS	5
IP version	60

¹¹ Only sent to collectors configured for the Version 10 format

Field name	Field ID
ICMP Type & Code	32
Direction	61
BGP Source ASN	16
BGP Dest ASN	17
Source IPv4 Prefix Length	9
Dest IPv4 Prefix Length	13
MPLS Top Label Type	46
MPLS Top Label IPv4 Addr	47
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
MPLS Label 4	73
MPLS Label 5	74
MPLS Label 6	75

Table 63: Basic IPv6 template

Field name	Field ID
IPv6 Src Addr	27
IPv6 Dest Addr	28
IPv6 Nexthop	62
IPv6 BGP Nexthop	63
IPv4 Nexthop	15
IPv4 BGP Nexthop	18
Ingress Interface	10
Egress Interface	14
Packet Count	2
Byte Count	1
Start Time	22

Field name	Field ID
End Time	21
Flow Start Milliseconds ¹²	152
Flow End Milliseconds1	153
Src Port	7
Dest Port	11
Forwarding Status	89
TCP control Bits (Flags)	6
Protocol	4
IPv6 Extension Hdr	64
IPv6 Next Header	193
IPv6 Flow Label	31
TOS	5
IP version	60
IPv6 ICMP Type & Code	139
Direction	61
BGP Source ASN	16
BGP Dest ASN	17
IPv6 Src Mask	29
IPv6 Dest Mask	30

Table 64: MPLS-IPv6 template

Field name	Field ID
IPv6 Src Addr	27
IPv6 Dest Addr	28
IPv6 Nexthop	62
IPv6 BGP Nexthop	63

¹² Only sent to collectors configured for the Version 10 format

Field name	Field ID
IPv4 Nexthop	15
IPv4 BGP Nexthop	18
Ingress Interface	10
Egress Interface	14
Packet Count	2
Byte Count	1
Start Time	22
End Time	21
Flow Start Milliseconds 13	152
Flow End Milliseconds1	153
Src Port	7
Dest Port	11
Forwarding Status	89
TCP control Bits (Flags)	6
Protocol	4
IPv6 Extension Hdr	64
IPv6 Next Header	193
IPv6 Flow Label	31
TOS	5
IP version	60
IPv6 ICMP Type & Code	139
Direction	61
BGP Source ASN	16
BGP Dest ASN	17
IPv6 Src Mask	29

¹³ Only sent to collectors configured for the Version 10 format

Field name	Field ID
IPv6 Dest Mask	30
MPLS_TOP_LABEL_ TYPE	46
MPLS_TOP_LABEL_ ADDR	47
MPLS Top Label Type	46
MPLS Top Label IPv6 Addr	47
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
MPLS Label 4	73
MPLS Label 5	74
MPLS Label 6	75
MPLS_TOP_LABEL_ TYPE	46
MPLS_TOP_LABEL_ ADDR	47

Table 65: Basic MPLS template

Field name	Field ID
Start Time	22
End Time	21
Flow Start Milliseconds ¹⁰	152
Flow End Milliseconds ¹⁰	153
Ingress Interface	10
Egress Interface	14
Packet Count	2
Byte Count	1
Direction	61

Field name	Field ID
MPLS_TOP_LABEL_TYPE	46
MPLS_TOP_LABEL_ADDR	47
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
MPLS Label 4	73
MPLS Label 5	74
MPLS Label 6	75

Table 66: MPLS-IP template

Field name	Field ID
IPv4 Src Addr	8
IPv4 Dest Addr	12
IPv4 Nexthop	15
IPv6 Src Addr	27
IPv6 Dest Addr	28
Ingress Interface	10
Egress Interface	14
Packet Count	2
Byte Count	1
Start Time	22
End Time	21
Flow Start Milliseconds ¹⁰	152
Flow End Milliseconds ¹⁰	153
Src Port	7
Dest Port	11
TCP control Bits (Flags)	6

Field name	Field ID
IPv4 Protocol	4
IPv4 TOS	5
IP version	60
ICMP Type & Code	32
Direction	61
MPLS_TOP_LABEL_ TYPE	46
MPLS_TOP_LABEL_ ADDR	47
MPLS Top Label Type	46
MPLS Top Label IPv4 Addr	47
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
MPLS Label 4	73
MPLS Label 5	74
MPLS Label 6	75

Table 67: Ethernet (L2-IP) flow template

Field name ¹⁴	Field ID
MAC Src Addr	56
MAC Dest Addr	80
Ingress Physical Interface	252
Egress Physical Interface	253
Dot1q VLAN ID	243
Dot1q Customer VLAN ID	245
Post Dot1q VLAN ID	254

¹⁴ The Ethernet (L2-IP) flow template is only supported and exported to IPFIX (version 10) collectors.

Field name ¹⁴	Field ID
Post Dot1q Customer VLAN Id	255
IPv4 Src Addr	8
IPv4 Dest Addr	12
IPv6 Src Addr	27
IPv6 Dest Addr	28
Packet Count	2
Byte Count	1
Flow Start Milliseconds	152
Flow End Milliseconds	153
Src Port	7
Dest Port	11
TCP control Bits (Flags)	6
Protocol	4
IPv6 Option Header	64
IPv6 Next Header	196
IPv6 Flow Label	31
TOS	5
IP Version	60
ICMP Type Code	32

5.7.6 Specifying cflowd options on an IP interface

When cflowd is enabled on an interface, all packets forwarded by the interface are subject to analysis according to the global cflowd configuration and sorted according to the collector configurations.

See [Table 68: Cflowd configuration dependencies](#) for configuration combinations.

When the **cflowd interface** option is configured in the **config>router>interface** context, the following requirements must be met to enable traffic sampling on the specific interface:

- Cflowd must be enabled.
- At least one cflowd collector must be configured and enabled.

- The **interface>cflowd interface** option must be selected. For configuration information, see [Filter policy configuration overview](#).

5.7.6.1 Interface configurations

Use the following CLI syntax to enable traffic sampling on an interface.

```
config>router>if>cflowd-parameters#  
  sampling {unicast|multicast} type {interface} [direction {ingress-only}]  
  no sampling {unicast|multicast}
```

When the **interface** option is configured, cflowd extracts traffic flow samples from an interface for analysis. All packets forwarded by the interface are analyzed in accordance with the cflowd configuration.

Configure the **interface** option to enable traffic sampling on an interface. If cflowd is not enabled (**no cflowd**), traffic sampling does not occur on the interface.

5.7.6.2 Service interfaces

Use the following CLI syntax to enable traffic sample on a service interface.

```
config>service>ies>if>cflowd-parameters# sampling {unicast|multicast} type {interface}  
  [direction {ingress-only}]  
  config>service>vprn>if>cflowd-parameters# sampling {unicast|multicast} type {interface}  
  [direction {ingress-only}]  
  no sampling {unicast|multicast}
```

When enabled on a service interface, cflowd collects routed traffic flow samples through a router for analysis. Cflowd is supported on IES and VPRN service interfaces only. Layer 2 traffic is excluded. All packets forwarded by the interface are analyzed according to the **cflowd** configuration. On the interface level, cflowd can be associated with an IP interface.

5.7.7 Dependencies

For cflowd to be operational, the following requirements must be met:

- Cflowd must be enabled on a global level. If cflowd is disabled, any traffic sampling instances are also disabled.
- At least one collector must be configured and enabled for traffic sampling to occur on an enabled entity.
- If a specific collector UDP port is not identified, flows are sent to port 2055 by default.

The following table displays the expected results when specific features are enabled and disabled.

Table 68: Cflowd configuration dependencies

Interface setting	router>interface cflowd [interface] setting	Command ip-filter entry setting	Expected results
Interface mode ¹⁵	Interface	none	All IP traffic ingressing the interface is subject to sampling

5.8 Cflowd configuration management tasks

This section describes cflowd configuration management tasks.

5.8.1 Modifying global cflowd components

Cflowd parameter modifications apply to all instances where cflowd or traffic sampling is enabled. Changes are applied immediately. Use the following syntax to modify global cflowd parameters.

```
config>cflowd#
  active-timeout minutes
  no active-timeout
  cache-size num-entries
  no cache-size
  inactive-timeout seconds
  no inactive-timeout
  overflow percent
  no overflow
  rate sample-rate
  no rate
  [no] shutdown
  template-retransmit seconds
  no template-retransmit
```

The following example shows the cflowd command usage to modify configuration parameters.

Example

```
config>cflowd# active-timeout 60 config>cflowd# no inactive-timeout config>cflowd# overflow 2
config>cflowd# rate 10
```

Output example

The following is a sample cflowd component configuration output.

```
A:ALA-1>config>cflowd# info
#-----
      active-timeout 60
      overflow 2
      rate 10
#-----
```

¹⁵ See [Configuration notes](#) for more information.

```
A:ALA-1>config>cflowd#
```

5.8.2 Modifying cflowd collector parameters

Use the following syntax to modify cflowd collector and aggregation parameters.

```
config>cflowd#
  collector ip-address[:port] [version version]
  no collector ip-address[:port]
    [no] aggregation
        [no] as-matrix
        [no] destination-prefix
        [no] protocol-port
        [no] raw
        [no] source-destination-prefix
        [no] source-prefix
    [no] autonomous-system-type [origin | peer]
    [no] description description-string
    [no] shutdown
    template-set {basic | mpls-ip}
```

If a specific collector UDP port is not identified, flows are sent to port 2055 by default.

Output example

The following sample output shows basic cflowd modifications.

```
A:ALA-1>config>cflowd# info
-----
  active-timeout 60
  overflow 2
  rate 10
  collector 10.10.10.1:2000 version 5
    description "AS info collector"
  exit
  collector 10.10.10.2:5000 version 8
    aggregation
      source-prefix
      raw
    exit
    description "Test collector"
  exit
-----
A:ALA-1>config>cflowd#
```

5.9 Cflowd configuration command reference

5.9.1 Command hierarchies

- [Configuration commands](#)
- [Show commands](#)
- [Tools commands](#)

- [Clear commands](#)

5.9.1.1 Configuration commands

```
config
- [no] cflowd
  - active-timeout minutes
  - no active-timeout
  - cache-size num-entries
  - no cache-size
  - collector ip-address[:port] [version version]
  - no collector ip-address[:port]
    - [no] aggregation
      - [no] as-matrix
      - [no] destination-prefix
      - [no] protocol-port
      - [no] raw
      - [no] source-destination-prefix
      - [no] source-prefix
    - autonomous-system-type {origin | peer}
    - description description-string
    - no description
    - [no] shutdown
    - template-set {basic | mpls-ip | l2-ip}
  - export-mode {automatic | manual}
  - inactive-timeout seconds
  - no inactive-timeout
  - overflow percent
  - no overflow
  - rate sample-rate
  - no rate
  - [no] shutdown
  - template-retransmit seconds
  - no template-retransmit
  - [no] use-vrtr-if-index
```

5.9.1.2 Show commands

```
show
- cflowd
  - collector [ip-address[:port]] [detail]
  - interface [ip-int-name]
  - status
```

5.9.1.3 Tools commands

```
tools
- dump
  - cflowd
    - cache aggregate {src-dst-proto | src-dst-proto-port} family {ipv4 | ipv6}
    - cache all family {ipv4 | ipv6}
    - packet-size protocol [clear]
    - top-flows protocols [clear]
    - top-protocols protocols [clear]
```

5.9.1.4 Clear commands

```
clear  
- cflowd
```

5.9.2 Command descriptions

5.9.2.1 Global commands

cflowd

Syntax

[no] cflowd

Context

config>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure cflowd.

The **no** form of this command removes all configuration under cflowd, including all configured collectors. The **no** form can only be executed if cflowd is shut down.

Default

no cflowd

active-timeout

Syntax

active-timeout *minutes*

no active-timeout

Context

config>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the maximum amount of time before an active flow is aged out of the active cache. If a specific flow is active for the configured amount of time, the flow is aged out and a new flow is created on the next packet sampled for that flow.

If the *minutes* parameter is changed while cflowd is active, the existing flows do not inherit the new active timeout value. The active timeout value for a flow is set when the flow is first created in the active cache table; the value does not change dynamically.

The **no** form of this command resets the inactive timeout back to default value.

Default

active-timeout 30

Parameters

minutes

Specifies the value, expressed in minutes, before an active flow is exported.

Values 1 to 600

cache-size

Syntax

cache-size *num-entries*

no cache-size

Context

config>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command specifies the maximum number of active entries maintained in the flow cache table.

The **no** form of this command reverts the number of active entries to the default value.

Default

cache-size 65536

Parameters

num-entries

Specifies the maximum number of entries maintained in the cflowd cache.

Values 1000 to 131072

collector

Syntax

collector *ip-address[:port]* [**version** *version*]

no collector *ip-address[:port]*

Context

config>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command defines a flow data collector for cflowd data. The IP address of the flow collector must be specified.

If the optional UDP port number parameter is not configured, default port 2055 is used for all collector versions. To connect to an IPFIX (version 10) collector using the IPFIX default port, specify port 4739 when configuring the collector. The version must be specified. A maximum of five collectors can be configured.

The **no** form of this command removes the flow collector definition from the configuration and stops the export of data to the collector. The collector must be shut down before it can be deleted.

Parameters

ip-address

Specifies the address of a remote cflowd collector host that will receive the exported cflowd data.

Values

<ip-address[:port]>	ip-address - a.b.c.d[:port]	(IPv4)
	x:x:x:x:x:x:x	(IPv6)
	[x:x:x:x:x:x]:port	(IPv6)
	x - [0..FFFF]H	

port

Specifies the UDP port number on the remote cflowd collector host that will receive the exported cflowd data.

Values 1 to 65535

Default 2055

version

Specifies the version of the flow data collector.

Values 5, 8, 9, 10

Default 5

aggregation

Syntax

[no] aggregation

Context

config>cflowd>collector

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables data aggregation for the collector and commands in this context configure the aggregation types.

To configure aggregation, you must choose the aggregation scheme: autonomous system, destination prefix, protocol port, raw, source destination, or source prefix.

This command can only be configured if the collector version is configured as Version 8.

The **no** form of this command removes all aggregation types from the collector configuration.

Default

no aggregation

as-matrix

Syntax

[no] as-matrix

Context

config>cflowd>collector>aggregation

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables cflowd aggregation based on autonomous system (AS) information. An AS matrix contains packet and byte counters for traffic from either source-destination ASs or last-peer to next-peer ASs.

The **no** form of this command removes this type of aggregation from the collector configuration.

Default

no as-matrix

destination-prefix

Syntax

[no] destination-prefix

Context

config>cflowd>collector>aggregation

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables cflowd aggregation based on destination prefix information.

The **no** form removes this type of aggregation from the collector configuration.

Default

no destination-prefix

protocol-port

Syntax

[no] protocol-port

Context

config>cflowd>collector>aggregation

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables cflowd aggregation based on the IP protocol, source port number, and destination port number.

The **no** form of this command removes this type of aggregation from the collector configuration.

Default

no protocol-port

raw

Syntax

[no] raw

Context

config>cflowd>collector>aggregation

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables the sending of raw (unaggregated) flow data in Version 5.

The **no** form of this command removes this type of aggregation from the collector configuration.

Default

no raw

source-destination-prefix

Syntax

[no] source-destination-prefix

Context

config>cflowd>collector>aggregation

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures cflowd aggregation based on source and destination prefixes.

The **no** form of this command removes this type of aggregation from the collector configuration.

Default

no source-destination-prefix

source-prefix

Syntax

[no] source-prefix

Context

config>cflowd>collector>aggregation

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures cflowd aggregation based on source prefix information.

The **no** form of this command removes this type of aggregation from the collector configuration.

Default

no source-prefix

autonomous-system-type

Syntax

autonomous-system-type {**origin** | **peer**}

Context

config>cflowd>collector

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures whether the AS information included in the flow data is based on the originating AS or external peer AS of the routes.

This option is supported only if the collector is configured as Version 5 or Version 8.

Default

autonomous-system-type origin

Parameters

origin

Keyword to specify that the AS information included in the flow data is based on the originating AS.

peer

Keyword to specify that the AS information included in the flow data is based on the peer AS.

description

Syntax

description *description-string*

no description

Context

config>cflowd>collector

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command creates a text description stored in the configuration file for a configuration context.

The **no** form of this command removes the description string from the context.

Parameters

description-string

Specifies the description character string, up to 80 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.

shutdown

Syntax

[no] shutdown

Context

config>cflowd

config>cflowd>collector

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

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 can be deleted.

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.

The **no** form of this command administratively enables an entity.

Default

no shutdown

template-set

Syntax

template-set {**basic** | **mpls-ip** | **I2-ip**}

Context

config>cflowd>collector

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the set of templates sent to the collector when using cflowd Version 9 or Version 10.

Default

template-set basic

Parameters

basic

Keyword to send basic flow data.

mpls-ip

Keyword to send extended flow data that includes IP and MPLS flow information.

I2-ip

Keyword to send extended flow data that includes Layer 2 (Ethernet) and IP flow information.

export-mode

Syntax

export-mode [**automatic** | **manual**]

Context

config>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures how exports are generated by the cflowd process.

The default behavior is for flow data to be exported automatically based on the active and inactive time-out values. In manual mode, flow data is exported only when the **tools perform cflowd manual-export** command is issued. The only exception is if the cflowd cache overflows, in which case the normal automatic export process is used.

Default

export-mode automatic

Parameters

automatic

Keyword to automatically generate cflowd flow data.

manual

Keyword to export cflowd flow data only when manually triggered.

inactive-timeout

Syntax

inactive-timeout *seconds*

no inactive-timeout

Context

config>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the amount of time, in seconds, that must elapse without a packet matching a flow before the flow is considered inactive.

If the *seconds* parameter is changed while cflowd is active, the existing flows do not inherit the new inactive timeout value. The inactive timeout value for a flow is set when the flow is first created in the active cache table; the value does not change dynamically.

The **no** form of this command reverts the inactive timeout to the default value.

Default

inactive-timeout 15

Parameters

seconds

Specifies the amount of time, in seconds, that must elapse without a packet matching before the flow is considered inactive

Values 10 to 600

overflow

Syntax

overflow *percent*

no overflow

Context

config>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command specifies the percentage of the flow cache entries removed when the maximum number of entries is exceeded. Entries that have not been updated for the longest amount of time are removed.

The **no** form of this command reverts the number of entries cleared from the flow cache on overflow to the default value.

Default

overflow 1

Parameters

percent

Specifies the percentage of the flow cache entries removed when the maximum number of entries is exceeded.

Values 1 to 50

rate

Syntax

rate *sample-rate*

no rate

Context

config>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command specifies the rate (N) at which traffic is sampled and sent for flow analysis. A packet is sampled every N packets. For example, if *sample-rate* is configured as 1, all packets are sent to the cache. If *sample-rate* is configured as 100, one out of every 100 packets is sent to the cache.



Note:

On the 7210 SAS, when cflowd is enabled on an IP interface, the sampling rate is applied to a port and only the samples that match the IP interface for which cflowd is enabled are processed further to update or create flow records in the flow cache. Samples received that do not match the IP interface for which cflowd is enabled are not processed further, and flow records are not created for them.

The **no** form of this command reverts the sample rate to the default value.

Default

rate 1000

Parameters

sample-rate

Specifies the rate at which traffic is sampled.

Values 1 to 10000

template-retransmit

Syntax

template-retransmit *seconds*

no template-retransmit

Context

config>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command specifies the interval for sending template definitions.

Default

template-retransmit 600

Parameters

seconds

Specifies the interval, in seconds, between the sending of template definitions.

Values 10 to 600

use-vrtr-if-index

Syntax

[no] use-vrtr-if-index

Context

config>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command exports flow data using interface indexes (ifIndex values), which can be used directly as the index into the IF-MIB tables for retrieving interface statistics.

Specifically, if this command is enabled, the ingressInterface (ID=10) and egressInterface (ID= 14) fields in IP flow templates, which are used to export the flow data to cflowd Version 9 and Version 10 collectors, is populated with the IF-MIB ifIndex of that interface. In addition, for Version 10 templates, two fields are available in the IP flow templates to present the virtual router ID associated with the ingress and egress interfaces.

The **no** form of this command removes the command from the active configuration and causes cflowd to revert to the default behavior of populating the ingress and egress interface ID with the global IF index ID.

Default

no use-vrtr-if-index

5.9.2.2 Show commands

collector

Syntax

collector [*ip-address[:port]*] [**detail**]

Context

show>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays the administrative and operational status of the configured data collectors.

Parameters

ip-address

Displays information about the specified collector IP address.

Default all collectors

Values

ip-address	a.b.c.d[:port]	(IPv4)
	x:x:x:x:x:x:x	(IPv6)
	[x:x:x:x:x:x] :port	(IPv6)
	x - [0 to FFFF]H	

:port

Displays information about the collector on the specified UDP port.

Default all UDP ports

Values 1 to 65535

detail

Keyword to display informational details about either all collectors or the specified collector.

Output

The following outputs are examples of cflowd collector information, and the associated tables describe the output fields.

- Standard output: [Sample output 1, Table 69: Output fields: cflowd collector](#)
- Detailed output: [Sample output 2, Table 70: Output fields: cflowd collector detailed](#)

Sample output 1

```
A:R51-CfmA# show cflowd collector
```

```
=====
Cflowd Collectors
=====
```

Host Address	Port	Version	AS Type	Admin	Oper	Sent
138.120.135.103	2055	v5	peer	up	up	1380 records
138.120.135.103	9555	v8	origin	up	up	90 records
138.120.135.103	9996	v9	-	up	up	0 packets
138.120.214.224	2055	v5	origin	up	up	1380 records

```
-----
Collectors : 4
=====
```

Table 69: Output fields: cflowd collector

Label	Description
Host Address	Displays the IP address of a remote cflowd collector host to receive the exported cflowd data
Port	Displays the UDP port number on the remote cflowd collector host to receive the exported cflowd data
AS Type	Displays the style of AS reporting used in the exported flow data origin — Reflects the endpoints of the AS path that the flow is following peer — Reflects the AS of the previous and next hops for the flow
Version	Displays the configured version for the associated collector
Admin	Displays the desired administrative state for this cflowd remote collector host
Oper	Displays the current operational status of this cflowd remote collector host
Recs Sent	Displays the number of cflowd records that have been transmitted to this remote collector host
Collectors	Displays the total number of collectors using this IP address

Sample output 2

```

A:R51-CfmA# show cflowd collector detail
=====
Cflowd Collectors (detail)
=====
Address           : 138.120.135.103
Port              : 2055
Description       : Test v5 Collector
Version           : 5
AS Type           : peer
Admin State       : up
Oper State        : up
Records Sent      : 1260
Last Changed      : 09/03/2009 17:24:04
Last Pkt Sent     : 09/03/2009 18:07:10
-----
                                     Sent      Open      Errors
-----
                                     42       0        0
=====
Address           : 138.120.135.103
Port              : 9555
Description       : Test v8 Collector
Version           : 8
AS Type           : origin
    
```

```

Admin State      : up
Oper State      : up
Records Sent    : 82
Last Changed    : 09/03/2009 17:24:04
Last Pkt Sent   : 09/03/2009 18:06:41
-----
Aggregation Type      Status      Sent      Open      Errors
-----
as-matrix             Disabled      0         0         0
protocol-port         Disabled      0         0         0
source-prefix         Enabled       21        0         0
destination-prefix    Enabled       21        0         0
source-destination-prefix Disabled      0         0         0
raw                   Disabled      0         0         0
=====
Address              : 138.120.135.103
Port                 : 9996
Description           : Test v9 Collector
Version              : 9
Admin State          : up
Oper State           : up
Packets Sent         : 51
Last Changed         : 09/03/2009 17:24:04
Last Pkt Sent        : 09/03/2009 18:07:10
Template Set          : Basic
-----
Traffic Type          Template Sent      Sent      Open      Errors
-----
IPv4                  09/03/2009 18:07:29 51        1         0
MPLS                  No template sent    0         0         0
IPv6                  No template sent    0         0         0
=====
A:R51-CfmA#
    
```

Table 70: Output fields: cflowd collector detailed

Label	Description
Address	Displays the IP address of a remote cflowd collector host to receive the exported cflowd data
Port	Displays the UDP port number on the remote cflowd collector host to receive the exported cflowd data
Description	Displays a user-provided descriptive string for this cflowd remote collector host
Version	Displays the version of the flow data sent to the collector
AS Type	Displays the style of AS reporting used in the exported flow data origin — Reflects the endpoints of the AS path which the flow is following peer — Reflects the AS of the previous and next hops for the flow
Admin State	Displays the desired administrative state for this cflowd remote collector host

Label	Description
Oper State	Displays the current operational status of this cflowd remote collector host
Records Sent	Displays the number of cflowd records that have been transmitted to this remote collector host
Last Changed	Displays the time when this row entry was last changed
Last Pkt Sent	Displays the time when the last cflowd packet was sent to this remote collector host
Aggregation Type	Displays the bit mask that specifies the aggregation schemes used to aggregate multiple individual flows into an aggregated flow for export to this remote host collector. none — No data will be exported for this remote collector host raw — Flow data is exported without aggregation in version 5 format All other aggregation types use version 8 format to export the flow data to this remote host collector.
Collectors	Displays the total number of collectors using this IP address
Sent	Displays the number of packets with flow data sent to the associated collector
Open	Displays the number of partially filled packets that have some flow data but are not yet filled or have been timed out (60 seconds maximum)
Error	Increments when an error occurs during export of the collector packet. The most common reason is a UDP unreachable destination for the configured collector.

interface

Syntax

interface [*ip-int-name*]

Context

show>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays the administrative and operational status of the interfaces in which cflowd is enabled.

Parameters

ip-int-name

Displays only information for the specified IP interface name, up to 32 characters.

Default all interfaces with cflowd enabled

Output

The following output is an example of cflowd interface information, and [Table 71: Output fields: cflowd interface](#) describes the output fields.

Sample output

```
# show cflowd interface [ip-int-name]
=====
Cflowd Interfaces
=====
Interface          Router      IF Index Type/Dir Samp Admin
IPv4 Address              Oper IPv4
IPv6 Address              Oper IPv6
-----
test                Base        1      intf/ingr  Up
1.1.1.1/24                uni         Down
N/A                        uni         Down
-----
Interfaces : 1
```

Table 71: Output fields: cflowd interface

Label	Description
Interface	Displays the physical port identifier
IPv4 Address	Displays the primary IPv4 address for the associated IP interface
IPv6 Address	Displays the primary IPv6 address for the associated IP interface
Router	Displays the virtual router index (Base = 0)
IF Index	Displays the global IP interface index
Type/Dir Samp	Displays the cflowd sampling type and direction intf — Interface based sampling acl — ACL based sampling ingr — Ingress sampling egr — Egress sampling both — Both ingress and egress sampling

Label	Description
Admin	Displays the administrative state of the interface
Opr-IPv4	Displays the operational state for IPv4 sampling
Opr-IPv6	Displays the operational state for IPv6 sampling

status

Syntax

status

Context

show>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays administrative and operational status information for cflowd.

Output

The following output is an example of cflowd status information, and [Table 72: Output fields: cflowd status](#) describes the output fields.

Sample output

```

sr1# show cflowd status
=====
Cflowd Status
=====
Cflowd Admin Status : Enabled
Cflowd Oper Status : Enabled
Active Timeout : 1 minutes
Inactive Timeout : 30 seconds
Template Retransmit : 60 seconds
Cache Size : 65536 entries
Overflow : 1%
Sample Rate : 1
Active Flows : 34000
Overflow events 10
Dropped Flows: 0
Pkts Rcvd : 801600
Total Pkts Dropped : 0

Times flow created      Raw
Times flow matched     160000
Total flows flushed    224428382
                       150000
=====
Version Info
=====
Version      Status      Sent      Open      Errors
    
```

```

-----
5          Enabled      92      0      0
8          Enabled      46      0      0
9          Enabled      56      1      0
10         Enabled      39      1      0
=====

Cflowd Status
=====
Cflowd Admin Status : Enabled
Cflowd Oper Status  : Enabled
Active Timeout      : 1 minutes
Inactive Timeout    : 30 seconds
Template Retransmit : 60 seconds
Cache Size          : 65536 entries
Overflow            : 1%
Sample Rate         : 1
Active Flows        : 34
Total Pkts Rcvd    : 801600
Total Pkts Dropped : 0

=====

Version Info
=====
Version          Status      Sent      Open      Errors
-----
5                Enabled    92        0         0
8                Enabled    46        0         0
9                Enabled    56        1         0
10               Enabled    39        1         0
=====
    
```

Table 72: Output fields: cflowd status

Label	Description
Cflowd Admin Status	Displays the desired administrative state for this cflowd remote collector host
Cflowd Oper Status	Displays the current operational status of this cflowd remote collector host
Active Timeout	Displays the maximum amount of time, in minutes, before an active flow is exported. If an individual flow is active for this amount of time, the flow is exported and a new flow is created.
Inactive Timeout	Displays the inactive timeout in seconds
Template Retransmit	Displays the time, in seconds, before template definitions are sent
Cache Size	Displays the maximum number of active flows to be maintained in the flow cache table
Overflow	Displays the percentage number of flows to be flushed when the flow cache size has been exceeded

Label	Description
Sample Rate	Displays the rate at which traffic is sampled and forwarded for cflowd analysis one (1) — All packets are analyzed 1000 (default) — One in every one thousand packet is analyzed
Active Flows	Displays the current number of active flows being collected
Total Pkts Rcvd	Displays the total number of packets sampled and forwarded for cflowd analysis
Total Pkts Dropped	Displays the total number of packets dropped
Aggregation Info:	
Type	Displays the type of data to be aggregated and to the collector
Status	enabled — Specifies that the aggregation type is enabled disabled — Specifies that the aggregation type is disabled
Sent	Displays the number of packets with flow data sent to the associated collector
Open	Displays the number of partially filled packets which have some flow data but are not yet filled or have been timed out (60 seconds maximum)
Error	Counter increments when an error occurs during export of the collector packet. The most common reason is a UDP unreachable destination for the configured collector.
Overflow events	Displays the number of times the active cache overflowed
Dropped Flows	Displays the total number of flows dropped due to cache overflow events

5.9.2.3 Tools commands

cache

Syntax

```
cache aggregate {src-dst-proto | src-dst-proto-port} family {ipv4 | ipv6}
cached all family {ipv4 | ipv6}
```

Context

```
tools>dump>cflowd
```


Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays the contents of the cflowd active cache. This information can be displayed either in raw form, where every flow entry is displayed, or in an aggregated form.

Parameters

all

Displays the raw active cache flow data with no aggregation.

aggregate

Displays the aggregated active cache flow data.

src-dst-proto — Aggregates the active flow cache based on the source and destination IP address and the IP protocol value.

src-dst-proto-port — Aggregates the active flow cache based on the source and destination IP address, IP protocol value, and the source and destination port numbers.

family

Specifies the IP address family flow for which data should be displayed.

ipv4 — Displays the IPv4 flow data.

ipv6 — Displays the IPv6 flow data.

Output

The following output is an example of cflowd cache information, [Table 73: Output fields: tools dump cflowd cache](#) describes the output fields.

Sample output

```

INFO: |18:59:55 +00:00.153| tools dump cflowd cache all family ipv4
Current time: 08/26/2021 13:29:55
-----
Intf/Ingr SrcIP          Intf/Egr DstIP          Prot ToS Flgs  Pkts
vRtr-ID   S-Port Msk AS        D-Port Msk AS        NextHop      Pkt-Size Active
-----
1          150.1.1.2          2          150.2.1.2      17  0x00 0x00  25
1          7          /24 200        7          /24 300      150.2.1.2     46  0
-----

```

Table 73: Output fields: tools dump cflowd cache

Label	Description
Proto/Protocol	Displays the IPv4 or IPv6 protocol type
Source Address/Src-IP	Displays the source IP address of the flow (IPv4 or IPv6)
Destination Address/Dst-IP	Displays the destination IP address of the flow (IPv4 or IPv6)

Label	Description
Intf/Ingr	Displays the ingress interface associated with the sampled flow (only displayed with the raw (all) output)
Intf/Egr	Displays the egress interface associated with the sampled flow (only displayed with the raw (all) output)
S-Port	Displays the source protocol port number
D-Port	Displays the destination protocol port number
Pkt-Cnt	Displays the total number of packets sampled for the associated flow
Byte-Cnt	Displays the total number of bytes of traffic sampled for the associated flow
Start-Time	Displays the system time when the first packet was sampled for the associated flow
Flags	Displays the IP flag value from the sampled IP flow header (only displayed with the raw (all) output)
ToS	Displays the ToS byte values from the sampled IP flow header (only displayed with the raw (all) output)
(Src) Mask	Displays the IP route mask for the route to the flow source IP address associated with the flow (only displayed with the raw (all) output)
(Dst) Mask	Displays the IP route mask for the route to the flow destination IP address associated with the flow (only displayed with the raw (all) output)
(Src) AS	Displays the ASN associated with the route to the flow source IP address associated with the flow (only displayed with the raw (all) output)
(Dst) AS	Displays the ASN associated with the route to the flow destination IP address associated with the flow (only displayed with the raw (all) output)
vRtr-ID	Displays the virtual router ID associated with the reported IP flow (only displayed with the raw (all) output)

packet-size

Syntax

packet-size *protocol* [**clear**]

Context

tools>dump>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays packet size distribution for sampled IP traffic. Values are displayed in decimal format (1.0 = 100%, .500 = 50%). Separate statistics are maintained and displayed for IPv4 and IPv6 traffic.

Parameters

protocol

Displays packet size information for the specified protocol.

Values ipv4, ipv6, mcast-ipv4, mcast-ipv6

clear

Keyword to clear statistics.

Output

The following output is an example of cflowd packet size information.

Sample output

```
A:Dut-A#  
INFO: |18:57:06 +00:00.100| tools dump cflowd packet-size ipv4  
IPv4 unicast packet size distribution (30 total packets):  
    Current Time: 08/26/2021 13:27:05  
    Last Cleared Time: 08/26/2021 13:26:32  
1-32  64  96  128  160  192  224  256  288  320  352  384  416  448  480  
.033 .067 .100 .033 .033 .033 .033 .033 .033 .033 .033 .033 .033 .033 .033  
512  544  576  1024  1536  2048  2560  3072  3584  4096  4608  9000  
.033 .033 .033 .033 .033 .033 .033 .033 .033 .033 .033 .033 .033 .033
```

top-flows

Syntax

top-flows *protocols* [**clear**]

Context

tools>dump>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays the top 20 (highest traffic volume) flows for IPv4, IPv6, or MPLS traffic types collected since the cflowd top-flow table was last cleared or initialized.

Parameters

protocol

Displays top-flow information for the specified protocol.

Values ipv4, ipv6, mpls, l2, mcast-ipv4, mcast-ipv6

clear

Keyword to clear statistics.

Output

The following output is an example of cflowd top-flow information, and [Table 74: Output fields: tools dump cflowd top-flows](#) describes the output fields.

Sample output

```
INFO: |18:57:05 +00:00.086| tools dump cflowd top-flows ipv4
The top 20 IPv4 unicast flows seen by cflowd are:
  Current Time: 08/26/2021 13:27:05
  Last Cleared Time: 08/26/2021 13:26:32
  ifIndexContext: global
Intf/Ingr SrcIP          Intf/Egr DstIP          Pro ToS Flgs          Pkts
vRtr-ID  S-Port Msk AS         D-Port Msk AS      NextHop        Pkt-Size      Time
-----
1         150.1.1.2          2         150.2.1.2          6  0x00 0x00          26
1         10 /24 200         20 /24 300 150.2.1.2          1079          18
1         1.20.1.2          0         1.20.1.3           6  0xc0 0x18          1
1         179 /0 0           51201 /0 0 0.0.0.0           71            0
1         150.1.1.2          2         150.2.1.2          2  0x00 0x00          1
1         0 /24 200         0 /24 300 150.2.1.2          28            0
-----
```

Table 74: Output fields: tools dump cflowd top-flows

Label	Description
Ingress	Displays the ingress interface ID
Src IP	Displays the source IP address of the flow (IPv4 or IPv6)
Egress	Displays the egress interface ID
Dest IP	Displays the destination IP address of the flow (IPv4 or IPv6)
Pr	Displays the protocol type for flow
TOS	Displays the Type of Service/DSCP bits filed markings
Flgs	Displays the protocol flag markings

Label	Description
Pkts	Displays the total number of packets sampled for this flow (since stats were last cleared)
vRtr-ID	Displays the vRouter context the flow was sampled in
S-Port	Displays the source protocol port number
Msk	Displays the route prefix length for route to source IP address
AS	Displays the AS number for the source route (the AS is either originating or peer, depending on the cflowd configuration)
DstIP	Displays the destination protocol port number
Msk	Displays the route prefix length for route to destination IP address (Forwarding route)
AS	Displays the AS number for the destination route (the AS is either originating or peer, depending on the cflowd configuration)
Nexthop	Displays the next-hop address used to forward traffic associated with the flow
Avg pkt size	Displays the average packet size of a sampled traffic associated with this flow (total number of packets sampled / total number of packets sampled)
Active	Displays the number of seconds the flow has been active

top-protocols

Syntax

top-protocols *protocols* [**clear**]

Context

tools>dump>cflowd

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays the summary information for the top 20 protocols traffic in the cflowd cache. All statistics are calculated based on data collected since the cflowd statistics were last cleared using the **clear** keyword.

If the **clear** optional keyword is configured, the top flows are displayed and then this cache is cleared.

Parameters

protocol

Displays top protocol information for the specified protocol.

Values ipv4, ipv6, mcast-ipv4, mcast-ipv6

clear

Keyword to clear statistics.

Output

The following output is an example of cflowd top protocol traffic information, and [Table 75: Output fields: tools dump cflowd top-protocols](#) describes the output fields.

Sample output

```
A:Dut-A#
INFO: |18:57:05 +00:00.095| tools dump cflowd top-protocols ipv4
The top 20 IPv4 unicast protocols seen by cflowd are:
    Current Time: 08/26/2021 13:27:05
    Last Cleared Time: 08/26/2021 13:26:32
Protocol (ID)      Total      Flows      Pkts      Bytes      Pkts      Secs      % Total
-----          Flows      /Sec      /Flow     /Pkt      /Sec      /Flow     Bandwidth
-----
TCP                2           0          13       1041         0         9         99%
IGMP               1           0           1         28         0         0          0%
-----
TOTALS            3           0           9       1005         0         6        100%
```

Table 75: Output fields: tools dump cflowd top-protocols

Label	Description
Protocol ID	Displays the IPv4 or IPv6 protocol type Prints either the well-known protocol name or the decimal protocol number
Total Flows	Displays the total number of flows recorded since the cflowd statistics were last cleared with this protocol type
Flows/Sec	Displays the average number of flows detected for the associated protocol type (Total flows / number of seconds since last clear)
Packets/Flow	Displays the average number of packets per flow (Total number of packets / total flows)
Bytes/Pkts	Displays the average number of bytes per packet for the associated protocol type (Total number of bytes for the associated protocol / total number of packets seen for the associated protocol)
Packets/Sec	Displays the average number of packets seen for the associated protocol type

Label	Description
	(Number of packets / time since last clear)
Duration/Flow	Displays the average lifetime of a flow for the associated protocol type (Number of seconds since last clear / total flows)
Bandwidth Total (%)	Displays the percentage of bandwidth consumed by the associated protocol type (Total protocol bytes / total bytes of all flows)

5.9.2.4 Clear commands

cflowd

Syntax

cflowd

Context

clear

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command clears the raw and aggregation flow caches that are sending flow data to the configured collectors. This action triggers all flows to be discarded. The cache restarts flow data collection from a fresh state. This command also clears global statistics collector statistics listed in the cflowd **show** commands.

6 Common CLI command descriptions

This section provides information about Command Line Interface (CLI) syntax and command usage for common service commands.

6.1 Command descriptions

6.1.1 SAP syntax

sap

Syntax

[no] sap *sap-id*

Context

Various

Platforms

Supported on all 7210 SAS platforms as described in this document.

Description

This command specifies the physical port identifier portion of the SAP definition.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition.

Values The *sap-id* can be configured in one of the formats shown in the following table.

Table 76: Formats of *sap-id*

Type	Syntax	Example
port-id	<i>slot/mda/port[.channel]</i>	1/1/5
null	[<i>port-id</i> <i>lag-id</i>]	<i>port-id</i> : 1/1/3 <i>lag-id</i> : lag-3
dot1q	[<i>port-id</i> <i>lag-id</i>]:qtag1	<i>port-id</i> :qtag1: 1/1/3:100

Type	Syntax	Example
		<i>lag-id:lag-1:102</i>

qtag1, qtag2

Specifies the encapsulation value used to identify the SAP on the port or sub-port. If this parameter is not specifically defined, the default value is 0.

Values

qtag1: * | 0 to 4094

qtag2: * | 0 to 4094

The values depend on the encapsulation type configured for the interface. The following table describes the allowed values for the port and encapsulation types.

Table 77: Port and encapsulation types

Port type	Encap-type	Allowed values	Comments
Ethernet	Null	0	The SAP is identified by the port.
Ethernet	Dot1q	0 to 4094	The SAP is identified by the 802.1Q tag on the port. Note that a 0 qtag1 value also accepts untagged packets on the dot1q port.

7 Standards and protocol support



Note:

The information provided in this chapter is subject to change without notice and may not apply to all platforms.

Nokia assumes no responsibility for inaccuracies.

The following conventions are used in this section:

- T(A,N) indicates 7210 SAS-T in both Access-uplink mode and Network mode. Similarly, T(N) indicates 7210 SAS-T in network mode only.
- K5 indicates 7210 SAS-K 2F2T1C.
- K12 indicates 7210 SAS-K 2F4T6C.
- K30 indicates 7210 SAS-K 3SFP+ 8C.
- Sx/S-1/10GE indicates all variants of 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE platforms.
- Sx-10/100GE indicates only the variants of 7210 SAS-Sx 10/100GE. Sx-10/100GE QSFP28 indicates the 7210 SAS-Sx 10/100GE 64 SFP+ 4QSFP28 variant.
- Sx/S-1/10GE-VC indicates 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE platforms in standalone VC mode.
- R6 indicates 7210 SAS-R6.
- R12 indicates 7210 SAS-R12.
- D indicates 7210 SAS-D and 7210 SAS-D ETR. If a line item applies only to 7210 SAS-D ETR, then it is indicated as D-ETR.
- Dxp indicates 7210 SAS-Dxp-12p (2SFP+ 4F6T) 7210 SAS-Dxp-12p ETR (2SFP+ 4F6T), 7210 SAS-Dxp 16p (2SFP+ 4F10T), and 7210 SAS-Dxp-24p (2SFP+ 6F16T). If a line item applies only to a particular variant, the variant name will be called out explicitly against that item.
- This standards list is not applicable to platforms in the satellite mode of operation, as most of the features are supported on 7x50 SR platforms. For this reason, the host platforms standards compliance must be consulted for the satellite mode of operation.

7.1 BGP

draft-ietf-idr-add-paths-04, Advertisement of Multiple Paths in BGP is supported on T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-sidr-origin-validation-signaling-04, BGP Prefix Origin Validation State Extended Community is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1772, Application of the Border Gateway Protocol in the Internet is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1997, BGP Communities Attribute is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2385, Protection of BGP Sessions via the TCP MD5 Signature Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2439, BGP Route Flap Damping is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2545, Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2858, Multiprotocol Extensions for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2918, Route Refresh Capability for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3107, Carrying Label Information in BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3392, Capabilities Advertisement with BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4271, A Border Gateway Protocol 4 (BGP-4) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4360, BGP Extended Communities Attribute is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4456, BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only R6 and R12 supports RR server functionality. Rest of the platforms support only client function.

RFC 4659, BGP/MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4684, Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4724, Graceful Restart Mechanism for BGP (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4760, Multiprotocol Extensions for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4798, Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4893, BGP Support for Four-octet AS Number Space is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5004, Avoid BGP Best Path Transitions from One External to Another is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5291, Outbound Route Filtering Capability for BGP-4 is supported on K12, K30, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5668, 4-Octet AS Specific BGP Extended Community is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6811, Prefix Origin Validation is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.2 Ethernet

IEEE 802.1AB, Station and Media Access Control Connectivity Discovery is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

TIA-1057, LLDP for Media endpoint devices is supported on Dxp, Sx/S-1/10GE, and Sx/S-1/10GE-VC

IEEE 802.1ad, Provider Bridges is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1ag, Connectivity Fault Management is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1ah, Provider Backbone Bridges is supported on T(N)

IEEE 802.1ax, Link Aggregation is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1D, MAC Bridges is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1p, Traffic Class Expediting is supported on D, E, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1Q, Virtual LANs is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1s, Multiple Spanning Trees is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1w, Rapid Reconfiguration of Spanning Tree is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1X, Port Based Network Access Control is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ab, 1000BASE-T is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ac, VLAN Tag is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ad, Link Aggregation is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ae, 10 Gb/s Ethernet is supported on Dxp, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ah, Ethernet in the First Mile is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ba, 40 Gb/s and 100 Gb/s Ethernet is supported on R6, R12, and Sx-10/100GE

IEEE 802.3i, Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3u, Fast Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3z, Gigabit Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3af, Power Over Ethernet (PoE) is supported on Dxp, T-ETR, Mxp-ETR, and Sx/S-1/10GE



Note:

Sx/S-1/10GE only on PoE variant and Sx-1/10GE fiber variant with two fixed copper ports. Only on Dxp-16p and Dxp-24p.

IEEE 802.3at, Power Over Ethernet (PoE+) is supported on Dxp, T-ETR, Mxp-ETR, and Sx/S-1/10GE



Note:

Sx/S-1/10GE only on PoE variant and Sx-1/10GE fiber variant with two fixed copper ports. Only on Dxp-16p and Dxp-24p.

IEEE 802.3bt, Power Over Ethernet (PoE++/HPoE) is supported on Dxp



Note:

Only on Dxp-16p and Dxp-24p.

ITU-T G.8032, Ethernet Ring Protection Switching is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ITU-T Y.1731, OAM functions and mechanisms for Ethernet based networks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.3 EVPN

draft-snr-bess-evpn-proxy-arp-nd-00, Proxy-ARP/ND function in EVPN networks is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12



Note:

Sx/S-1/10GE standalone mode only.

RFC 7432, BGP MPLS-Based Ethernet VPN is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12



Note:

Sx/S-1/10GE standalone mode only.

draft-ietf-bess-evpn-vpws-14, Virtual Private Wire Service support in Ethernet VPN is supported on Mxp

7.4 Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, Operational management of Loop Free Alternates is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:
With Segment Routing.

RFC 5286, Basic Specification for IP Fast Reroute: Loop-Free Alternates is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:
With Segment Routing.

RFC 7490, Remote Loop-Free Alternate (LFA) Fast Reroute (FRR) is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:
With Segment Routing.

7.5 Internet Protocol (IP) — General

draft-grant-tacacs-02, The TACACS+ Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-rrp-unified-spec-02, Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6 is supported on Mxp

RFC 768, User Datagram Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 793, Transmission Control Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 854, TELNET Protocol Specifications is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 951, Bootstrap Protocol (BOOTP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1034, Domain Names - Concepts and Facilities is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1035, Domain Names - Implementation and Specification is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1350, The TFTP Protocol (revision 2) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1534, Interoperation between DHCP and BOOTP is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1542, Clarifications and Extensions for the Bootstrap Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2131, Dynamic Host Configuration Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2347, TFTP Option Extension is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2348, TFTP Blocksize Option is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2349, TFTP Timeout Interval and Transfer Size Options is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2428, FTP Extensions for IPv6 and NATs is supported on D, Dxp, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2865, Remote Authentication Dial In User Service (RADIUS) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2866, RADIUS Accounting is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3046, DHCP Relay Agent Information Option (Option 82) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3596, DNS Extensions to Support IP version 6 is supported on D, Dxp, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3768, Virtual Router Redundancy Protocol (VRRP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4250, The Secure Shell (SSH) Protocol Assigned Numbers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4251, The Secure Shell (SSH) Protocol Architecture is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4252, The Secure Shell (SSH) Authentication Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

All 7210 platforms support password and publickey based user authentication. 7210 SAS-D supports only password based authentication.

RFC 4253, The Secure Shell (SSH) Transport Layer Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4254, The Secure Shell (SSH) Connection Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4632, Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5880, Bidirectional Forwarding Detection (BFD) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 5881, Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 5883, Bidirectional Forwarding Detection (BFD) for Multihop Paths is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 6528, Defending against Sequence Number Attacks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7130, Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces is supported on T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

7.6 IP — Multicast

RFC 1112, Host Extensions for IP Multicasting is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IGMP v1, v2, v3 is supported.

RFC 2236, Internet Group Management Protocol, Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3306, Unicast-Prefix-based IPv6 Multicast Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3376, Internet Group Management Protocol, Version 3 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3446, Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

MSDP supported only on Sx/S-1/10GE standalone.

RFC 3618, Multicast Source Discovery Protocol (MSDP) is supported on Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 4601, Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4604, Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

MLD not supported.

RFC 4607, Source-Specific Multicast for IP is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4608, Source-Specific Protocol Independent Multicast in 232/8 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4610, Anycast-RP Using Protocol Independent Multicast (PIM) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5059, Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5384, The Protocol Independent Multicast (PIM) Join Attribute Format is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6513, Multicast in MPLS/BGP IP VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6514, BGP Encodings and Procedures for Multicast in MPLS/IP VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6515, IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6625, Wildcards in Multicast VPN Auto-Discover Routes is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6826, Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 7246, Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 7385, IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

7.7 IP — Version 4

RFC 791, Internet Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 792, Internet Control Message Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 826, An Ethernet Address Resolution Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1519, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1812, Requirements for IPv4 Routers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1981, Path MTU Discovery for IP version 6 is supported on K12, K30, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2401, Security Architecture for Internet Protocol is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Supported only for OSPFv3 authentication. Not supported for services.

RFC 2460, Internet Protocol, Version 6 (IPv6) Specification is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

7.8 IP — Version 6

RFC 2464, Transmission of IPv6 Packets over Ethernet Networks is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 3021, Using 31-Bit Prefixes on IPv4 Point-to-Point Links is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3122, Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification is supported on K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3587, IPv6 Global Unicast Address Format is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4007, IPv6 Scoped Address Architecture is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4193, Unique Local IPv6 Unicast Addresses is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4291, Internet Protocol Version 6 (IPv6) Addressing Architecture is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4443, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4861, Neighbor Discovery for IP version 6 (IPv6) is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4862, IPv6 Stateless Address Autoconfiguration (Router Only) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5095, Deprecation of Type 0 Routing Headers in IPv6 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5952, A Recommendation for IPv6 Address Text Representation is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 6106, IPv6 Router Advertisement Options for DNS Configuration is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6164, Using 127-Bit IPv6 Prefixes on Inter-Router Links is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.9 IPsec

RFC 2401, Security Architecture for the Internet Protocol is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only for use with OSPFv3 authentication. Not supported for services.

RFC 2406, IP Encapsulating Security Payload (ESP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only for use with OSPFv3 authentication. Not supported for services.

7.10 IS-IS

draft-ietf-isis-mi-02, IS-IS Multi-Instance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-segment-routing-extensions-04, IS-IS Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-kaplan-isis-ext-eth-02, Extended Ethernet Frame Size Support is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ISO/IEC 10589:2002, Second Edition, Nov. 2002, Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3359, Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3719, Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3787, Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4971, Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5120, M-ISIS: Multi Topology (MT) Routing in IS-IS is supported on K12, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5130, A Policy Control Mechanism in IS-IS Using Administrative Tags is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5301, Dynamic Hostname Exchange Mechanism for IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5302, Domain-wide Prefix Distribution with Two-Level IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5303, Three-Way Handshake for IS-IS Point-to-Point Adjacencies is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5304, IS-IS Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5305, IS-IS Extensions for Traffic Engineering TE is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5306, Restart Signaling for IS-IS (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5308, Routing IPv6 with IS-IS is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5310, IS-IS Generic Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6232, Purge Originator Identification TLV for IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6233, IS-IS Registry Extension for Purges is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-mi-02, IS-IS Multi-Instance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

K12, K30 support only a single instance and can operate in multi-instance deployment as it supports the processing of TLVs for multi-instance support.

draft-ietf-isis-segment-routing-extensions-04, IS-IS Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

7.11 Management

draft-ietf-snmv3-update-mib-05, Management Information Base (MIB) for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-idr-bgp4-mib-05, Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-wg-mib-06, Management Information Base for Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-mboned-msdp-mib-01, Multicast Source Discovery protocol MIB is supported on Sx/S-1/10GE



Note:

Only in standalone mode.

draft-ietf-mpls-ldp-mib-07, Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-mpls-lsr-mib-06, Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-mpls-te-mib-04, Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-mib-update-08, OSPF Version 2 Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ianaaddressfamilynumbers-mib, IANA-ADDRESS-FAMILY-NUMBERS-MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ianaiftype-mib, IANAifType-MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ianaiprouteprotocol-mib, IANA-RTPROTO-MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE8021-CFM-MIB, IEEE P802.1ag(TM) CFM MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE8021-PAE-MIB, IEEE 802.1X MIB is supported on D, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE8023-LAG-MIB, IEEE 802.3ad MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

LLDP-MIB, IEEE P802.1AB(TM) LLDP MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1157, A Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1215, A Convention for Defining Traps for use with the SNMP is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1724, RIP Version 2 MIB Extension is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2021, Remote Network Monitoring Management Information Base Version 2 using SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2138, Remote Authentication Dial In User Service (RADIUS) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2206, RSVP Management Information Base using SMIv2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2213, Integrated Services Management Information Base using SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2571, An Architecture for Describing SNMP Management Frameworks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2573, SNMP Applications is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2575, View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP) is supported on D, E, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2578, Structure of Management Information Version 2 (SMIv2) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2579, Textual Conventions for SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2787, Definitions of Managed Objects for the Virtual Router Redundancy Protocol is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2819, Remote Network Monitoring Management Information Base is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2856, Textual Conventions for Additional High Capacity Data Types is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2863, The Interfaces Group MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2864, The Inverted Stack Table Extension to the Interfaces Group MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2933, Internet Group Management Protocol MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3014, Notification Log MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3164, The BSD syslog Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3165, Definitions of Managed Objects for the Delegation of Management Scripts is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3231, Definitions of Managed Objects for Scheduling Management Operations is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3273, Remote Network Monitoring Management Information Base for High Capacity Networks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3414, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3416, Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3417, Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3419, Textual Conventions for Transport Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3584, Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3593, Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals is supported on Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3635, Definitions of Managed Objects for the Ethernet-like Interface Types is supported on D, E, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3826, The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3877, Alarm Management Information Base (MIB) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4001, Textual Conventions for Internet Network Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4022, Management Information Base for the Transmission Control Protocol (TCP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4113, Management Information Base for the User Datagram Protocol (UDP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4220, Traffic Engineering Link Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4292, IP Forwarding Table MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4293, Management Information Base for the Internet Protocol (IP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5101, Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information is supported on Mxp, Sx/S-1/10GE, and R6

RFC 6241, Network Configuration Protocol (NETCONF) is supported on K5, K12, T(N), Mxp, Sx/S-1/10GE, R6, and R12

RFC 6242, Using the NETCONF Protocol over Secure Shell (SSH) is supported on K5, K12, T(N), Mxp, Sx/S-1/10GE, R6, and R12

7.12 MPLS — General

RFC 3031, Multiprotocol Label Switching Architecture is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3032, MPLS Label Stack Encoding is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3443, Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4182, Removing a Restriction on the use of MPLS Explicit NULL is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5332, MPLS Multicast Encapsulations is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12

7.13 MPLS — GMPLS

draft-ietf-ccamp-rsvp-te-srlg-collect-04, RSVP-TE Extensions for Collecting SRLG Information is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.14 MPLS — LDP

draft-ietf-mpls-ldp-ipv6-15, Updates to LDP for IPv6 is supported on Mxp

draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-pdutta-mpls-ldp-v2-00, LDP Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-pdutta-mpls-ldp-hello-reduce-04, Targeted LDP Hello Reduction is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3037, LDP Applicability is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5036, LDP Specification is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5443, LDP IGP Synchronization is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5561, LDP Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
P2MP LSPs only.

7.15 MPLS — MPLS-TP

RFC 5586, MPLS Generic Associated Channel is supported on T(N), R6, and R12

RFC 5921, A Framework for MPLS in Transport Networks is supported on T(N), R6, and R12

RFC 5960, MPLS Transport Profile Data Plane Architecture is supported on T(N), R6, and R12

RFC 6370, MPLS Transport Profile (MPLS-TP) Identifiers is supported on T(N), R6, and R12

RFC 6378, MPLS Transport Profile (MPLS-TP) Linear Protection is supported on T(N), R6, and R12

RFC 6426, MPLS On-Demand Connectivity and Route Tracing is supported on T(N), R6, and R12

RFC 6428, Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile is supported on T(N), R6, and R12

RFC 6478, Pseudowire Status for Static Pseudowires is supported on T(N), R6, and R12

RFC 7213, MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing is supported on T(N), R6, and R12

7.16 MPLS — OAM

RFC 6424, Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6425, Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping is supported on T(N), Mxp, R6, and R12

7.17 MPLS — RSVP-TE

RFC 2702, Requirements for Traffic Engineering over MPLS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2747, RSVP Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2961, RSVP Refresh Overhead Reduction Extensions is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3097, RSVP Cryptographic Authentication -- Updated Message Type Value is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3209, RSVP-TE: Extensions to RSVP for LSP Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3477, Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4090, Fast Reroute Extensions to RSVP-TE for LSP Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4561, Definition of a Record Route Object (RRO) Node-Id Sub-Object is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4875, Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs) is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4950, ICMP Extensions for Multiprotocol Label Switching is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5817, Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.18 OSPF

draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1765, OSPF Database Overflow is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2328, OSPF Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3101, The OSPF Not-So-Stubby Area (NSSA) Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3509, Alternative Implementations of OSPF Area Border Routers is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3623, Graceful OSPF Restart Graceful OSPF Restart (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3630, Traffic Engineering (TE) Extensions to OSPF Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4222, Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4552, Authentication/Confidentiality for OSPFv3 is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4576, Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4577, OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4970, Extensions to OSPF for Advertising Optional Router Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5185, OSPF Multi-Area Adjacency is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5187, OSPFv3 Graceful Restart (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5243, OSPF Database Exchange Summary List Optimization is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5250, The OSPF Opaque LSA Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5340, OSPF for IPv6 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5838, Support of Address Families in OSPFv3 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6987, OSPF Stub Router Advertisement is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

7.19 Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, OAM Procedures for VPWS Interworking is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3916, Requirements for Pseudo- Wire Emulation Edge-to-Edge (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3985, Pseudo Wire Emulation Edge-to-Edge (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4385, Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4446, IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4447, Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4448, Encapsulation Methods for Transport of Ethernet over MPLS Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5659, An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6073, Segmented Pseudowire is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6310, Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6391, Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network is supported on K12, K30, Mxp, R6, and R12

RFC 6718, Pseudowire Redundancy is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6870, Pseudowire Preferential Forwarding Status bit is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7023, MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7267, Dynamic Placement of Multi-Segment Pseudowires is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.20 Quality of Service

RFC 2430, A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2598, An Expedited Forwarding PHB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3140, Per Hop Behavior Identification Codes is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3260, New Terminology and Clarifications for Diffserv is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.21 RIP

RFC 1058, Routing Information Protocol is supported on Mxp and Sx/S-1/10GE



Note:
Only in standalone mode.

RFC 2082, RIP-2 MD5 Authentication is supported on Mxp and Sx/S-1/10GE



Note:
Only in standalone mode.

RFC 2453, RIP Version 2 is supported on Mxp and Sx/S-1/10GE



Note:
Only in standalone mode.

7.22 Timing

GR-1244-CORE, Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:
For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

GR-253-CORE, SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:
For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

IEEE 1588-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, IEEE default profile is supported only includes the Dxp-12p ETR, Dxp-16p, Dxp-24p. Sx-10/100GE does not support IEEE default profile.

ITU-T G.781, Synchronization layer functions, issued 09/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.813, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8261, Timing and synchronization aspects in packet networks, issued 04/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8262, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8264, Distribution of timing information through packet networks, issued 10/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8265.1, Precision time protocol telecom profile for frequency synchronization, issued 10/2010 is supported on D-ETR, K5, K12, K30, T(A,N), Mxp, Sx-1/10GE, R6, and R12

ITU-T G.8275.1, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014 is supported on Dxp, K12, K30, T(A,N), Mxp, Sx-1/10GE, Sx-10/100GE, R6, and R12



Note:

For 7210 SAS-Sx 10/100GE, the support only includes the Sx 10/100GE QSFP28 variant. For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

IEC/IEEE 61850-9-3-2016, Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation is supported on Dxp-16p and Dxp-24p

IEEE C37.238-2017 - IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications is supported on Dxp-16p and Dxp-24p

RFC 5905, Network Time Protocol Version 4: Protocol and Algorithms Specification is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.23 VPLS

RFC 4761, Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

On 7210 platforms, only BGP-AD is supported with TLDP signalling for PW. No BGP signalling is supported for PW establishment.

RFC 4762, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5501, Requirements for Multicast Support in Virtual Private LAN Services is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6074, Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)