



7210 SERVICE ACCESS SYSTEM

7210 SAS-D, DXP, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C ROUTER CONFIGURATION GUIDE

RELEASE 22.9.R1

3HE 18207 AAAB TQZZA
Issue 01

September 2022

© 2022 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.

© 2022 Nokia.

Table of contents

List of tables.....	8
List of figures.....	11
1 Getting started.....	12
1.1 About this guide.....	12
1.1.1 Document structure and content.....	12
1.2 7210 SAS modes of operation.....	13
1.3 7210 SAS port modes.....	15
1.4 7210 SAS-series router configuration process.....	17
1.5 Conventions.....	18
1.5.1 Precautionary and information messages.....	18
1.5.2 Options or substeps in procedures and sequential workflows.....	18
2 IP router configuration.....	20
2.1 Configuring IP router parameters.....	20
2.1.1 Interfaces.....	20
2.1.1.1 Network interface on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	20
2.1.1.2 System interface on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	20
2.1.1.3 System interface on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	20
2.1.1.4 Unicast Reverse Path Forwarding check on 7210 SAS-K 3SFP+ 8C.....	21
2.1.1.5 Unnumbered interfaces.....	21
2.1.1.6 Router ID.....	22
2.1.1.7 Autonomous systems (AS).....	22
2.1.1.8 Proxy ARP.....	23
2.1.1.9 Internet Protocol versions.....	23
2.1.1.9.1 IPv6 applications for 7210 SAS-D and 7210 SAS-Dxp.....	25
2.1.1.9.2 IPv6 applications for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	26
2.1.1.9.3 IPv6 Provider Edge router over MPLS (6PE).....	26
2.1.1.9.4 DNS.....	27
2.1.1.10 Bidirectional forwarding detection for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	27
2.1.1.10.1 BFD control packet.....	28
2.1.1.10.2 Control packet format.....	28
2.1.1.10.3 BFD echo support.....	29

2.1.10.4 BFD support on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C platforms.....	30
2.1.11 DHCP.....	30
2.1.11.1 DHCP principles.....	31
2.1.12 DHCP relay.....	33
2.1.13 DHCP relay agent options.....	33
2.1.13.1 Option 82.....	33
2.1.13.2 Local DHCP server on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	34
2.1.13.3 DHCP server options.....	35
2.1.13.4 Trusted and untrusted.....	36
2.1.13.5 DHCP snooping.....	36
2.1.14 IGP-LDP and static route-LDP synchronization on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	36
2.2 Process overview on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	37
2.3 Process overview on the 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	37
2.4 Configuration notes.....	38
2.4.1 Configuration guidelines for DHCP relay and snooping.....	38
2.5 Configuring an IP router with CLI.....	39
2.5.1 Router configuration overview of 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T....	39
2.5.1.1 System interface on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	39
2.5.2 Router configuration overview on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	39
2.5.2.1 System interface on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	40
2.5.2.2 Network interface.....	40
2.5.3 Basic configuration.....	40
2.5.4 Common configuration tasks.....	41
2.5.4.1 Configuring a system name.....	41
2.5.4.2 Configuring interfaces.....	41
2.5.4.3 Configuring an unnumbered interface.....	44
2.5.4.4 Router advertisement on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.....	44
2.5.4.5 Configuring proxy ARP.....	45
2.5.4.6 ECMP considerations.....	46
2.5.4.7 Deriving the router ID.....	47
2.5.4.8 Configuring an autonomous system.....	47
2.5.4.9 Configuring Option 82 handling.....	48
2.5.4.10 Configuring a local DHCP server.....	49
2.5.5 Service management tasks.....	50
2.5.5.1 Changing the system name.....	50

2.5.5.2	Modifying interface parameters.....	51
2.5.5.3	Deleting a logical IP interface.....	51
2.6	IP router command reference.....	52
2.6.1	Command hierarchies.....	52
2.6.1.1	Configuration commands.....	53
2.6.1.2	Show commands.....	58
2.6.1.3	Clear commands.....	59
2.6.1.4	Debug commands.....	60
2.6.2	Command descriptions.....	60
2.6.2.1	Configuration commands.....	60
2.6.2.2	Show commands.....	138
2.6.2.3	Clear commands.....	187
2.6.2.4	Debug commands.....	193
3	Filter policies.....	198
3.1	Filter policy configuration overview.....	198
3.1.1	Service-based filtering.....	198
3.1.2	Filter policy entities.....	199
3.1.2.1	Applying filter policies.....	199
3.1.2.2	ACL on range SAPs.....	201
3.2	Creating and applying filter policies.....	204
3.2.1	Packet matching criteria.....	205
3.2.1.1	DSCP values.....	207
3.2.2	Ordering filter entries.....	209
3.2.3	Applying filters.....	211
3.2.3.1	Applying a filter to a SAP.....	211
3.2.3.2	Applying a filter to an IES interface.....	212
3.2.3.3	Applying a filter to a network IP interface.....	212
3.3	Configuration notes.....	212
3.3.1	MAC filters.....	213
3.3.2	IP filters.....	214
3.3.3	IPv6 filters.....	214
3.3.3.1	Resource usage for ingress filter policies for 7210 SAS-D and 7210 SAS-Dxp.....	214
3.3.3.2	Resource usage for egress filter policies (supported only for 7210 SAS-D and 7210 SAS-Dxp).....	216
3.3.4	Ingress filter policy resource usage: 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.....	217

3.4 Configuring filter policies with CLI.....	217
3.5 Basic configuration.....	218
3.6 Common configuration tasks.....	219
3.6.1 Allocating resources for filter policies (ingress and egress).....	219
3.6.2 Creating an IP filter policy.....	220
3.6.2.1 IP filter policy.....	220
3.6.2.2 IP filter entry.....	220
3.6.2.3 IP entry matching criteria.....	221
3.6.3 Creating an IPv6 filter policy (applicable only for 7210 SAS-D and 7210 SAS-Dxp).....	221
3.6.3.1 IPv6 filter entry.....	221
3.6.4 Creating a MAC filter policy.....	222
3.6.4.1 MAC filter policy.....	222
3.6.4.2 MAC filter entry.....	222
3.6.4.3 MAC entry matching criteria.....	223
3.6.4.4 Apply IP and MAC filter policies.....	223
3.6.4.5 Apply filter policies to an IES interface.....	224
3.7 Filter management tasks.....	224
3.7.1 Renumbering filter policy entries.....	224
3.7.2 Modifying an IP filter policy.....	226
3.7.3 Modifying a MAC filter policy.....	227
3.7.4 Deleting a filter policy.....	228
3.7.4.1 From an ingress SAP.....	228
3.7.4.2 From an egress SAP.....	228
3.7.4.3 From the filter configuration.....	229
3.7.5 Copying filter policies.....	229
3.8 Filter command reference.....	230
3.8.1 Command hierarchies.....	230
3.8.1.1 Configuration commands.....	230
3.8.1.2 Show commands.....	234
3.8.1.3 Clear commands.....	234
3.8.1.4 Monitor commands.....	235
3.8.2 Command descriptions.....	235
3.8.2.1 Configuration commands.....	235
3.8.2.2 Show commands.....	268
3.8.2.3 Clear commands.....	286
3.8.2.4 Monitor commands.....	289

4 Common CLI command descriptions.....	293
4.1 In this chapter.....	293
4.2 Common service commands.....	293
4.2.1 SAP commands.....	293
4.2.1.1 SAP command description.....	293
5 Standards and protocol support.....	295
5.1 BGP.....	295
5.2 Ethernet.....	297
5.3 EVPN.....	298
5.4 Fast Reroute.....	298
5.5 Internet Protocol (IP) — General.....	299
5.6 IP — Multicast.....	300
5.7 IP — Version 4.....	302
5.8 IP — Version 6.....	303
5.9 IPsec.....	304
5.10 IS-IS.....	304
5.11 Management.....	306
5.12 MPLS — General.....	309
5.13 MPLS — GMPLS.....	309
5.14 MPLS — LDP.....	309
5.15 MPLS — MPLS-TP.....	310
5.16 MPLS — OAM.....	310
5.17 MPLS — RSVP-TE.....	311
5.18 OSPF.....	311
5.19 Pseudowire.....	312
5.20 Quality of Service.....	313
5.21 RIP.....	314
5.22 Timing.....	314
5.23 VPLS.....	315

List of tables

Table 1: Supported modes of operation and configuration methods.....	14
Table 2: Supported port modes by mode of operation.....	16
Table 3: 7210 SAS platforms supporting port modes.....	17
Table 4: Configuration process.....	18
Table 5: IPv6 header field descriptions.....	25
Table 6: BFD control packet field descriptions.....	28
Table 7: Default route preferences.....	72
Table 8: Output fields: ARP.....	140
Table 9: Output fields: neighbor.....	142
Table 10: Output fields: declined addresses.....	145
Table 11: Output fields: free addresses.....	146
Table 12: Output fields: lease.....	147
Table 13: Output fields: server statistics.....	149
Table 14: Output fields: subnet extended statistics.....	152
Table 15: Output fields: subnet statistics.....	153
Table 16: Output fields: DHCP server summary.....	155
Table 17: Output fields: DHCP servers.....	157
Table 18: Output fields: DHCP statistics.....	159
Table 19: Output fields: DHCP summary.....	160
Table 20: Output fields: DHCP statistics.....	162
Table 21: Output fields: router ECMP.....	163

Table 22: Output fields: FIB BGP PIC.....	165
Table 23: Output fields: ICMP6.....	166
Table 24: Output fields: ICMP6 interface.....	167
Table 25: Output fields: router interface.....	170
Table 26: Output fields: router interface detail.....	171
Table 27: Output fields: route table.....	174
Table 28: Output fields: route table summary.....	174
Table 29: Output fields: route-table BGP PIC.....	175
Table 30: Output fields: router advertisement.....	179
Table 31: Output fields: static ARP.....	181
Table 32: Output fields: static route.....	184
Table 33: Output fields: router status.....	185
Table 34: Output fields: tunnel table.....	187
Table 35: Applying filter policies for 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	200
Table 36: Applying filter policies for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	200
Table 37: Applying ACLs support on Epipe and VPLS services on 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C variants when using range SAPs.....	201
Table 38: Packet fields for match in QoS classification policy and ACL policy.....	202
Table 39: DSCP name to DSCP value table.....	207
Table 40: MAC match criteria exclusivity rules.....	214
Table 41: IP protocol IDs and descriptions.....	243
Table 42: 3-bit mask format.....	259
Table 43: 48-bit mask format.....	260

Table 44: 48-bit mask format.....	265
Table 45: Output fields: filter IP.....	269
Table 46: Output fields: filter IP with filter ID specified.....	271
Table 47: Output fields: filter IP associations.....	274
Table 48: Output fields: filter IP counters.....	275
Table 49: Output fields: filter IPv6.....	277
Table 50: Output fields: filter IPv6 with filter ID specified.....	278
Table 51: Output fields: filter IPv6 associations.....	280
Table 52: Output fields: filter IPv6 counters.....	281
Table 53: Output fields: MAC filter.....	283
Table 54: Output fields: filter MAC counters.....	285
Table 55: Output fields: filter MAC associations.....	286
Table 56: Formats of sap-id.....	294
Table 57: Port and encapsulation types.....	294

List of figures

Figure 1: IPv6 header format.....	25
Figure 2: Example of a 6PE topology within one AS.....	26
Figure 3: IP address assignment with DHCP.....	32
Figure 4: Creating and applying filter policies.....	205
Figure 5: Filtering process example.....	211
Figure 6: Applying an IP filter to an ingress interface.....	219

1 Getting started

This chapter provides process flow information to configure routing entities, VRRP, IP and MAC filters. It also provides an overview of the document organization and content, and describes the terminology used in this guide.

1.1 About this guide

This guide describes the logical IP routing interfaces, VRRP, 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.



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.

- 7210 SAS-D
- 7210 SAS-Dxp 12p (2SFP+ 4SFP 6Tx)
- 7210 SAS-Dxp 16p (2SFP+ 4SFP 10Tx)
- 7210 SAS-Dxp 24p (2SFP+ 6SFP 16Tx)
- 7210 SAS-K 2F1C2T
- 7210 SAS-K 2F6C4T
- 7210 SAS-K 3SFP+ 8C

See section [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.

- **access-uplink mode of operation**

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

- **network mode of operation**

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

When configured in the access-uplink or network mode of operation, the 7210 SAS platform implicitly operates in the standalone mode.

1.1.1 Document structure and content

This guide uses the following structure to describe features and configuration content.



Note:

This guide generically covers Release 22.x.Rx content and may include some content that will be released in later maintenance loads. See the *7210 SAS Software Release Notes 22.x.Rx*, part number 3HE 18217 000x TQZZA, for information about features supported in each load of the Release 22.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 22.x.Rx*, part number 3HE 18217 000x TQZZA, and 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 specific *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-K 2F1C2T		Implicit	Implicit		

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

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

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

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 specific *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 1CFP4 or IMM-c 1QSFP28)	Yes	Yes	Yes	No
7210 SAS-R12 IMM-b	Yes	Yes	Yes	No
7210 SAS-R12 IMM-c 100GE (IMM-c 1CFP4 or IMM-c 1QSFP28)	Yes	Yes	Yes	No
7210 SAS-Sx 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-series 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 and addresses and router IDs.	IP router configuration
	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 are information symbols 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 device, 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.

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 Network interface on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

2.1.2 System interface on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

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 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 System interface on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The system interface is associated with the network entity (such as a specific router or switch), and not a specific interface.

The system interface is associated during the configuration of the following entities:

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

2.1.4 Unicast Reverse Path Forwarding check on 7210 SAS-K 3SFP+ 8C

The Unicast Reverse Path Forwarding Check (Unicast RPF) feature helps mitigate problems caused by the introduction of malformed or forged (spoofed) IP source addresses into a network. The feature works by discarding IP packets that lack a verifiable IP source address. For example, common types of denial-of-service (DoS) attacks, including smurf and tribe flood network (TFN), can use forged or rapidly changing source addresses to thwart efforts to locate or filter the attacks. ISPs that provide public access can use Unicast RPF to deflect such attacks by only forwarding packets with source IP addresses that are valid and consistent with the IP routing table. This protects the network of the ISP, its customer, and the rest of the Internet.

Unicast RPF is supported for both IPv4 and IPv6 on access ports only. It is supported on any IP interface configured in the IES and VPRN services.

Unicast RPF has two modes: strict and loose, but the 7210 SAS-K 3SFP+ 8C supports only strict mode in this release.

In strict mode, Unicast RPF checks whether there is a matching prefix entry for the source address of the incoming packet in the routing table, and whether the interface expects to receive a packet with this source address prefix. If **urpf-check** is enabled on the interface, all interfaces are assumed to be enabled for strict mode Unicast RPF check.

In the case of ECMP, the 7210 SAS-K 3SFP+ 8C allows a packet received on an IP interface configured in strict mode Unicast RPF to be forwarded, if the IP interface on which the packet is received matches any one of the interfaces used by that ECMP route.

If there is a default route, the following is included in the Unicast RPF check:

- If there is a default route, a strict mode Unicast RPF check only succeeds if the source address matches any route (including the default route) where the next-hop is on the incoming interface for the packet. A per node option exists to ignore the default route for a strict mode Unicast RPF check.
- If a match is not found, the Unicast RPF check fails.

If the source IP address matches a discard/blackhole route, the packet is treated as if it failed the Unicast RPF check.

2.1.5 Unnumbered interfaces

Unnumbered interfaces are point-to-point interfaces that are not explicitly configured with a dedicated IP address and subnet; instead, they borrow (or link to) an IP address from another interface on the system (the system IP address, another loopback interface, or any other numbered interface) and use it as the source IP address for packets originating from the interface.

The benefits of using unnumbered interfaces are the following:

- ISP backhaul can be enabled with a single IP address allocated to the CE nodes (the network interface address is coupled with the system IP address).
- Nodes can be added to or deleted from a network without address changes; unnumbered interfaces are linked to a centralized IP address and therefore do not require any address change if the nodes are

relocated. After a topology change, the ARP table is updated to ensure reachability, and the upper layer protocols re-establish the peering sessions.

Unnumbered interfaces are supported on:

- network interfaces
- IES interfaces
- VPRN interfaces

Only IPv4 addresses are supported.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, unnumbered interfaces are supported for the IS-IS and OSPF routing protocols and for MPLS routing (RSVP-TE and LDP). This feature is supported through both dynamic and static ARP. Any Ethernet port with null, dot1q, or QinQ encapsulation supports IP unnumbered interfaces.

See the *7210 SAS-K 2F6C4T, K 3SFP+ 8C Routing Protocols Guide* for more information about IS-IS and OSPF unnumbered interface support. See “Unnumbered Point-to-Point Interface in RSVP” and “Unnumbered Interface Support in LDP” in the *7210 SAS-K 2F6C4T, K 3SFP+ 8C MPLS Guide* for more information about MPLS unnumbered interface support.



Note:

Unnumbered interfaces do not support PIM routing or IGMP listener capabilities.

2.1.6 Router ID



Note:

This feature is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

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, the system interface acts as the router ID.
- If neither the system interface or router ID are implicitly specified, the router ID is inherited from the last four bytes of the MAC address.
- The router can be derived on the protocol level.

2.1.7 Autonomous systems (AS)

**Note:**

- AS is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.
- 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 AS that have been administratively assigned to the same group. An area's 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.

2.1.8 Proxy ARP

**Note:**

This feature is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

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 supports 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 Nokia 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 specific 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.

2.1.9 Internet Protocol versions



Note:

IPv4 and IPv6 support on the different platforms is as follows:

- IPv6 is supported on all 7210 SAS platforms as described in this document, except the 7210 SAS-K 2F1C2T.
- The 7210 SAS-D and 7210 SAS-Dxp platforms only support the use of IPv6 for management purposes. IPv6 cannot be used to deliver a service.
- The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C platforms can be used as dual-stack IPv6 and IPv4 routers capable of IPv6 forwarding and the provision of IPv6 services, including IPv6 VPN (6VPE) services. IPv6 support can also be used for management of the node (in-band management is available).

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 anycast address is defined that is used to send a packet to any one of a group of nodes.



Note:

The 7210 SAS-Dxp supports a maximum 64-bit prefix length for IPv6 addresses. This restriction applies when configuring static routes; for example, a static route can be configured with a /64 prefix using the **configure router static-route 2001::0/64 next-hop 4001::5** command.

- **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. This field uses the same values as the IPv4 protocol field.
Hop Limit	8-bit unsigned integer. Decrement 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).

2.1.9.1 IPv6 applications for 7210 SAS-D and 7210 SAS-Dxp

The IPv6 applications for 7210 SAS-D and 7210 SAS-Dxp are:

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

2.1.9.2 IPv6 applications for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C


The IPv6 applications for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C are:

- IPv6 services
- IPv6 dual-stack router with capability to route/forward IPv6 packets
- IPv6 inband management of the node

2.1.9.3 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 with one AS.

Figure 2: Example of a 6PE topology within one AS



Fig_30

2.1.9.3.1 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.

2.1.9.3.2 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.

2.1.9.4 DNS

The DNS client is extended to use IPv6 as transport and to handle the IPv6 address in the DNS AAAA resource record from an IPv4 or IPv6 DNS server. An assigned name can be used instead of an IPv6 address as IPv6 addresses are more difficult to remember than IPv4 addresses.

2.1.10 Bidirectional forwarding detection for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 and 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.

2.1.10.1 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 networks is specified in *draft-ietf-bfd-v4v6-1hop-04.txt, BFD for IPv4 (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.

2.1.10.2 Control packet format

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



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's reason for the last transition of the session from Up to some other state. Possible values are:

Field	Description
	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	
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.
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.

2.1.10.3 BFD 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's 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.

2.1.10.4 BFD support on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C platforms

BFD support on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C platforms is as follows:

BFD in a VPRN service can be used for:

- OSPFv2 PE-CE routing protocol
- Static routes (IPv4)
- BGP for PE-CE protocol (IPv4)

BFD in IES service can be used for:

- OSPFv2
- IS-IS for IPv4 interfaces
- Static routes (IPv4)

BFD in Base routing instance can be used for:

- OSPFv2 on network IPv4 interfaces
- IS-IS on network IPv4 interfaces
- MP-BGP for vpn-ipv4 families (only multi-hop)
- Static routes (IPv4)
- RSVP-TE
- TLDP (IPv4)
- Interface LDP (link-level) (IPv4)



Note:

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, BFD processing is supported in hardware enabling faster detection (minimum timer supported is 10ms).

2.1.11 DHCP



Note:

DHCP server support on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C is designed to be used for IP address assignment used for local management access to the node or to the devices connected to the node for maintenance activities.

DHCP is a configuration protocol used to communicate network information and configuration parameters from a DHCP server to a DHCP-aware client. DHCP is based on the BOOTP protocol, with additional configuration options and the added capability of allocating dynamic network addresses. DHCP-capable devices are also capable of handling BOOTP messages.

A DHCP client is an IP-capable device (typically a computer or base station) that uses DHCP to obtain configuration parameters such as a network address. A DHCP server is an Internet host or router that returns configuration parameters to DHCP clients. A DHCP/BOOTP Relay agent is a host or router that passes DHCP messages between clients and servers.

Home computers in a residential high-speed Internet application typically use the DHCP protocol to have their IP address assigned by their Internet service provider.

The following is supported on different 7210 SAS platforms:

- The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C platforms can act as a DHCP Relay agent, or a local DHCP server.
- The 7210 SAS-K 2F1C2T platform can act as a DHCP relay agent only.
- The 7210 SAS-D and 7210 SAS-Dxp platforms can act as a DHCP relay agent only.

The following paragraphs describe the functionality available on 7210 SAS as DHCP server, and as a relay agent.

For DHCP, the DHCP protocol requires the client to transmit a request packet with a destination broadcast address of 255.255.255.255 that is processed by the DHCP server. Because IP routers do not forward broadcast packets, this would suggest that the DHCP client and server must reside on the same network segment. However, for various reasons, it is sometimes impractical to have the server and client reside in the same IP network. When the 7210 is acting as a DHCP Relay agent, it processes these DHCP broadcast packets and relays them to a preconfigured DHCP server. Therefore, DHCP clients and servers do not need to reside on the same network segment.

When the 7210 SAS is acting as a local DHCP server, it processes these DHCP broadcast packets and allocates IP addresses for the DHCP client as needed.

2.1.11.1 DHCP principles



Note:


The references to spoke-SDP and mesh-SDP in this section are only applicable to the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

In a Triple Play network, client devices (such as a routed home gateway, a session initiation protocol (SIP) phone or a set-top box) use DHCP to dynamically obtain their IP address and other network configuration information. The 7210 SAS auto-init procedure also uses DHCP to dynamically obtain the BOF used for first-time booting of the system (along with IP address required to retrieve the BOF, the configuration file and the TIMOS software image from the network). DHCP is defined and shaped by several RFCs and drafts in the IETF DHC working group including the following:

- RFC 2131, Dynamic Host Configuration Protocol
- RFC 3046, DHCP Relay Agent Information Option

The DHCP operation is shown in the following figure.

Figure 3: IP address assignment with DHCP



1. During boot-up, the client device sends a DHCP discover message to get an IP address from the DHCP Server. The message contains the following:

- destination MAC address - broadcast
- source MAC address - MAC of client device
- client hardware address - MAC of client device

If this message passes through a DSLAM or other access node (possibly a 7210 SAS device), typically the Relay information option (Option 82) field is added, indicating shelf, slot, port, VPI, VCI and other fields, to identify the subscriber.

DHCP relay is enabled on the first IP interface in the upstream direction. Depending on the scenario, the DSLAM, 7210 SAS access node, or the BSR relays the discover message as a unicast packet toward the configured DHCP server. DHCP relay is configured to insert the giaddr to indicate to the DHCP server in which subnet an address should be allocated.

2. The DHCP server looks up the client MAC address and Option 82 information in its database. If the client is recognized and authorized to access the network, an IP address is assigned and a DHCP offer message returned. The BSA or BSR relays this back to the client device.
3. It is possible that the discover reached more than one DHCP server, and therefore that more than one offer was returned. The client selects one of the offered IP addresses and confirms it needs to use this in a DHCP request message, sent as unicast to the DHCP server that offered it.
4. The DHCP server confirms that the IP address is still available, updates its database to indicate it is now in use, and replies with a DHCP ACK message back to the client. The ACK also contains the Lease Time of the IP address.

2.1.12 DHCP relay

The 7210 SAS provides DHCP/BOOTP Relay agent services for DHCP clients. DHCP is used for IPv4 network addresses. DHCP is known as stateful protocols because they use dedicated servers to maintain parameter information. In the stateful auto-configuration model, hosts obtain interface addresses and configuration information and parameters from a server. The server maintains a database that keeps track of which addresses have been assigned to which hosts.

DHCP relay on different 7210 SAS platforms is as follows:

- 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T support DHCP Relay on the base router, and on access IP interfaces associated with IES service used for management.
- 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C support DHCP Relay on the base router, and on access IP interfaces associated with IES service and VPRN service.

2.1.13 DHCP relay agent options

DHCP options are codes that the router inserts in packets being forwarded from a DHCP client to a DHCP server. Some options have more information stored in sub-options.

The 7210 SAS supports Option 60 and Option 61 as specified in RFC 2132. Option 60 is the vendor class identifier, which can contain information such as the client's hardware configuration. Option 61 is the client identifier.

The 7210 SAS supports the Relay Agent Information Option 82 as specified in RFC3046. The following sub-options are supported for the base router:

- action
- circuit ID
- copy-82
- remote ID

2.1.13.1 Option 82

Option 82, or the relay information option is specified in RFC 3046, DHCP Relay Agent Information Option, allows the router to append some information to the DHCP request that identifies where the original DHCP request arrives from.

There are two sub-options under Option 82:

- Agent Circuit ID Sub-option (RFC 3046, section 3.1): This sub-option specifies data which must be unique to the box that is relaying the circuit.
- Remote ID Sub-option (RFC 3046 section 3.2): This sub-option identifies the host at the other end of the circuit. This value must be globally unique.

Both sub-options are supported by the 7210 SAS and can be used separately or together.

Inserting Option 82 information is supported independently of DHCP relay.

When the circuit ID sub-option field is inserted by the 7210 SAS, it can take following values:

- **sap-id**

This is the SAP index (only under an IES service).

- **ifindex**

This is the index of the IP interface (only under an IES service).

- **ascii-tuple**

This is an ASCII-encoded concatenated tuple, consisting of [system-name|serviceid| interface-name] (for IES) or [system-name|service-id|sap-id] (for VPLS).

- **vlan-ascii-tuple**

This is an ASCII-encoded concatenated tuple, consisting of the ascii-tuple followed by dot1p bits and dot1q tags.

When a DHCP packet is received with Option 82 information already present, the system can do one of three things. The available actions are:

- **Replace**

On ingress the existing information-option is replaced with the information-option parameter configured on the 7210 SAS. On egress (toward the customer) the information-option is stripped (per the RFC).

- **Drop**

The DHCP packet is dropped and a counter is incremented.

- **Keep**

The existing information is kept on the packet and the router does not add any more information. On egress the information option is not stripped and is sent on to the downstream node.

In accordance with the RFC, the default behavior is to keep the existing information; except if the giaddr of the packet received is identical to a local IP address on the router, then the packet is dropped and an error incremented regardless of the configured action.

The maximum packet size for a DHCP relay packet is 1500 bytes. If adding the Option 82 information would cause the packet to exceed this size, the DHCP relay request is forwarded without the Option 82 information. This packet size limitation exists to ensure that there is no fragmentation on the end Ethernet segment where the DHCP server attaches.

In the downstream direction, the inserted Option 82 information should not be passed back toward the client (as per RFC 3046, DHCP Relay Agent Information Option). To enable downstream stripping of the option 82 field, DHCP snooping should be enabled on the SDP or SAP connected to the DHCP server.

2.1.13.2 Local DHCP server on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C supports local DHCP server functionality on the base router and on access IP interfaces associated with VPRN, by dynamically assigning IPv4 addresses to access devices that request them. This standards-based, full DHCP server implementation allows a service provider the option. The 7210 SAS can support public and private addressing in the same router, including overlapped private addressing in the form of VPRNs in the same router. The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C acts as a DHCP server.

An administrator creates pools of addresses that are available for assigned hosts. Locally attached hosts can obtain an address directly from the server. Routed hosts receive addresses through a relay point in the customer's network. When a DHCP server receives a DHCP message from a DHCP Relay agent, the server looks for a subnet to use for assigning an IP address. If configured with the **use-pool-from-client** command, the server searches Option 82 information for a pool name. If a pool name is found,

an available address from any subnet of the pool is offered to the client. If configured with the **use-gi-address** command, the server uses the gateway IP address (GIADDR) supplied by the Relay agent to find a matching subnet. If a subnet is found, an address from the subnet is offered to the client. If no pool or subnet is found, then no IP address is offered to the client.

IPv4 address assignments are temporary and expire when the configured lease time is up. The server can reassign addresses after the lease expires.

If both the **no use-pool-from-client** command and the **no use-gi-address** command or **no use-link-address** command are specified, the server does not act.

2.1.13.3 DHCP server options

Options and identification strings can be configured on several levels.

DHCP servers support the following options, as defined in RFC 2132:

- Option 1 - Subnet Mask
- Option 3 - Default Routers
- Option 6 - DNS Name Servers
- Option 12 - Host Name
- Option 15 - Domain Name
- Option 44 - Netbios Name Server
- Option 46 - Netbios Node Type Option
- Option 50 - IP Address
- Option 51 - IP Address Lease Time
- Option 53 - DHCP Message Type
- Option 54 - DHCP Server IP Address
- Option 55 - Parameter Request List
- Option 58 - Renew (T1) Timer
- Option 59 - Renew (T2) Timer
- Option 60 - Class Identifier
- Option 61 - Client Identifier

DHCP servers also support Sub-option 13 Relay Agent Information Option 82 as specified in RFC 3046, to enable the use of a pool indicated by the DHCP client.

These options are copied into the DHCP reply message, but if the same option is defined several times, the order of priority is the following:

1. subnet option
2. pool options
3. options from the DHCP client request

A local DHCP server must be bound to a specified interface by referencing the server from that interface. The DHCP server is then addressable by the IP address of that interface. A normal interface or a loop-back interface can be used.

A DHCP client is defined by the MAC address and the circuit identifier. This implies that for a specific combination of MAC and circuit identifier, only one IP address can be returned; if more than one request is made, the same address is returned.

2.1.13.4 Trusted and untrusted

There is a case where the relay agent could receive a request where the downstream node added Option 82 information without also adding a giaddr (giaddr of 0). In this case the default behavior is for the router to drop the DHCP request. This behavior is in line with the RFC.

The 7210 SAS supports a command `trusted`, which allows the router to forward the DHCP request even if it receives one with a giaddr of 0 and Option 82 information attached. This could occur with older access equipment. In this case the relay agent would modify the request's giaddr to be equal to the ingress interface. This only makes sense when the action in the information option is `keep`, and the service is IES. In the case where the Option 82 information gets replaced by the relay agent, either through explicit configuration or the VPLS DHCP Relay case, the original Option 82 information is lost, and the reason for enabling the trusted option is lost.

2.1.13.5 DHCP snooping

To support DHCP based address assignment in L2 aggregation network, 7210 supports DHCP snooping. 7210 can copy packets designated to the standard UDP port for DHCP (port 67) to its control plane for inspection, this process is called DHCP snooping.

DHCP snooping can be performed in two directions:

- From the client to the DHCP server (Discover or Request messages) to insert Option 82 information; For these applications, DHCP snooping must be enabled on the SAP toward the subscriber.
- From the DHCP server (ACK messages), to remove the Option 82 field toward the client. For these applications, DHCP snooping must be enabled on both the SAP toward the network and the SAP toward the subscriber.

2.1.14 IGP-LDP and static route-LDP synchronization on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 provides LDP with 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 may not be up, which means that the LDP SDP remains down. In this case, the MPLS traffic is 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 re-advertised. IGP announces a new best next-hop and LDP uses 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.

2.2 Process overview on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 IP address.

- **autonomous system**

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

2.3 Process overview on the 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

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 (if supported by the platform).

2.4 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.
- On 7210 SAS-D and 7210 SAS-Dxp 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 CLI command **config>system>resource-profile>max-ipv6-routes**. 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. Please see the following example and the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide for more information.
- On 7210 SAS-D, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, a separate route table (or a block in the route table) is used for IPv6 /128-bit prefix route lookup. A limited amount of IPv6 /128 prefixes route lookup entries is supported. The software enables lookups in this table by default (that is, no user configuration is required to enable IPv6 /128-bit route lookup).
- On 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, IPv6 interfaces are allowed to be created without allocating IPv6 route entries. With this only IPv6 hosts on the same subnet are reachable.

2.4.1 Configuration guidelines for DHCP relay and snooping

The following configuration guidelines must be followed to configure DHCP relay and snooping:

- 7210 SAS devices do not support the ARP populate based on the DHCP lease, assigned to the DHCP client
- 7210 SAS devices do not maintain the DHCP lease assigned to the client
- 7210 SAS devices do not perform IP spoofing checks and MAC spoofing checks based on the DHCP parameters assigned to the client
- MAC learning must be enabled in the VPLS service, for DHCP snooping.
- DHCP snooping is not supported for B-SAPs in B-VPLS services and I-SAPs in I-VPLS services.
- Ingress ACLs cannot be used to drop DHCP control packet.
- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, DHCP packets received over an SDP are identified and option-82 inserted by the node can be removed by the node, in the downstream

direction. SAP or a SDP, as applicable. DHCP snooping configuration on an SDP is supported only on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

2.5 Configuring an IP router with CLI

This section provides information to configure an IP router.

2.5.1 Router configuration overview of 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

In 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 long, 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 router, the basic configuration tasks that must be performed are as follows:

- Assign a name to the interface.
- Associate an IP address with the interface.
- Associate the interface with a system or a loop-back interface.

A system interface should be configured.

2.5.1.1 System interface on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

The system interface is associated with the network entity, not a specific interface.

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.

2.5.2 Router configuration overview on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

In a 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, an interface is a logical named entity. An interface is created by specifying an interface name under the **configure>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 long, 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 router, the basic configuration tasks that must be performed 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.

2.5.2.1 System interface on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The system interface is associated with the network entity, 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:

- 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 used as the router identifier. A system interface must have an IP address with a 32-bit subnet mask.

2.5.2.2 Network interface

**Note:**

Network port and network IP interfaces are only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

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

2.5.3 Basic configuration

The most basic router configuration must have the following:

- system name
- system address

Example: Router configuration for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
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#
```

2.5.4 Common configuration tasks

The following sections describe the basic system configuration tasks.

2.5.4.1 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#
```

Example: 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
-----
```

2.5.4.2 Configuring interfaces

The following command sequences create a system IP interface.

**Note:**

The system interface cannot be deleted.

2.5.4.2.1 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]}
```

Example: IP configuration output showing interface information

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

2.5.4.2.2 Configure a network interface on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Use the following to configure a network interface on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

```
config>router
  interface interface-name
    address ip-addr{/mask-length | mask} [broadcast {all-ones | host-ones}]
    egress
      filter ip ip-filter-id
    ingress
      filter ip ip-filter-id
    port port-name
```

Example: 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#
```

2.5.4.2.3 Configuring IPv6 parameters on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

On the 7210 SAS-D and 7210 SAS-Dxp IPv6 interfaces with static routing can be configured. On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, IPv6 interfaces can be configured with static routing or using routing protocols such as OSPFv3 and IS-IS Ipv6 can be configured.

Before configuring the use of IPv6, system resource must be allocated for IPv6 routes on the 7210 SAS-D and 7210 SAS-Dxp using the following command.

```
configure>system>resource-profile>router> max-ipv6-routes num-routes
```

Example

The following is a sample interface configuration 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::/64
      no dad-disable
      no reachable-time
      no neighbor-limit
      no qos-route-lookup
      no local-proxy-nd
      no tcp-mss
-----
```

Example

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
```

Example: Configuration 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#
```

2.5.4.3 Configuring an unnumbered interface

Use the following syntax to configure an unnumbered interface.

```
config>router  
  interface interface-name  
    unnumbered [ip-int-name | ip-address]
```

Example

```
config>router> interface "to-ALU-3"  
config>router>if# unnumbered "system"  
config>router>if# exit
```

2.5.4.4 Router advertisement on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C



Note:

This feature is supported on all 7210 SAS platforms as described in this document, except the 7210 SAS-K 2F1C2T.

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
```

Example: Router advertisement configuration output

```
*A:sim131>config>router>router-advert# info  
-----  
  interface "n1"  
    prefix 2001:db8:3::/64
```

```

        exit
        no shutdown
    exit
-----
*A:sim131>config>router>router-advert# interface n1
*A:sim131>config>router>router-advert>if# prefix 2001:db8:3::/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#

```

2.5.4.5 Configuring proxy ARP



Note:

This feature is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

To configure proxy ARP, you can configure:

- 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 as follows:
 - In the policy statement **entry>to** context, specify the host source address(es) 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. For more information about route policies, see the *7210 SAS-D, Dxp, K 2F1C2T Routing Protocols Guide* and *7210 SAS-K 2F6C4T, K 3SFP+ 8C Routing Protocols Guide*.
- 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)]

```

Example: 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
```

Example: 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#
```

2.5.4.6 ECMP considerations



Note:

- These ECMP considerations only apply to the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.
- IPv4 ECMP is supported.
- LDP LSR and LDP LER ECMP are supported.

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

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

2.5.4.6.1 Configuration notes

IPv6 ECMP is not supported. Only a single IPv6 route for an 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).

2.5.4.7 Deriving the router ID

**Note:**

This feature is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

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]
```

Example: 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#
```

2.5.4.8 Configuring an autonomous system

**Note:**

This feature is supported only on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

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

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

Example: 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#
```

2.5.4.9 Configuring Option 82 handling

Option 82, or "Relay Information Option" is a field in DHCP messages used to identify the subscriber. The Option 82 field can already be filled in when a DHCP message is received at the router, or it can be empty. If the field is empty, the router should add identifying information (circuit ID, remote ID or both). If the field is not empty, the router can decide to replace it.

Example

The following is a sample of a partial BSA configuration with Option 82 adding on a VPLS service. Note that snooping must be enabled explicitly on a SAP or a SDP, as applicable. DHCP snooping configuration on an SDP is supported only on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

```
*A:7210SAS>config>service#
-----
vpls 2 customer 1 create
  shutdown
  stp
    shutdown
  exit
sap 1/1/12:100 create
  dhcp //Configuration example to add option 82
    option
      action replace
      circuit-id
      no remote-id
```

```

        exit
        no shutdown
    exit
    exit
    no shutdown
    exit
-----
*A:7210SAS>config>service#

The following example displays an example of a partial BSA configuration to remove
the Option 82 on a VPLS service.

vpls 2 customer 1 create
    stp
        shutdown
    exit
    sap 1/1/14:100 create //Configuration example to remove option 82
        dhcp
            snoop
            no shutdown
        exit
    exit

```

2.5.4.10 Configuring a local DHCP server

Local DHCP servers provide a standards-based full DHCP server implementation which allows a service provider the option of decentralizing IP address management into the network. Local DHCP servers is supported for IP address assignment when a front-facing port on the 7210 SAS is used for local craft terminal access.

Example

The following CLI syntax shows an example of configuring a local DHCP server ("DHCP").

```

config>service# vprn 100 customer 2 create
config>service>vprn$ autonomous-system 65535
config>service>vprn$ route-distinguisher 100:100
config>service>vprn$ interface "WAN" create
config>service>vprn>if$ address 192.168.17.1/31
config>service>vprn>if$ sap 1/1/1:101 create
config>service>vprn>if>sap$ exit
config>service>vprn>if$ exit

config>service>vprn$ dhcp
config>service>vprn>dhcp$ local-dhcp-server "DHCP" create
config>service>vprn>dhcp>server$ description "DHCP-T0-CUSTOMER"
config>service>vprn>dhcp>server$ use-gi-address
config>service>vprn>dhcp>server$ pool "DHCP-POOL" create
config>service>vprn>dhcp>server>pool$ subnet 10.0.0.0/24 create
config>service>vprn>dhcp>server>pool>subnet$ address-range 10.0.0.5 10.0.0.50
config>service>vprn>dhcp>server>pool>subnet$ exit
config>service>vprn>dhcp>server>pool$ exit
config>service>vprn>dhcp>server$ no shutdown
config>service>vprn>dhcp>server$ exit
config>service>vprn>dhcp$ exit

```

Example

The following CLI syntax shows an example of mapping the configured local DHCP server ("DHCP") to an IP address, and configuring the customer-facing interface ("LAN") to relay DHCP requests to the local DHCP server.

```
config>service>vprn$ interface "DHCP" create
config>service>vprn>if$ address 10.2.2.2/32
config>service>vprn>if$ local-dhcp-server "DHCP"
config>service>vprn>if$ loopback
config>service>vprn>if$ exit

config>service>vprn$ interface "LAN" create
config>service>vprn>if$ address 10.0.0.1/31
config>service>vprn>if$ dhcp
config>service>vprn>if>dhcp$ server 10.2.2.2
config>service>vprn>if>dhcp$ trusted
config>service>vprn>if>dhcp$ no shutdown
config>service>vprn>if>dhcp$ exit
config>service>vprn>if$ exit
```

2.5.5 Service management tasks

This section discusses the service management tasks.

2.5.5.1 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: Command usage to change the system name

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

Example: System name change 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#
```

2.5.5.2 Modifying interface parameters

Starting at the **config>router** level, navigate down to the router interface context. Follow the steps shown in the examples in this section to modify an IP address and a port.

Example: Modifying 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: Modifying 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
```

Example: 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#
```

2.5.5.3 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#
```

2.6 IP router command reference

2.6.1 Command hierarchies

- [Configuration commands](#)
 - [Router commands for 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T](#)
 - [Router commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
 - [Router interface commands for 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T](#)
 - [Router interface commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
 - [Router DHCP local user database commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
 - [Router interface IPv6 commands \(supported only on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C\)](#)
 - [IPv6 router advertisement commands \(supported only on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C\)](#)
- [Show commands](#)
 - [Router show commands](#)
 - [DHCP show commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
- [Clear commands](#)
 - [Router clear commands](#)
 - [DHCP clear commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
- [Debug commands](#)

2.6.1.1 Configuration commands

2.6.1.1.1 Router commands for 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

```

config
- router [router-name]
  - interface interface-name
  - no interface interface-name
  - router-id ip-address
  - no router-id
  - [no] static-route {ip-prefix/prefix-length | ip-prefix netmask}
[preference preference] [metric metric] [enable | disable] next-hop ip-address
  - [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] [tag tag] [enable | disable] next-hop ip-int-name
| ip-address [{cpe-check cpe-ip-address [interval seconds] [drop-count count] [log]} |
{prefix-list prefix-list-name [all | none]}] [description description]
  - [no] triggered-policy

```

2.6.1.1.2 Router commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```

config
- router [router-name]
  - allow-icmp-redirect
  - no allow-icmp-redirect
  - allow-icmp6-redirect
  - no allow-icmp6-redirect
  - autonomous-system autonomous-system
  - no autonomous-system
  - bgp
  - no bgp
  - dhcp
  - ecmp max-ecmp-routes
  - no ecmp
  - 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
  - interface ip-int-name
  - isis
  - no isis
  - mpls-labels
    - static-label-range static-range
    - no static-label-range
    - sr-labels start start-value end end-value
    - no sr-labels
  - ospf
  - no ospf
  - ospf3
  - no ospf3
  - 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 | tunnel}
- no nh-type
- protection-type {link | node}
- no protection-type
- [no] srlg-enable
- router-id ip-address
- no router-id
- [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 [bfd-enable | {cpe-check cpe-ip-address [interval seconds] [drop-count count]
[log]} | {prefix-list prefix-list-name [all | none]}] [ldp-sync] [description description]
- [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]} | {prefix-list prefix-
list-name [all | none]}] [description description]
- [no] static-route {ip-prefix/prefix-length | ip-prefix netmask}
[preference preference] [metric metric] [tag tag] [enable | disable] black-hole [prefix-
list prefix-list-name [all | none]}] [description description]
- [no] static-route {ip-prefix/prefix-length | ip-prefix netmask} black-hole
[preference preference] [metric metric] [tag tag] [enable | disable] [description description]
- no interface interface-name
- [no] triggered-policy

```

```

- router
- sgt-qos
- 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

```

**Note:**

See the 7210 SAS-K 2F1C2T, K 2F6C4T, K 3SFP+ 8C *Quality of Service Guide* section "Self-Generated Traffic Commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C" for information about self-generated traffic and applicable command descriptions.

2.6.1.1.3 Router interface commands for 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

```

config
- router [router-name]
- [no] interface interface-name
- address {ip-address/mask | ip-address netmask} [broadcast {all-ones | host-
ones}]
- no address
- delayed-enable
- no delayed-enable
- description long-description-string
- no description
- icmp
- redirects [number seconds]
- no redirects
- ttl-expired [number seconds]
- no ttl-expired

```

- **unreachables** *[number seconds]*
- **no unreachables**
- **[no] loopback**
- **[no] shutdown**

2.6.1.1.4 Router interface commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```

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 interface-name
    - 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
    - delayed-enable
    - no delayed-enable
    - description long-description-string
    - no description
    - egress
      - filter ip ip-filter-id
      - no filter
    - icmp
      - [no] mask-reply
      - redirects [number seconds]
      - no redirects
      - ttl-expired [number seconds]
      - no ttl-expired
      - unreachables [number seconds]
      - no unreachables
    - ingress
      - filter ip ip-filter-id
      - 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)]
    - [no] remote-proxy-arp
    - [no] shutdown
    - srlg-group group-name [group-name...(up to 5 max)]
    - [no] srlg-group group-name
    - static-arp ip-address ieee-address
    - static-arp ieee-address unnumbered
    - no static-arp ip-address
    - no static-arp unnumbered
    - no unnumbered [ip-int-name | ip-address]

```

- no unnumbered
- urpf-check
- [no] ignore-default

2.6.1.1.5 Router DHCP local user database commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```

config
- router
  - dhcp
    - local-dhcp-server server-name [create]
    - no local-dhcp-server server-name
      - description description-string
      - no description
      - [no] force-renews
      - lease-hold-time [lease-hold-time]
      - no lease-hold-time
      - pool pool-name [create]
      - no pool pool-name
        - description description-string
        - no description
        - max-lease-time [max-lease-time]
        - no max-lease-time
        - min-lease-time [min-lease-time]
        - no min-lease-time
        - minimum-free minimum-free [percent] [event-when-depleted]
        - no minimum-free
        - [no] nak-non-matching-subnet
        - offer-time [min minutes] [sec seconds]
        - no offer-time
        - options
          - custom-option option-number address [ip-address (up to 4 max)]
          - custom-option option-number hex hex-string
          - custom-option option-number string ascii-string
          - no custom-option option-number
          - dns-server [ip-address (up to 4 max)]
          - domain-name domain-name
          - no domain-name
          - lease-rebind-time [lease-rebind-time]
          - no lease-rebind-time
          - lease-renew-time [lease-renew-time]
          - no lease-renew-time
          - lease-time [lease-time]
          - no lease-time
        - subnet {ip-address/mask | ip-address netmask} [create]
        - no subnet {ip-address/mask | ip-address netmask}
          - [no] address-range start-ip-address end-ip-address
          - [no] exclude-addresses start-ip-address [end-ip-address]
          - maximum-declined maximum-declined
          - no maximum-declined
          - minimum-free minimum-free [percent] [event-when-depleted]
          - no minimum-free
          - options
            - custom-option option-number address [ip-address...(up to 4 max)]
            - custom-option option-number hex hex-string
            - custom-option option-number string ascii-string
            - no custom-option option-number
            - default-router ip-address [ip-address...(up to 4 max)]
            - no default-router
            - subnet-mask ip-address
            - no subnet-mask
    
```

```

- use-gi-address [scope scope]
- no use-gi-address
- user-db local-user-db-name
- no user-db

```

2.6.1.1.6 Router interface IPv6 commands (supported only on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C)

```

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
        - unreachablees [number seconds]
        - no unreachablees
      - 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
    - urpf-check ipv6
      - [no] ignore-default

```

2.6.1.1.7 IPv6 router advertisement commands (supported only on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C)

```

config
- router [router-name]
  - router-advertisement
  - no router-advertisement
  - interface ip-int-name
  - no interface ip-int-name
    - current-hop-limit number
    - no current-hop-limit
    - managed-configuration
    - no managed-configuration
    - max-advertisement-interval seconds
    - no max-advertisement-interval
    - min-advertisement-interval seconds
    - no min-advertisement-interval
    - mtu
    - no mtu
    - other-stateful-configuration
    - no other-stateful-configuration
    - prefix ipv6-prefix/prefix-length
    - no prefix

```

```

- autonomous
- no autonomous
- on-link
- 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
- shutdown

```

2.6.1.2 Show commands

2.6.1.2.1 Router 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]
  - bgp
  - ecmp
  - fib fib slot-number [family] [ip-prefix/prefix-length] [longer] [secondary]
  - fib fib slot-number [family] [summary]
  - fib fib slot-number [nh-table-usage]
  - icmp6
  - interface [[ip-address | ip-int-name]
  - interface {[ip-address | ip-int-name] [detail] [family]} | summary | exclude-
services]
  - interface [ip-address | ip-int-name] statistics
  - isis
  - neighbor [ip-address | ip-int-name | mac ieee-mac-address | summary] [dynamic |
static | managed]
  - ospf
  - policy
  - route-table [family] [ip-prefix[/prefix-length] [longer | exact]] [protocol protocol-
name | [summary]
  - rsvp
  - rtr-advertisement [interface interface-name] [prefix ipv6-prefix[/prefix-length]]
  - 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 summary [ipv4 | ipv6]
  - tunnel-table [protocol protocol] [ipv4 | ipv6]
  - tunnel-table [ip-prefix [/mask]] [alternative] [ipv4 | ipv6] detail
  - tunnel-table [ip-prefix [/mask]] [alternative]
  - tunnel-table [ip-prefix [/mask]] protocol protocol
  - tunnel-table [ip-prefix [/mask]] sdp sdp-id

```

**Note:**

For descriptions of the **show>router>sgt-qos** commands, see the "Network QoS Policy Command Reference, Show Commands" section in the *7210 SAS-K 2F1C2T, K 2F6C4T, K 3SFP + 8C Quality of Service Guide*.

2.6.1.2.2 DHCP show commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
show
- router
  - dhcp
    - local-dhcp-server server-name
      - declined-addresses ip-address[/mask] [detail]
      - declined-addresses pool pool-name
      - free-addresses ip-address[/mask]
      - free-addresses summary [subnet ip-address[/mask]]
      - free-addresses pool pool-name
      - leases [detail]
      - leases ip-address[/mask] address-from-user-db [detail]
      - leases ip-address[/mask] dhcp-host dhcp-host-name [detail]
      - leases ip-address[/mask] [detail] [state]
      - server-stats
      - subnet-ext-stats ip-address[/mask]
      - subnet-ext-stats pool pool-name
      - subnet-stats ip-address[/mask]
      - subnet-stats pool pool-name
      - summary
    - servers
    - servers all
    - statistics [interface ip-int-name | ip-address]
    - summary
```

2.6.1.3 Clear commands

2.6.1.3.1 Router clear commands

```
clear
- router [router-instance]
  - arp {all | ip-addr | interface {ip-int-name | ip-addr}}
  - 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]
```

2.6.1.3.2 DHCP clear commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
clear
- router
  - dhcp
    - local-dhcp-server server-name
```

```
- declined-addresses ip-address[/mask]  
- declined-addresses pool pool-name  
- leases ip-address[/mask] [state]  
- leases all [state]  
- server-stats  
- statistics [ip-int-name | ip-address]
```

2.6.1.4 Debug commands

```
debug  
- trace  
- router router-instance  
  - ip  
    - [no] arp  
    - [no] icmp  
    - icmp6 [ip-int-name]  
    - no icmp6  
    - 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
```

2.6.2 Command descriptions

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

2.6.2.1 Configuration commands

- [Generic commands](#)
- [Router global commands](#)
- [Router DHCP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
- [Route next-hop policy commands](#)
- [Router interface commands](#)
- [Router interface filter commands](#)
- [Router interface ICMP commands](#)
- [Interface attribute commands](#)
- [Router interface IPv6 commands](#)
- [IPv6 router advertisement commands](#)

2.6.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

config>router>interface

config>router>router-advertisement>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description



Note:

The **config>router>router-advertisement>interface** context is not supported on the 7210 SAS-K 2F1C2T.

This command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

The **no** form of this command puts an entity into the administratively enabled state.

Default

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 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.

2.6.2.1.2 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.

allow-icmp-redirect

Syntax

allow-icmp-redirect

no allow-icmp-redirect

Context

config>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables or disables ICMP redirects received on the management interface.

The **no** form of this command disables ICMP redirects.

Default

no allow-icmp-redirect

allow-icmp6-redirect

Syntax

allow-icmp6-redirect

no allow-icmp6-redirect

Context

config>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables or disables IPv6 ICMP redirects received on the management interface.

The **no** form of this command disables IPv6 ICMP redirects.

Default

no allow-icmp6-redirect

autonomous-system

Syntax

autonomous-system autonomous-system

no autonomous-system

Context

config>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the autonomous system (AS) number for the router. A router can only belong to one AS. An AS number is a globally unique number with an AS. This number is used to exchange exterior routing information with neighboring ASs and as an identifier of the AS itself.

If the AS number is changed on a router with an active BGP instance, the new AS number is not used until the BGP instance is restarted either by administratively disabling or enabling (**shutdown** or **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

ecmp

Syntax

ecmp *max-ecmp-routes*

no ecmp

Context

config>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables ECMP and configures the number of routes for path sharing. For example, a value of 2 means two equal-cost routes are used for cost sharing.

ECMP can only be used for routes learned with the same preference and same protocol. When more ECMP routes are available at the best preference than configured in *max-ecmp-routes*, the lowest next-hop IP address algorithm is used to select the number of routes configured in *max-ecmp-routes*.

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, route selection is as follows:

- IGP selects the next-hop based on the lowest router ID
- static-route chooses the next-hop based on lowest next-hop IP address

Default

no ecmp

Parameters

max-ecmp-routes

Specifies the maximum number of equal cost routes allowed on this routing table instance, expressed as a decimal integer. Setting ECMP *max-ecmp-routes* to 1 yields the same result as entering **no ecmp**.

Values 1 to 4

mpls-labels

Syntax

mpls-labels

Context

config>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

static-label-range

Syntax

static-label-range *static-range*

no static-label-range

Context

config>router>mpls-labels

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the range of MPLS static label values shared among static LSP, MPLS-TP LSP, and static service VC labels. When 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.

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

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 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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the range of the segment routing global block (SRGB). It is a label block that is used for assigning labels to SR prefix SIDs originated by the router. The range is carved from the system dynamic label range and is not instantiated by default.

This is a reserved label and when configured it cannot be used by other protocols such as RSVP, LDP, and BGP to assign a label dynamically.

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

Default

no sr-labels

Parameters

start *start-value*

Specifies the start label value in the SRGB.

Values 18432 to 131071

end *end-value*

Specifies the end label value in the SRGB.

Values 18432 to 131071

router-id

Syntax

router-id *ip-address*

no router-id

Context

config>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

The system uses the system interface address (which is also the loopback address)

If a system interface address is not configured, use the last 32 bits of the chassis MAC address.

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]  
[enable | disable] next-hop ip-address
```

```
[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]  
[tag tag] [enable | disable] next-hop ip-int-name | ip-address [{cpe-check cpe-ip-address [interval  
seconds] [drop-count count] [log]}] | {prefix-list prefix-list-name [all | none]}] [description description]
```

Context

```
config>router
```

Platforms

7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

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.

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

If a static-route command is issued with no **cpe-check** target but the destination prefix/netmask and next-hop match a static route that did have an associated cpe-check, the cpe-check test is 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 to uniquely identify the static route must be entered.

Parameters

ip-prefix/prefix-length

Specifies the destination address of the static route.

Values

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

ipv4-prefix-length 0 to 32

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 (hexadecimal)

d — 0 to 255 (decimal)

ipv6-prefix-length 0 to 128 (7210 SAS-D and
7210 SAS-K 2F1C2T)

0 to 64 (7210 SAS-Dxp)

ip-address

Specifies the IP address of the IP interface. The *ip-addr* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values

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

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 is not changed unless specified.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is according to the route preference defaults listed in [Table 7: Default route preferences](#).

If multiple routes are learned with an identical preference using the same protocol, the lowest-cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the route to use is determined by the next-hop with the lowest address.

Values 1 to 255

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 does not change unless specified. This value also determines which static route to install in the forwarding table.

If there are multiple routes with different preferences, the lower preference route is installed. If there are multiple static routes with the same preference but different metrics, the lower cost (metric) route is installed. If there are multiple static routes with the same preference and metric, the route with the lowest next-hop IP address is installed.

Values 0 to 65535

Default 1

black-hole

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

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

next-hop *ip-address*

Specifies the directly connected next hop IP address used to reach the destination.

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

The *ip-address* 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)

enable

Static routes can be administratively enabled or disabled. The **enable** parameter reenables 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

Static routes can be administratively enabled or disabled. The **disable** parameter disables 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

cpe-check cpe-ip-address

Specifies the IP address of the target CPE device. ICMP pings are sent to this target IP address. This parameter must be configured to enable CPE connectivity for the associated static route. To avoid possible circular references, the target IP address cannot be in the same subnet as the static route subnet.

Default no cpe-check enabled

interval seconds

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

Values 1 to 255

Default 1

drop-count 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

log

Sets the ability to log transitions between active and inactive based on the CPE connectivity check. Events should be sent to the system log, syslog, and SNMP traps.

static-route**Syntax**

```
[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 [bfd-enable | {cpe-check cpe-ip-address
```

```

[interval seconds] [drop-count count] [log] | {prefix-list prefix-list-name [all | none]}}] [ldp-sync]
[description description]
[no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [tag
tag] [enable | disable] indirect ip-address [mcast-family] [community comm-id] [{cpe-check cpe-
ip-address [interval seconds] [drop-count count] [log] | {prefix-list prefix-list-name [all | none]}}]
[description description]
[no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [tag
tag] [enable | disable] black-hole] [prefix-list prefix-list-name [all | none]] [description description]
[no] static-route {ip-prefix/prefix-length | ip-prefix netmask} black-hole [preference preference] [metric
metric] [tag tag] [enable | disable] [description description]

```

Context

config>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command creates static route entries for both the network and access routes. When configuring a static route, the **next-hop**, **indirect**, or **black-hole** parameter, indicating the type of static route, must be configured. Multiple types of static routes (**next-hop**, **indirect**, **black-hole**) can be applied to the same IP prefix. If a static route that is forwarding traffic goes down, the default route is 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 preference 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, the **cpe-check** parameters must match. If they do not, the new configuration command is rejected.

If a **static-route** command is issued with no **cpe-check** target, but the destination prefix/netmask and next hop matches a static route that did have an associated CPE check, the **cpe-check** test is 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 to uniquely identify the static route must be entered.

Default

no static-route

Parameters

ip-prefix/prefix-length

Specifies the destination address of the static route.

Values

<i>ipv4-prefix</i>	a.b.c.d (host bits must be 0)
<i>ipv4-prefix-length</i>	0 to 32

<i>ipv6-prefix</i>	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)
<i>ipv6-prefix-length</i>	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 is 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[-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 OSPF, IS-IS, or BGP expressed as a decimal integer. When modifying the preference of an existing static route, the metric is not changed unless specified.

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 7: 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

Route type	Preference	Configurable
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

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

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 does not change unless specified. This value is also used to determine which static route to install in the forwarding table.

If there are multiple routes with different preferences, the lower preference route is installed. If there are multiple static routes with the same preference but different metrics, the lower cost (metric) route is installed. If there are multiple static routes with the same preference and metric, the route with the lowest next-hop IP address is installed.

Values 0 to 65535

Default 1

black-hole

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

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

next-hop ip-int-name | ip-address

Specifies the directly connected next-hop interface name or 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** and **black-hole** keywords are mutually exclusive. If an identical command is entered (with the exception of the **black-hole** keyword), this static route is replaced with the newly entered command, and unless specified, the respective defaults for preference and metric is applied.

The *ip-int-name* is the interface name of the next hop. 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.

The *ip-address* configured for the **next-hop** parameter can be either on the network side or the access side on this node. This address must be associated with a network that is directly connected to a network configured on this node.

Values		
	<i>ip-int-name</i>	32 chars max (must start with a letter)
	ipv4-address	a.b.c.d
	ipv6-address	x:x:x:x:x:x[-interface] (eight 16-bit pieces)
		x:x:x:x:x:x.d.d.d.d[-interface]
		x - 0 to FFFF (hexadecimal)
		d - 0 to 255 (decimal)
		<i>interface</i> : mandatory for link local addresses, up to 32 characters

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 a disabled static route will be reenabled. 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 the static route will be disabled 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** keyword 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 is associated with a BFD session between the local system and the configured next hop. This keyword cannot be configured if the next hop is configured as **indirect** or **black-hole**. For more information about the protocols and platforms that support BFD, see the BFD section in the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*.

cpe-check cpe-ip-address

Specifies the IP address of the target CPE device. ICMP pings are sent to this target IP address. This parameter must be configured to enable CPE connectivity for the associated static route. To avoid possible circular references, the target IP address cannot be in the same subnet as the static route subnet. This parameter option and BFD support are mutually exclusive on a specific static route.

Default no cpe-check enabled

interval seconds

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

Values 1 to 255

Default 1

drop-count 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](#)).

log

Sets the ability to log transitions between active and in-active based on the CPE connectivity check. Events should be sent to the system log, syslog, and SNMP traps.

triggered-policy

Syntax

triggered-policy
no triggered-policy

Context

config>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command triggers 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.

2.6.2.1.3 Router DHCP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

local-dhcp-server

Syntax

local-dhcp-server *server-name* [**create**]
no local-dhcp-server *server-name*

Context

config>router>dhcp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command instantiates a local DHCP server. A local DHCP server can serve multiple interfaces but is limited to the routing context it was which it was created.

Parameters

server-name

Specifies the name of local DHCP server.

create

Specifies that the local DHCP server is created. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

force-renews

Syntax

[no] **force-renews**

Context

config>router>dhcp>server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the sending of force-renew messages.

The **no** form of this command disables the sending of force-renew messages.

Default

no force-renews

lease-hold-time

Syntax

lease-hold-time [*lease-hold-time*]

no lease-hold-time

Context

config>router>dhcp>server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the time to remember this lease. This lease-hold-time is for unsolicited release conditions such as lease timeout and normal solicited release from DHCP client.

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

Default

sec 0

Parameters

lease-hold-time

Specifies the amount of time to remember the lease.

Values		
days <i>days</i>		0 to 3650
hrs <i>hours</i>		0 to 23
min <i>minutes</i>		0 to 59
sec <i>seconds</i>		0 to 59

pool

Syntax

pool *pool-name* [**create**]

no pool *pool-name*

Context

config>router>dhcp>server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a DHCP address pool on the router.

Parameters

pool name

Specifies the name of this IP address pool. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters.

create

Specifies that the pool is created. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

max-lease-time

Syntax

max-lease-time [*max-lease-time*]

no max-lease-time

Context

config>router>dhcp>server>pool

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the maximum lease time.

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

Default

10 days

Parameters

time

Specifies the maximum lease time.

Values		
days	<i>days</i>	0 to 3650
hrs	<i>hours</i>	0 to 23
min	<i>minutes</i>	0 to 59
sec	<i>seconds</i>	0 to 59

min-lease-time

Syntax

min-lease-time [*min-lease-time*]

no min-lease-time

Context

config>router>dhcp>server>pool

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the minimum lease time.

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

Default

10 minutes

Parameters

time

Specifies the minimum lease time.

Values		
days	<i>days</i>	0 to 3650
hrs	<i>hours</i>	0 to 23
min	<i>minutes</i>	0 to 59
sec	<i>seconds</i>	0 to 59

minimum-free

Syntax

minimum-free *minimum-free* [**percent**] [**event-when-depleted**]

no minimum-free

Context

config>router>dhcp>server>pool

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the minimum number of free addresses in this pool.

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

Default

1

Parameters

minimum-free

Specifies the minimum number of free addresses.

0 to 255

percent

Specifies that the value indicates a percentage.

event-when-depleted

Enables a system-generate event when all available addresses in the pool or subnet of local DHCP server are depleted.

nak-non-matching-subnet

Syntax

[no] nak-non-matching-subnet

Context

config>router>dhcp>server>pool

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the system to return a DHCP NAK message if the following conditions are met:

- the local DHCPv4 server receives a DHCP request with option 50 (meaning the client is trying to request a previously allocated message as described in section 3.2 of RFC 2131, Dynamic Host Configuration Protocol)
- the address allocation algorithm uses a pool and the address in option 50 is not in the pool

If the conditions are not met, the system drops the DHCP packet.

Default

no nak-non-matching-subnet

offer-time

Syntax

offer-time [min *minutes*] [sec *seconds*]

no offer-time

Context

config>router>dhcp>server>pool

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the offer time.

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

Default

1 minute

Parameters***time***

Specifies the offer time.

Values		
min	<i>minutes</i>	0 to 10
sec	<i>seconds</i>	0 to 59

options**Syntax****options****Context**

config>router>dhcp>server>pool

config>router>dhcp>server>pool>subnet

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure pool options. The options defined here can be overruled by defining the same option in the local user database.

custom-option**Syntax****custom-option** *option-number* **address** [*ip-address...*(up to 4 max)]**custom-option** *option-number* **hex** *hex-string***custom-option** *option-number* **string** *ascii-string***no custom-option** *option-number***Context**

config>router>dhcp>server>pool>options

config>router>dhcp>server>pool>subnet>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures specific DHCP options. The options defined here can overrule options in the local user database.

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

Parameters

option-number

Specifies the option number that the DHCP server uses to send the identification strings to the DHCP client.

Values 1 to 254

address ip-address

Specifies the IP address of this host.

hex hex-string

Specifies the hex value of this option.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

string ascii-string

Specifies the value of this option, up to 127 characters.

dns-server

Syntax

dns-server address [*ip-address...*(up to 4 max)]

no dns-server

Context

config>router>dhcp>server>pool>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the IP address of the DNS server.

Parameters

ip-address

Specifies the IP address of the DNS server in dotted-decimal notation. Up to four addresses can be entered.

Values a.b.c.d

domain-name

Syntax

domain-name *domain-name*

no domain-name

Context

config>router>dhcp>server>pool>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the default domain for a DHCP client that the router uses to complete unqualified hostnames (without a dotted-decimal domain name).

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

Parameters

domain-name

Specifies the domain name for the client.

Values 127 character maximum

lease-rebind-time

Syntax

lease-rebind-time [*lease-rebind-time*]

no lease-rebind-time

Context

config>router>dhcp>server>pool>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the time the client transitions to a rebinding state.

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

Parameters

time

Specifies the lease rebind time.

Values		
	days <i>days</i>	0 to 3650
	hrs <i>hours</i>	0 to 23
	min <i>minutes</i>	0 to 59
	sec <i>seconds</i>	0 to 59

lease-renew-time

Syntax

lease-renew-time [*lease-renew-time*]
no lease-renew-time

Context

config>router>dhcp>server>pool>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the time the client transitions to a renew state.
 The **no** form of this command removes the time from the configuration.

Parameters

time

Specifies the lease renew time.

Values		
	days:	0 to 3650
	hours:	0 to 23
	minutes:	0 to 59
	seconds	0 to 59

lease-time

Syntax

lease-time [*lease-time*]
no lease-time

Context

```
config>router>dhcp>server>pool>options
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the amount of time that the DHCP server grants to the DHCP client permission to use a particular IP address.

The **no** form of this command removes the lease time parameters from the configuration.

Parameters

time

Specifies the lease time.

Values		
days <i>days</i>		0 to 3650
hrs <i>hours</i>		0 to 23
min <i>minutes</i>		0 to 59
sec <i>seconds</i>		0 to 59

subnet

Syntax

```
subnet {ip-address/mask | ip-address netmask} [create]
```

```
no subnet {ip-address/mask | ip-address netmask}
```

Context

```
config>router>dhcp>server>pool
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command creates a subnet of IP addresses to be served from the pool. The subnet cannot include any addresses that were assigned to subscribers without those addresses specifically excluded. When the subnet is created no IP addresses are made available until a range is defined.

Parameters

ip-address

Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values a.b.c.d (no multicast address)

mask

Specifies the subnet mask in Classless Inter-Domain Routing (CIDR) notation, expressed as a decimal integer.

Values 8 to 30

netmask

Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.

Values a.b.c.d (any mask expressed as dotted quad)

create

Specifies that the subnet is created. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

address-range

Syntax

[no] **address-range** *start-ip-address end-ip-address* [**failover** {**local** | **remote**}]

Context

config>router>dhcp>server>pool>subnet

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a range of IP addresses to be served from the pool. All IP addresses between the start and end IP addresses are included (other than specific excluded addresses).

Parameters

start-ip-address

Specifies the start address of this range to include. This address must be unique within the subnet and specified in dotted decimal notation.

Values a.b.c.d

end-ip-address

Specifies the end address of this range to include. This address must be unique within the subnet and specified in dotted decimal notation.

Values a.b.c.d

failover local

Specifies that the DHCP server failover control type is in control under normal operation.

failover remote

Specifies that the remote DHCP server failover system is in control under normal operation.

exclude-addresses

Syntax

[no] **exclude-addresses** *start-ip-address* [*end-ip-address*]

Context

config>router>dhcp>server>pool>subnet

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies a range of IP addresses that excluded from the pool of IP addresses in this subnet.

Parameters

start-ip-address

Specifies the start address of this range to exclude. This address must be unique within the subnet and specified in dotted decimal notation.

Values a.b.c.d

end-ip-address

Specifies the end address of this range to exclude. This address must be unique within the subnet and specified in dotted decimal notation.

Values a.b.c.d

maximum-declined

Syntax

maximum-declined *maximum-declined*

no maximum-declined

Context

config>router>dhcp>server>pool>subnet

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the maximum number of declined addresses allowed.

Default

64

Parameters

maximum-declined

Specifies the maximum number of declined addresses allowed.

Values 0 to 4294967295

minimum-free

Syntax

minimum-free *minimum-free* [percent] [event-when-depleted]

no minimum-free

Context

config>router>dhcp>server>pool>subnet

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the minimum number of free addresses in this subnet. If the actual number of free addresses in this subnet falls below this configured minimum, a notification is generated.

Default

1

Parameters

minimum-free

Specifies the minimum number of free addresses in this subnet.

Values 0 to 255

percent

Specifies that the value indicates a percentage.

event-when-depleted

This parameter enables a system-generate event when all available addresses in the pool or subnet of local DHCP server are depleted.

default-router

Syntax

default-router *ip-address* [*ip-address...*(up to 4 max)]

no default-router

Context

config>router>dhcp>server>pool>subnet>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the IP address of the default router for a DHCP client. Up to four IP addresses can be specified.

The **no** form of this command removes the address or addresses from the configuration.

Parameters

ip-address

Specifies the IP address of the default router. This address must be unique within the subnet and specified in dotted decimal notation.

Values a.b.c.d

subnet-mask

Syntax

subnet-mask *ip-address*

no subnet-mask

Context

config>router>dhcp>server>pool>subnet>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the subnet-mask option to the client. The mask can either be defined (for super-netting) or taken from the pool address.

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

Parameters

ip-address

Specifies the IP address of the subnet mask. This address must be unique within the subnet and specified in dotted decimal notation.

Values a.b.c.d

use-gi-address

Syntax

use-gi-address [**scope** *scope*]

Context

config>router>dhcp>local-dhcp-server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the use of gi-address matching. If the gi-address flag is enabled, a pool can be used even if a subnet is not found. If the *local-user-db-name* is not used, the gi-address flag is used and addresses are handed out by GI only. If a user must be blocked from getting an address, the server maps to a local user database and configures the user with no address.

A pool can include multiple subnets. Because the GI is shared by multiple subnets in a subscriber interface, the pool may provide IP addresses from any of the subnets included when the GI is matched to any of its subnets. This allows a pool to be created that represents a sub-int.

Default

no use-gi-address

Parameters

scope *scope*

Specifies if addresses are assigned for a specific subnet where the GI address belongs to only or for all subnets part of the pool.

Values **subnet** — Addresses are only assigned for the subnet where the GI address belongs.

pool — All subnets that are part of the pool containing the subnet to which the GI address can assign addresses.

user-db

Syntax

user-db *local-user-db-name* [**create**]

no user-db

Context

config>router>dhcp>server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a local user database for authentication.

Default

no user-db

Parameters

local-user-db-name

Specifies the name of a local user database.

create

Specifies that the local user database is created. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

2.6.2.1.4 Route next-hop policy commands

route-next-hop-policy

Syntax

route-next-hop-policy

Context

config>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure route next-hop policies.

abort

Syntax

abort

Context

config>router>route-next-hop-policy

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the editing mode for 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

template

Syntax

[no] **template-name** *name*

Context

config>router>route-next-hop-policy

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command creates a template to configure the attributes of a Loop-Free Alternate (LFA) Shortest Path First (SPF) policy. An LFA SPF policy allows the user to apply specific criteria, such as admin group and SRLG constraints, to the selection of an LFA backup next-hop for a subset of prefixes which resolve to a specific primary next-hop.

First, the user 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 for the template name 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 is still stored temporarily in the template module until the **commit** command is executed under the **route-next-hop-policy** context. Any temporary parameter changes are 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 has no effect on the prior deletion or creation of a template.

After the **commit** command is executed, IS-IS or OSPF reevaluates the templates. If there are any net changes, IS-IS or OSPF 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 a template.

Parameters

template-name

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

description

Syntax

description *description-string*

no description

Context

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the description of the next-hop template.

Parameters

description-string

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

exclude-group

Syntax

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

Context

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

**Note:**

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* c d**

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

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

**Note:**

The admin group criteria are applied before running the LFA next-hop selection algorithm.

The **no** form deletes the admin group constraint from the route next-hop policy template.

Parameters

ip-admin-group-name

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

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 | tunnel}

no nh-type

Context

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

The user can select IP backup next-hop is preferred.

When the route next-hop policy template is applied to an IP interface, all prefixes using this interface as a primary next-hop follow the next-hop type preference specified in the template.

The **no** form deletes the next-hop type constraint from the route next-hop policy template.

Default

ip

Parameters

{ip | tunnel}

Specifies the two possible values for the next-hop type.

protection-type

Syntax

protection-type {link | node}

no protection-type

Context

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 falls 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 follow the protection type preference specified in the template.

The **no** form 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 attempts 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.



Note:

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.

2.6.2.1.5 Router interface commands

interface

Syntax

[no] **interface** *interface-name*

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a system or a loopback IP routing interface. When created, attributes like IP address, 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, not a specific interface. The system interface is also referred to as the loopback address.

The **no** form of this command removes the IP interface and all the associated configurations. The interface must be administratively shut down before issuing the **no interface** command.

Parameters

interface-name

Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** commands. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

If the *interface-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 occurs and the context is not changed to that IP interface. If *interface-name* does not exist, the interface is created and the context is changed to that interface for further command processing.

Values 1 to 32 alphanumeric characters.

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 to a system IP interface. Only one IP address can be associated with an IP interface.

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.

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

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.

Parameters

ip-address

Specifies the IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values a.b.c.d (no multicast or 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-address* from the *mask* 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.

Values 1 to 32

netmask

Specifies the subnet mask in dotted decimal notation.

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

broadcast {all-ones | host-ones}

Specifies an optional 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-address* 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

Parameters

seconds

The minimum number of seconds a learned ARP entry is stored in the ARP table, expressed as a decimal integer. A value of 0 specifies that the timer is inoperative and learned ARP entries are not 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

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 1000

receive receive-interval

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

Values 10 to 1000

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 1000

Default 100

type iom-hw

Specifies that IOM-based hardware BFD sessions are used. 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

local-proxy-arp

Syntax

local-proxy-arp

no local-proxy-arp

Context

config>router>interface

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables local proxy ARP on the interface.

The **no** form of this command disables local proxy ARP on the interface.

Default

no local-proxy-arp

ldp-sync-timer

Syntax

ldp-sync-timer *seconds*

no ldp-sync-timer

Context

config>router>interface

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the IGP-LDP synchronization timer. This timer enables synchronization of IGP and LDP, and synchronization of static routes and LDP. This command is not supported on RIP interfaces.

When a link is restored after a failure, IGP sets the link cost to infinity and advertises it; if it's a static route, the route activation is delayed until this timer expires. The supported IGP's are OSPF and IS-IS. 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).

If an interface belongs to both IS-IS and OSPF, a physical failure causes both IGP's to advertise infinite metric and to follow the IGP-LDP synchronization procedures. If only one IGP bounces on this interface or on the system, only the affected IGP advertises the infinite metric and follows the IGP-LDP synchronization procedures.

After IGP advertises the link cost, the LDP hello adjacency is brought up with the neighbor. IGP starts the LDP synchronization timer when the LDP session to the neighbor becomes operationally 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 announces a new best next-hop and LDP uses it if the label binding for the neighbor's 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 is only advertised after the timer expires. Also, if the currently advertised cost is different, the new cost value is advertised after the user executes any of the following commands:

- **tools>perform>router>ospf>ldp-sync-exit**
- **tools>perform>router>isis>ldp-sync-exit**
- **config>router>interface>no ldp-sync-timer**
- **config>router>ospf>disable-ldp-sync**
- **config>router>isis>disable-ldp-sync**

See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C OAM and Diagnostics Guide* for the tools commands and to the *7210 SAS-D, Dxp, K 2F1C2T Routing Protocols Guide* and the *7210 SAS-K 2F6C4T, K 3SFP+ 8C Routing Protocols Guide* for the OSPF and IS-IS commands.

If the user changes the value of the LDP synchronization timer parameter, the new value takes effect at the next synchronization event. That is, if the timer is still running, it continues using the previous value.

If parallel links exist to the same neighbor, 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, the 7210 SAS only considers this interface for forwarding after IGP re-advertises its actual cost value.

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

If the **ldp-sync-timer** value is configured on the interface but LDP is not running on the interface, the configuration causes the IGP route cost to increase to the maximum value.

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

Default

no ldp-sync-timer

Parameters

seconds

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

Values 1 to 1800

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

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 a 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 specific IP address to send it to the corresponding ARP address.

Default

no proxy-arp-policy

Parameters

policy-name

Specifies the export route policy name. Allowed values are any string up to 32 characters 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 names must already be defined.

remote-proxy-arp

Syntax

[no] remote-proxy-arp

Context

config>router>interface

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables remote proxy ARP on the interface.

Default

no remote-proxy-arp

static-arp

Syntax

static-arp *ip-address* *ieee-address*

no static-arp *ip-address*

static-arp *ieee-address* **unnumbered**

no static-arp **unnumbered**

Context

config>router>interface

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

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 7210 SAS configuration can state that if it has a packet that has a specific IP address to send the packet to the corresponding ARP address.

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

Default

no static-arp

Parameters

ip-address

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

ieee-address

Specifies the 48-bit MAC address for the static ARP in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

unnumbered

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

unnumbered

Syntax

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

no unnumbered

Context

config>router>interface

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-address* parameter configured.

An error message is generated when 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 the **no unnumbered** command is issued to delete the IP address from the interface.

Default

no unnumbered

Parameters

ip-int-name | *ip-address*

Specifies the IP interface name or IP address with which to associate the unnumbered IP interface, in dotted decimal notation. The configured IP address must exist on this node. Nokia recommends to use 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 *ip-int-name* or *ip-address* is not configured.

urpf-check

Syntax

urpf-check

Context

config>router

Platforms

7210 SAS-K 3SFP+ 8Cs

Description

This command enables the Unicast RPF check feature on this router.

ignore-default

Syntax

[no] ignore-default

Context

```
config>router>urpf-check  
config>router>urpf-check>ipv6>ignore-default
```

Platforms

7210 SAS-K 3SFP+ 8C

Description

This command configures the Unicast RPF check feature (if enabled) to ignore default routes for purposes of determining the validity of incoming packets.

The **no** form of this command considers the default route to be eligible when performing a Unicast RPF check.

Default

```
no ignore-default
```

2.6.2.1.6 Router interface filter commands

egress

Syntax

```
egress
```

Context

```
config>router>interface
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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*

no filter

Context

config>router>if>ingress

config>router>if>egress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command associates an IP filter policy with an IP interface.

Filter policies control packet forwarding and dropping based on IP match criteria.

The *ip-filter-id* must have been preconfigured before this **filter** command is executed. If the filter ID does not exist, an error occurs.

Only one filter ID can be specified.



Note:

For more information about service and IP interface support for different ACL match criteria for each platform, see [Filter policy entities](#).

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

Parameters

ip 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

2.6.2.1.7 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-D, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

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.

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval.

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

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. This value is expressed as a decimal integer.

Values 1 to 60

Default 10

tll-expired

Syntax

```
tll-expired [number seconds]
```

```
no tll-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.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

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

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. This value is 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.

By default, generation of ICMP destination unreachables messages is enabled at a maximum rate of 100 per 10 second time interval.

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

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. The value is expressed as a decimal integer.

Values 1 to 60

Default 10

2.6.2.1.8 Interface attribute commands

if-attribute

Syntax

if-attribute

Context

config>router

config>router>interface

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

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

Next, the user 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 computes a path that satisfies the inclusion and exclusion constraints of the admin group.

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 the admin group configuration. The system rejects 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.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 that 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.

First, the user configures, locally on each router, the name and identifier of each SRLG group. A maximum of 1024 SRLGs can be configured per system.

Next, the user 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.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 the admin groups bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the traffic-engineering option is enabled in IS-IS or OSPF.

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.

srlg-group

Syntax

srlg-group *group-name* [*group-name...* (up to 5 max)]

no srlg-group *group-name*

Context

config>router>interface>if-attribute

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

2.6.2.1.9 Router interface IPv6 commands

ipv6

Syntax

[no] **ipv6**

Context

config>router>interface

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

This command configures IPv6 for a router interface.

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

Default

not enabled

address

Syntax

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

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

Context

config>router>if>ipv6

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

This command assigns an IPv6 address to the interface.

Parameters

ipv6-address/prefix-length

Specifies the IPv6 address on the interface.

Values

ipv6-address/prefix:	ipv6-address	x:x:x:x:x:x:x (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	1 to 128 (7210 SAS-D, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C)
		1 to 64 (7210 SAS-Dxp)

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 addresses 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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-D, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

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 (when IPv6 is enabled on the interface)

Parameters

number

Specifies that the number of redirects 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 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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

This command configures the link local address.

local-proxy-nd

Syntax

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

Context

config>router>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

Parameters

policy-name

Specifies the neighbor discovery policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. The specified policy names must already be defined.

neighbor

Syntax

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

no neighbor [*ipv6-address*]

Context

config>router>if>ipv6

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

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

2.6.2.1.10 IPv6 router advertisement commands

router-advertisement

Syntax

```
router-advertisement  
no router-advertisement
```

Context

```
config>router
```

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command enables router advertisement on IPV6 interfaces. By default, it is disabled for all IPV6-enabled interfaces.

The **no** form of this command disables router advertisement on all IPV6 interfaces.

Default

```
no router-advertisement
```

interface

Syntax

```
interface ip-int-name  
no interface ip-int-name
```

Context

```
config>router>router-advertisement
```

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command configures router advertisement properties on a specified IPv6 interface. The interface name must already exist in the **config>router>interface>ipv6** context.

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

Parameters

template-name

Specifies the name of an existing IPv6 interface. An interface name cannot be in the form of an IPv6 address. If the string contains special characters (such as #, \$, spaces), the entire string must be enclosed within double quotes. The string must start with a letter and the max is 32 characters.

current-hop-limit

Syntax

current-hop-limit *number*

no current-hop-limit

Context

config>router>router-advertisement>if

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command enables the router to advertise the hop-limit in ICMPv6 Neighbor Discovery (ND) router advertisement messages.

The **no** form of this command disables the advertising of the hop-limit in ICMPv6 ND router advertisement messages by the router.

Default

no current-hop-limit

Parameters

number

Specifies the number of hop limits.

Values	0 to 255 (a value of 0 means that there are an unspecified number of hops)
---------------	--

Default	64
----------------	----

managed-configuration

Syntax

managed-configuration

no managed-configuration

Context

config>router>router-advertisement>if

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command enables the "managed-config-flag" to be advertised in ICMPv6 ND router advertisement messages.

The **no** form of this command disables the advertising of the "managed-config-flag" in ICMPv6 router advertisement messages.

Default

no managed-configuration

max-advertisement-interval

Syntax

max-advertisement-interval *seconds*

no max-advertisement-interval

Context

config>router>router-advertisement>if

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command configures the maximum interval between sending ICMPv6 ND router advertisement messages.

The **no** form of this command disables the setting of a maximum interval between sending ICMPv6 ND router advertisement messages.

Default

no max-advertisement-interval

Parameters

seconds

Specifies the maximum interval, in seconds.

Values 4 to 1800

Default 600

min-advertisement-interval

Syntax

min-advertisement-interval *seconds*
no min-advertisement-interval

Context

config>router>router-advertisement>if

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

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

The **no** form of this command disables the setting of a minimum interval between sending ICMPv6 ND router advertisement messages.

Default

no min-advertisement-interval

Parameters

seconds

Specifies the minimum interval, in seconds.

Values 3 to 1350

Default 200

mtu

Syntax

mtu *mtu-bytes*
no mtu

Context

config>router>router-advertisement>interface

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command configures the MTU for sending packets to the router.

The **no** form of this command disables the sending of MTU in the router advertisement messages

Default

no mtu

Parameters

mtu-bytes

Specifies the MTU for the nodes to use when sending packets.

Values 1280 to 9212

other-stateful-configuration

Syntax

other-stateful-configuration

no other-stateful-configuration

Context

config>router>router-advertisement>interface

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command enables the "other-config-flag" to be advertised in ICMPv6 ND router advertisement messages.

The **no** form of this command disables the advertising of the "other-config-flag" in ICMPv6 router advertisement messages.

Default

no other-stateful-configuration

prefix

Syntax

prefix *ipv6-prefix/prefix-length*

no prefix

Context

```
config>router>router-advertisement>interface
```

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command configures the IPv6 prefix to include in router advertisement messages. To support multiple IPv6 prefixes, use multiple prefix statements.

The **no** form of this command disables the inclusion of an IPv6 prefix in router advertisement messages.

Default

no prefix

Parameters***ipv6-prefix/prefix-length***

The Pv6 prefix.

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 (hexadecimal)
	d - 0 to 255 (decimal)
ipv6-prefix-length	0 to 128 (7210 SAS-D, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C)
	0 to 64 (7210 SAS-Dxp)

autonomous**Syntax**

autonomous

no autonomous

Context

```
config>router>router-advertisement>interface>prefix
```

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether the prefix can be used for a stateless address autoconfiguration.
The **no** form of this command disables the prefix to be used for a stateless address autoconfiguration.

Default

autonomous

on-link

Syntax

on-link
no on-link

Context

config>router>router-advertisement>interface>prefix

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether the prefix can be used for onlink determination.
The **no** form of this command disables the prefix to be used for onlink determination.

Default

on-link

preferred-lifetime

Syntax

preferred-lifetime [*seconds* | *infinite*]
no preferred-lifetime

Context

config>router>router-advertisement>interface>prefix

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command configures the remaining time, in seconds, that this prefix continues to be preferred (time until deprecation). The address generated from a deprecated prefix should not be used as a source

address in new communications. However, packets received on such an interface are processed as expected.

The **no** form of this command does disables the configuration of the time until deprecation for the prefix.

Default

no preferred-lifetime

Parameters

seconds

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

Values 1 to 4294967294

infinite

Specifies that the prefix will always be preferred. A value of 4294967295 represents infinity.

valid-lifetime

Syntax

valid-lifetime [*seconds* | *infinite*]

no valid-lifetime

Context

config>router>router-advertisement>interface>prefix

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command specifies the length of time, in seconds, that the prefix is valid for the purpose of onlink determination. The address generated from an invalidated prefix should not appear as the destination or source address of a packet.

The **no** form of this command disables configuration of the time that the prefix is valid for the purpose of onlink determination.

Default

no valid-lifetime

Parameters

seconds

Specifies the remaining length of time, in seconds, that this prefix will be valid.

Values 1 to 4294967294

infinite

Specifies that the prefix will always be valid. A value of 4294967295 represents infinity.

reachable-time

Syntax

reachable-time *milli-seconds*

no reachable-time

Context

config>router>router-advertisement>interface

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command configures how long the router should be considered reachable by other nodes on the link after receiving a reachability confirmation.

The **no** form of this command disables the configuration of how long the router should be considered reachable.

Default

no reachable-time

Parameters

milli-seconds

Specifies the amount of time, in milliseconds, that the router will be considered reachable.

Values 0 to 3600000

retransmit-time

Syntax

retransmit-time *milli-seconds*

no retransmit-time

Context

config>router>router-advertisement>interface

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command configures the retransmission frequency of neighbor solicitation messages.

The **no** form of this command disables the configuration of the retransmission frequency of neighbor solicitation messages.

Default

no retransmit-time

Parameters

milli-seconds

Specifies the length of time, in milliseconds, that a host should wait before retransmitting neighbor solicitation messages.

Values 0 to 1800000

router-lifetime

Syntax

router-lifetime *seconds*

no router-lifetime

Context

config>router>router-advertisement>interface

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command configures the router lifetime.

The **no** form of this command disables the configuration of the router lifetime.

Default

no router-lifetime

Parameters

seconds

Specifies the router lifetime, in seconds.

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

2.6.2.2 Show commands

aggregate

Syntax

aggregate [family] [active]

Context

show>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays aggregate routes.

Parameters

active

Specifies that 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 only ARP entries associated with the specified IP address and mask.

ip-int-name

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

mac *ieee-mac-addr*

Displays only ARP entries associated with the specified MAC address.

summary

Displays an abbreviated list of ARP entries.

[local | dynamic | static]

Displays only ARP information associated with the keyword.

Output

The following output is an example of router ARP table information, and [Table 8: Output fields: 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
=====
```

Table 8: Output fields: ARP

Label	Description
IP Address	Displays the IP address of the ARP entry
MAC Address	Displays the MAC address of the ARP entry
Expiry	Displays 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
Int	Specifies that the ARP entry is an internal ARP entry
[I]	Specifies that the ARP entry is in use
Interface	Displays the IP interface name associated with the ARP entry
No. of ARP Entries	Displays the number of ARP entries displayed in the list

neighbor

Syntax

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

Context

show>router

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

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*

Specify 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 neighbor information, and [Table 9: Output fields: 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
-----
2001:db8::5      REACHABLE  A_to_B2_17    Static    No
  00:00:1b:00:00:01
```

```

2001:db8::2
e4:81:84:24:1d:6c          STALE          A_to_B2_23
                          01h12m35s          Dynamic          Yes
-----
No. of Neighbor Entries: 2
=====
*A:Dut-A>config>router# show router neighbor dynamic
=====
Neighbor Table (Router: Base)
=====
IPv6 Address          Interface
MAC Address          Expiry          Type          RTR
-----
2001:db8::2          A_to_B2_23
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          Expiry          Type          RTR
-----
2001:db8::5          A_to_B2_17
00:00:1b:00:00:01          REACHABLE          -          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          Expiry          Type          RTR

```

Table 9: Output fields: 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

Label	Description
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

dhcp

Syntax

dhcp

Context

show>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays DHCP information for the specified service.

local-dhcp-server

Syntax

local-dhcp-server *server-name*

Context

show>router>dhcp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays local DHCP or DHCP 6server information.

Parameters

server-name

Specifies information about the local DHCP server.

declined-addresses

Syntax

declined-addresses *ip-address*[/*mask*] [**detail**]

declined-addresses pool *pool-name*

Context

show>router>dhcp>server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays information about declined addresses.

Parameters

pool *pool-name*

Specifies a DHCP pool name on the router.

ip-address

Specifies the IP address of the DNS server. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

mask

Specifies the subnet mask in Classless Inter-Domain Routing (CIDR) notation, expressed as a decimal integer.

Values 0 to 32

detail

Displays detailed information.

Output

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

Sample output

```
*A:ALA-48>show>router>dhcp>server# declined-addresses pool test
=====
Declined addresses for server test Base
=====
Pool                               Subnet      IP Address
PPPoE User Name/                   Time        MAC Address  Type
Option 82 Circuit ID
-----
No Matching Entries
```

```
=====
*A:ALA-48>show>router>dhcp>server#
```

Table 10: Output fields: declined addresses

Label	Description
Pool	Displays the name of the DHCP address pool
PPoe User Name/ Option 82 Circuit ID	Displays the PPOE username or Option 82 circuit ID
Subnet	Displays the subnet of the DHCP address pool
Time	Displays the time that the address was declined
IP Address	Displays the declined IP address
MAC Address	Displays the declined MAC address
Type	Displays the type of pool

free-addresses

Syntax

free-addresses *ip-address*[/mask]

free-addresses summary [subnet *ip-address*[/mask]

free-addresses pool *pool-name*

Context

show>router>dhcp>local-dhcp-server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays the free addresses in a subnet.

Parameters

pool *pool-name*

Specifies a DHCP pool name on the router.

subnet *subnet*

Specifies a subnet of IP addresses that are served from the pool.

summary

Displays summary output of the free addresses.

Output

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

Sample output

```
*A:ALA-48>show>router>dhcp>local-dhcp-server# free-addresses pool test subnet
1.0.0.0/24
=====
Free addresses in subnet 1.0.0.0/24
=====
IP Address
-----
No. of free addresses: 0
=====
*A:ALA-48>show>router>dhcp>local-dhcp-server#
```

Table 11: Output fields: free addresses

Label	Description
IP Address	The free IP address
No. of free addresses	Displays the number of free IP addresses

leases

Syntax

leases [**detail**]

leases *ip-address*[*/mask*] **address-from-user-db** [**detail**]

leases *ip-address*[*/mask*] **dhcp-host** *dhcp-host-name* [**detail**]

leases *ip-address*[*/mask*] [**detail**]

Context

show>router>dhcp>local-dhcp-server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays the DHCP leases.

Parameters

ip-address

Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

mask

Specifies the subnet mask, in dotted-decimal notation.

Values 0 to 32

address-from-user-db [detail]

Displays only leases that have ip-addresses from the local-user-db.

dhcp-host dhcp-host-name [detail]

Displays all leases that match a certain DHCP host from the local-user-db.

detail

Displays detailed information of all leases that fall into the indicated subnet.

The command with no parameters will show all leases from the local-user-db.

Output

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

Sample output

```
*A:ALA-48>show>router>dhcp>local-dhcp-server# leases ip-address 10.0.0.4
=====
Leases for DHCP server test router Base
=====
IP Address      Lease State      Mac Address      Remaining Clnt
  PPPoE user name/Opt82 Circuit Id      LifeTime  Type
-----
No leases found
*A:ALA-48>show>router>dhcp>local-dhcp-server#
```

Table 12: Output fields: lease

Label	Description
IP Address	The leased IP address
PPoE user name/ Opt82 Circuit Id	The PPoE username or Option 82 circuit ID
Lease State	The state of the lease. The state can be: <ul style="list-style-type: none"> advertised remove-pending held stable
Mac Address	The MAC address
Remaining LifeTime	The remaining time left in the lease

Label	Description
Clnt Type	The type of client

server-stats

Syntax

```
server-stats
```

Context

```
show>router>dhcp>server
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays DHCP or DHCP6 server statistics.

Output

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

Sample output

```
*A:SUB-Dut-A# show router dhcp local-dhcp-server dhcpS1 server-stats
=====
Statistics for DHCP Server dhcpS1 router Base
=====
Rx Discover Packets      : 0
Rx Request Packets      : 0
Rx Release Packets      : 0
Rx Decline Packets      : 0
Rx Inform Packets       : 0

Tx Offer Packets        : 0
Tx Ack Packets          : 0
Tx Nak Packets          : 0
Tx Forcerenew Packets   : 0

Client Ignored Offers    : 0
Leases Timed Out        : 0

Dropped Bad Packet      : 0
Dropped Invalid Type    : 0
Dropped No User Database : 0
Dropped Unknown Host    : 0
Dropped User Not Allowed : 0
Dropped Lease Not Ready : 0
Dropped Lease Not Found : 0
Dropped Not Serving Pool : 0
Dropped Invalid User    : 0
Dropped Overload        : 0
Dropped Persistence Overload : 0
```

```
Dropped Generic Error      : 0
Dropped Destined To Other  : 0
Dropped Address Unavailable : 0
Dropped Max Leases Reached : 0
Dropped Server Shutdown    : 0
Dropped No Subnet For Fixed IP: 0
```

```
=====
*A: SUB-Dut - A#
```

Table 13: Output fields: server statistics

Label	Description
Rx Discover Packets	The number of DHCPDISCOVER (option 53 with value 1) packets received by the DHCP server
Rx Request Packets	The number of DHCPREQUEST (option 53 with value 3) packets received by the DHCP server
Rx Release Packets	The number of DHCPRELEASE (option 53 with value 7) packets received by the DHCP server
Rx Decline Packets	The number of DHCPDECLINE (option 53 with value 4) packets received by the DHCP server
Rx Inform Packets	The number of DHCPINFORM (option 53 with value 8) packets received by the DHCP server
Tx Offer Packets	The number of DHCPOFFER (option 53 with value 2) packets sent by the DHCP server
Tx Ack Packets	The number of DHCPACK (option 53 with value 5) packets sent by the DHCP server
Tx Nak Packets	The number of DHCPNAK (option 53 with value 6) packets sent by the DHCP server
Tx Forcerenew Packets	The number of DHCPFORCERENEW (option 53 with value 9) packets sent by the DHCP server
Client Ignored Offers	The number of DHCPOFFER (option 52 with value 2) packets sent by the DHCP server that were ignored by the clients
Leases Timed Out	The number of DHCP leases that timed out without renewal
Dropped Bad Packet	The number of DHCP packets received that were corrupt
Dropped Invalid Type	The number of DHCP packets received that had an invalid message type (option 53)
Dropped No User Database	The number of DHCP packets dropped because the user-db value of the server was not equal to the default value and a local user database with that name could not be found

Label	Description
Dropped Unknown Host	The number of DHCP packets dropped from hosts that were not found in the user database when use-gi-address was disabled
Dropped User Not Allowed	The number of DHCP packets dropped from hosts, which have no specified address or pool, that were found in the user database while use-gi-address was disabled
Dropped Lease Not Ready	The number of DHCP packets dropped by the server before the lease database was ready
Dropped Lease Not Found	The number of DHCP packets dropped by the server because no valid lease was found
Dropped Not Serving Pool	The number of DHCP packets dropped by the server because there were no free addresses in the pool
Dropped Invalid User	The number of DHCP packets dropped by the server because the MAC address of the sender or the Option 82 did not match the host lease state
Dropped Overload	The number of DHCP packets dropped by the server because they were received in excess of what the server can process
Dropped Persistence Overload	The number of DHCP packets dropped by the server because they were received in excess of what the DHCP persistence system can process. If this occurs, only releases and declines are processed.
Dropped Generic Error	The number of DHCP packets dropped by the server because of a generic error
Dropped Destined to Other	The number of DHCP requests dropped by the server because the broadcast request was not addressed to this server
Dropped Address Unavailable	The number of DHCP requests dropped by the server because the requested address is not available
Dropped Max Leases Reached	The number of DHCP packets dropped by the server because the maximum number of leases was reached
Dropped Server Shutdown	The number of DHCP packets dropped by the server during server shutdown
Dropped No Subnet For Fixed IP	The number of DHCP packets dropped by the server for user-db hosts with a fixed address because the subnet to which the address belongs is not configured
Dropped Duplicate From Diff GI	The number of DHCP requests dropped by the server because they were received from a different Gateway IP address within an interval of 10 s after the previous DHCP request

subnet-ext-stats

Syntax

subnet-ext-stats *ip-address[/mask]*

subnet-ext-stats pool *pool-name*

Context

show>router>dhcp>server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays extended statistics per DHCPv4 subnet in local DHCPv4 server.

The following statistics are included in the output:

- the number of stable leases in the subnet
- the number of provisioned address in the subnet
- the number of used address in the subnet
- the number of free address in the subnet
- the percentage of used address
- the percentage of free address

For each statistic (except for provisioned addresses), there is current value and peak value, peak value is the highest value since subnet creation or last reset via the **clear router *rt-id* dhcp local-dhcp-server *svr-name* subnet-ext-stats** command.

When parameter pool is used, the statistics of each subnet in the pool will be displayed.

Parameters

ip-address[/mask]

Specifies the subnet.

pool-name

Specifies the name of local DHCPv4 server pool.

Output

The following output is an example of subnet extended statistics information, and [Table 14: Output fields: subnet extended statistics](#) describes the output fields.

Sample output

```
show router 500 dhcp local-dhcp-server "d4" subnet-ext-stats 10.10.10.0/24
=====
Extended statistics for subnet 10.10.10.0/24
=====
                Current          Peak          TimeStamp
```

```

-----
Local:
  Stable Leases           1           1           01/07/2013 19:38:36
  Provisioned Addresses  101
  Used Addresses         1           1           01/07/2013 19:38:36
  Free Addresses        100        100        01/07/2013 19:38:36
  Used Pct               1           1           01/07/2013 19:38:36
  Free Pct               99         99         01/07/2013 19:38:36
  Last Reset Time
-----
Number of entries       1
=====

```

Table 14: Output fields: subnet extended statistics

Label	Description
Current	The current value of the statistic
Peak	The highest value reached by the statistic since subnet creation or the last subnet statistics clearing operation
TimeStamp	The date and time of the current statistics capture
Offered Leases	The number of leases offered from the subnet
Stable Leases	The number of stable leases in the subnet
Provisioned Addresses	The number of provisioned addresses in the subnet
Used Addresses	The number of used addresses in the subnet
Free Addresses	The number of free addresses in the subnet
Used Pct	The percentage of used addresses in the subnet
Free Pct	The percentage of free addresses in the subnet
Last Reset Time	The date and time of the last subnet statistics clearing operation
Number of entries	The total number of subnet entries

subnet-stats

Syntax

subnet-stats *ip-address[/mask]*

subnet-stats pool *pool-name*

Context

show>router>dhcp>server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays subnet statistics.

Output

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

Sample output

```
*A:SUB-Dut-A# show router dhcp local-dhcp-server dhcpS2 subnet-stats pool P00L2
=====
Statistics for pool P00L2
=====
Subnet          Free      Offered   Stable
                FRPending RemPending Declined
-----
10.0.0.0/8      16384    0         0
                0        0         0
-----
No. of entries: 1
=====
*A:SUB-Dut-A#
```

Table 15: Output fields: subnet statistics

Label	Description
Subnet	The subnet of the pool
Free	The number of free leases in the subnet
FRPending	The number of leases in the subnet that are pending a force renew
Offered	The number of offered leases in the subnet
RemPending	The number of leases in the subnet that are pending removal
Stable	The number of stable leases in the subnet
Declined	The number of declined leases in the subnet

summary

Syntax

summary

Context

```
show>router>dhcp>server
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays DHCP server summary information.

Output

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

Sample output

```
*A:SUB-Dut-A# show router dhcp local-dhcp-server dhcpS2 summary
=====
DHCP server dhcpS2  router Base
=====
dhcpS2-P00L2
Admin State           : inService
Persistency State    : ok
User Data Base       : N/A
Use gateway IP address : disabled
Send force-renewals  : disabled
-----
Pool name : P00L2
-----
Subnet                Free          Stable      Declined    Offered     Remove-pending
-----
10.0.0.0/8           16384         0           0           0           0
-----
Totals for pool      16384         0           0           0           0
-----
Totals for server   16384         0           0           0           0
-----
Associations                Admin
-----
No associations found
=====
*A:SUB-Dut-A#

*A:vsim-2# show router 500 dhcp local-dhcp-server "d4" summary
=====
DHCP server d4  router 500
=====
Admin State           : inService
Operational State     : inService
Persistency State     : shutdown
User Data Base        : N/A
Use gateway IP address : enabled (scope subnet)
Use pool from client  : disabled
Send force-renewals  : disabled
Creation Origin       : manual
Lease Hold Time       : 0h0m0s
Lease Hold Time For   : N/A
User-ident            : mac-circuit-id
Failover Admin State  : outOfService
```

```

Failover Oper State      : shutdown
Failover Persist Key    : N/A
Administrative MCLT     : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time       : 0h2m0s
Partner down delay      : 23h59m59s
  Ignore MCLT           : disabled
-----
Pool name : v4-1
-----
Failover Admin State    : inService
Failover Oper State    : normal
Failover Persist Key    : N/A
Administrative MCLT     : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time       : 0h2m0s
Partner down delay      : 23h59m59s
  Ignore MCLT           : disabled
-----
Subnet                  Free      %      Stable  Declined Offered  Rem-pend Drain
-----
10.20.20.0/24          (L) 10      90%    1        0        0        0        N
                       (R) N/A      0        0        N/A      N/A      N/A      N
Totals for pool        10      90%    1        0        0        0
-----
Totals for server     10      90%    1        0        0        0
-----
Interface associations
Interface              Admin
-----
l1                    Up
-----
Local Address Assignment associations
Group interface        Admin
-----
=====
*A:vsim-2#

```

Table 16: Output fields: DHCP server summary

Label	Description
Admin State	The administrative state of the DHCP server
Persistency State	The persistence state of the DHCP server
User Data Base	Indicates whether the DHCP server uses a user database
Use gateway IP address	Indicates whether the DHCP server uses GIADDR
Send force-renewals	Indicates whether the DHCP server sends FORCERENEW messages
Operational State	The operational state of the DHCPv6 server
Persistency State	The persistence state of the DHCPv6 server

Label	Description
Use Link Address	Indicates whether use-link-address is enabled, and, if enabled, the scope
Use pool from client	Indicates whether use-pool-from-client is enabled
Creation Origin	The creation method of the DHCPv6 server
Lease Hold Time	The lease retention time configured using the lease-hold-time command
Lease Hold Time For	The lease being held by the DHCPv6 server
User-ident	The user identification method configured using the user-ident command
Interface-id-mapping	Indicates whether interface ID mapping is enabled
Ignore-rapid-commit	Indicates whether the DHCPv6 server is configured to ignore rapid committing
Allow-lease-query	Indicates whether the DHCPv6 server allows lease query messages
Pool	
Subnet	The subnet of the pool
Free	The number of free IP addresses in the subnet
Stable	The number of stable IP addresses in the subnet
Declined	The number of declined IP addresses in the subnet
Offered	The number of offered IP addresses in the subnet
Remove-pending	The number of IP addresses pending removal in the subnet
Associations	
Associations	The name of the associated interface
Admin	The administrative state of the interface

servers

Syntax

servers

servers all

Context

```
show>router>dhcp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command lists the local DHCP servers.

Output

The following output is an example of DHCP server information, and [Table 17: Output fields: DHCP servers](#) describes the output fields.

Sample output

```
*A:ALA-49>show>router>dhcp# servers
=====
Overview of DHCP Servers
=====
Active Leases:      0
Maximum Leases:    159744

Router              Server                               Admin State
-----
Router: Base        base_router_dhcp_server                outOfService
Service: 3          s1                                      inService
=====
*A:ALA-49>show>router>dhcp#
```

Table 17: Output fields: DHCP servers

Label	Description
Active Leases	The number of active leases
Maximum Leases	The maximum number of leases available
Router	The name of the router
Server	The name of the DHCP or DHCPv6 server
Admin State	The administrative state of the DHCP or DHCPv6 server

statistics

Syntax

```
statistics [sap sap-id] | [sdp [sdp-id[:vc-id]] | interface ip-int-name]]
```

Context

```
show>router>dhcp
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays statistics for DHCP relay and DHCP snooping.

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

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

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition.

sdp-id

Specifies the SDP ID to be shown.

Values 1 to 17407

vc-id

Specifies the virtual circuit ID on the ID to be shown.

Values 1 to 4294967295

ip-int-name | *ip-address*

Displays statistics for the specified IP interface.

Output

The following output is an example of DHCP statistics information, and [Table 18: Output fields: DHCP statistics](#) describes the output fields.

Sample output

```
A:ALA-A# show router 1000 dhcp statistics
=====
DHCP Global Statistics (Service: 1000)
=====
Rx Packets                : 16000
Tx Packets                : 15041
Rx Malformed Packets     : 0
Rx Untrusted Packets     : 0
Client Packets Discarded  : 423
Client Packets Relayed   : 0
Client Packets Snooped    : 0
Client Packets Proxied (RADIUS) : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded  : 0
Server Packets Relayed   : 0
Server Packets Snooped    : 0
DHCP RELEASEs Spoofed    : 0
DHCP FORCERENEWs Spoofed : 0
=====
A:ALA-A#
```

Table 18: Output fields: DHCP statistics

Label	Description
Received Packets	The number of packets received from the DHCP clients
Transmitted Packets	The number of packets transmitted to the DHCP clients
Received Malformed Packets	The number of malformed packets received from the DHCP clients
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded
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 the status of the DHCP relay and DHCP snooping functions on each interface.

Output

The following output is an example of DHCP summary information, and [Table 19: Output fields: DHCP summary](#) describes the output fields.

Sample output

```
A:ALA-48>show>router>dhcp# summary
=====
Interface Name           Arp      Used/    Info   Admin
                        Populate Provided Option  State
-----
ccaiesif                 No        0/0     Keep   Down
ccanet6                  No        0/0     Keep   Down
iesBundle                 No        0/0     Keep   Up
spokeSDP-test            No        0/0     Keep   Down
test                     No        0/0     Keep   Up
test1                    No        0/0     Keep   Up
test2                    No        0/0     Keep   Up
testA                    No        0/0     Keep   Up
testB                    No        0/0     Keep   Up
testIES                  No        0/0     Keep   Up
to-web                   No        0/0     Keep   Up
-----
Interfaces: 11
=====
A:ALA-48>show>router>dhcp#

*A:vsim-2# show router 500 dhcp summary
=====
DHCP Summary (Service: 500)
=====
Interface Name           Arp      Leases Per Interface/ Info   Admin
                        SapId/Sdp Populate Per Sap Limit  Option State
-----
g1                        No        1/1                       Keep   Up
  sap:1/1/7              1/1
l1                        No        0/0                       Keep   Down
-----
Interfaces: 2
=====
*A:vsim-2#
```

Table 19: Output fields: DHCP summary

Label	Description
Interface Name	Name of the router interface
ARP Populate	Indicates whether ARP populate is enabled
Used/Provided	Indicates the number of used and provided DHCP leases
Info Option	Indicates whether Option 82 processing is enabled on the interface
Admin State	Indicates the administrative state

statistics

Syntax

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

Context

show>router>dhcp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays DHCP statistics information.

Parameters

ip-int-name* | *ip-address

Displays statistics for the specified IP interface.

Output

The following output is an example of DHCP statistics information, and [Table 20: Output fields: 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 20: Output fields: 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
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

ecmp

Syntax

ecmp

Context

show>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays the ECMP settings for the router.



Note:

Weighted ECMP appears in the show output but is not supported on 7210 SAS platforms.

Output

The following output is an example of ECMP settings information, and [Table 21: 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 21: Output fields: router ECMP

Label	Description
Instance	Displays the router instance number
Router Name	Displays 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
Weight ECMP	False — Weighted ECMP is disabled

fib

Syntax

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

fib *slot-number* [*family*] **summary**

fib *slot-number* [*family*] [**nh-table-usage**]

Context

show>router

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

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

Parameters

slot-number

Displays only the routes matching the specified chassis slot number.

Values 1

family

Displays the router IP interface table for the specified family.

Values ipv4 — Displays only peers that have the IPv4 family enabled.
ipv6 — Displays only peers that have the IPv6 family enabled.

ip-prefix/prefix-length

Displays FIB entries matching only the specified *ip-prefix* and *prefix-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 (7210 SAS-D, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C) 0 to 64 (7210 SAS-Dxp)

longer

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

secondary

Displays secondary FIB information.

summary

Displays summary FIB information for the specified slot number.

nh-table-usage

Displays next-hop table usage.

Output

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

Sample output



Note:

The following output applies to the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C only.

```
*A:Dut-A# show router fib 1 10.77.77.77/32
=====
FIB Display
=====
Prefix [Flags]                                Protocol
NextHop
-----
10.77.77.77/32                                BGP
  2.2.2.2 (Transport:LDP)
  3.3.3.3 (Transport:LDP)
-----
Total Entries : 1
-----
=====
*A:Dut-A#
```

Table 22: Output fields: FIB BGP PIC

Label	Description
Prefix[Flags]	The route destination address and mask
Protocol	The active protocol (LOCAL, STATIC, OSPF, IS-IS, AGGREGATE, BGP, RIP, or BGP-VPN)
Next Hop	The next-hop or indirect next-hop IP address for the route destination
Total Entries	The total number of next-hop entries

icmp6

Syntax

icmp6

Context

show>router

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

Description

This command displays ICMPv6 statistics. ICMPv6 generates error messages to report errors during processing and other diagnostic functions. ICMPv6 packets can be used in the neighbor discovery protocol.

Output

The following output is an example of ICMP6 information, and [Table 23: Output fields: ICMP6](#) 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 23: Output fields: ICMP6

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

Label	Description
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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

Description

This command displays ICMPv6 statistics for a specified interface. If the *interface-name* parameter is not entered, ICMPv6 statistics for all interfaces are displayed.

Parameters

interface-name

Displays entries associated with the specified IP interface name.

Output

[Table 24: Output fields: ICMP6 interface](#) describes the router ICMP6 interface output fields.

Table 24: Output fields: ICMP6 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

Label	Description
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] [family]] | [summary] | [exclude-services]
interface [ip-address | ip-int-name] statistics
```

Context

```
show>router
```

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 2F1C2T, and 7210 SAS-K 3SFP+ 8C.

Description

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



Note:

The 7210 SAS-K 2F6C4T does not support IPv6 parameters and options.

Parameters

ip-address

Displays only 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

Only displays the interface information associated with the specified IP interface name.

detail

Displays detailed IP interface information.

family

Displays information for the specified IP interface family.

Values ipv4 — Displays only the peers that have the IPv4 family enabled.
 ipv6 — Displays only the peers that are IPv6-capable.

summary

Displays summary IP interface information.

exclude-services

Displays IP interface information, excluding IP interfaces configured for customer services. Only core network IP interfaces are displayed.

statistics

Displays the number of transmitted, received, and discarded packets and bytes at the Layer 3 level for IP interface statistics. The collection of IP interface statistics is supported on any IP interface, regardless of encapsulation. Supported IP interfaces are access (IES, VPRN, routed VPLS, and spoke SDP) and network (IPv4, IPv6, and MPLS) interfaces. Discard statistics are only displayed for IPv4 interfaces.

Output

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

- Standard output: [Sample output](#), [Table 25: Output fields: router interface](#)
- Detailed output: [Sample output — detailed](#), [Table 26: Output fields: router interface detail](#)

Sample output

```
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
10.22.24.169/32    n/a
-----
Interfaces : 1
=====
A:ALU-7210#

A:ALA-A# show router interface 10.6.6.2
=====
Interface Table (Router: Base)
```

```

=====
Interface-Name      Adm      Opr      Mode      Port/SapId
IP-Address          PfxState
-----
to-PE-E            Up       Up       IES       1/1/3:0.*
10.6.6.2/24        n/a
-----
Interfaces : 1
=====
A:ALA-A#

```

Table 25: 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 has been 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 disabled
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: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 224.1.1.1
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 225.0.0.0/32
join-tlv-packing : N/A
data-delay-interval: 3 seconds
data-threshold  : 224.0.0.0/4 --> 1 kbps
=====

```

Table 26: 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

Label	Description
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

route-table

Syntax

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

Context

show>router

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

Description



Note:

7210 SAS-K 2F6C4T does not support the IPv6 parameters and options.

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

Displays information for the specified IP interface family.

Values *ipv4* — Displays only the peers that have the IPv4 family enabled.
 ipv6 — Displays only the peers that are IPv6-capable.

ip-prefix/prefix-length

Displays only those entries matching the specified IP prefix and prefix length.

Values

<i>ipv4-prefix</i>	a.b.c.d (host bits must be 0)
<i>ipv4-prefix-length</i>	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 (7210 SAS-D, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C)
	0 to 64 (7210 SAS-Dxp)

longer

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

exact

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

protocol name

Displays routes learned from the specified protocol.

Values local | static | ospf | isis | aggregate | bgp | bgp-vpn

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 27: Output fields: route table](#)
- Summary output: [Sample output — summary, Table 28: Output fields: route table summary](#)
- Summary output: [Sample output — BGP PIC, Table 29: Output fields: route-table BGP PIC](#)

Sample output

```
A:ALA# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix          Type      Proto    Age          Pref
  Next Hop[Interface Name]      Metric
-----
10.1.1.1/32          Remote    Static   00h22m29s    5
   6.6.6.1                               1
10.2.2.2/32          Local     Local    00h22m52s    0
  system                               0
10.5.5.0/24          Remote    Static   00h22m29s    5
   6.6.6.1                               1
10.6.6.0/24          Local     Local    00h22m30s    0
  to-PE-E                               0
-----
No. of Routes: 4
=====
A:ALA#
```

```
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 27: Output fields: 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
Pref	The route preference 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 7 7
=====
A:ALA-A#
```

Table 28: Output fields: route table summary

Label	Description
Active	The number of installed active routes in the FIB
Available	The number of uninstalled routes available in the RIB
Static	The number of static routes in the FIB

Label	Description
Direct	The number of direct routes (local subnets, including loopback) in the routing FIB
Total	The total number of routes

Sample output — BGP PIC



Note:

The following output applies to the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C only.

```
*A:Dut-A# show router route-table 10.77.77.77/32
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]
Next Hop[Interface Name]
Type      Proto    Age      Pref
Metric
-----
10.77.77.77/32 [B]
2.2.2.2 (tunneled)      Remote  BGP      00h01m04s 170
0
-----
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
*A:Dut-A# show router route-table 10.77.77.77/32 alternative
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]
Next Hop[Interface Name]
Alt-NextHop
Type      Proto    Age      Pref
Metric
Alt-
Metric
-----
10.77.77.77/32
2.2.2.2 (tunneled)      Remote  BGP      00h01m09s 170
0
10.77.77.77/32 (Backup)
3.3.3.3 (tunneled)      Remote  BGP      00h01m09s 170
0
-----
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
      Backup = BGP backup route
      LFA = Loop-Free Alternate nexthop
      S = Sticky ECMP requested
=====
```

Table 29: Output fields: route-table BGP PIC

Label	Description
Dest Prefix	The route destination address and mask
[Flags]	n — Number of times nexthop is repeated B — BGP backup route

Label	Description
	L — Loop-free alternate next hop S — Sticky ECMP requested
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
Proto	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
Pref	The route preference value for the route
No. of Routes	The number of routes displayed in the list
Alt-NextHop	The LFA next hop to use if the primary next hop is not reachable
Alt-Metric	The metric value for secondary next hops

rtr-advertisement

Syntax

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

Context

show>router

Platforms

7210 SAS-Dxp, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays IPv6 router advertisement information.

If no command line arguments are specified, all IPv6 router advertisement information is displayed, sorted by prefix.

Parameters

ip-int-name

Only displays the router advertisement information associated with the specified IP interface name.

family

Displays information for the specified IP interface family.

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

ipv6-prefix/prefix-length

Displays only those entries matching the specified IPv6 prefix and prefix length.

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 (hexadecimal) d - 0 to 255 (decimal)
	ipv6-prefix-length	0 to 128 (7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C) 0 to 64 (7210 SAS-Dxp)

Output

The following output is an example of router advertisement information, and [Table 30: 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 30: 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 is 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
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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

Description

This command displays the router static ARP table information.

If no command line arguments are specified, all static ARP table information is displayed, sorted by prefix.

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 information, and [Table 31: Output fields: 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 12.200.1.1
=====
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 31: Output fields: 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 *[[ip-prefix Imask] | [preference preference] | [next-hop ip-address] tag tag]*

Context

show>router

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

**Note:**

7210 SAS-K 2F6C4T does not support the IPv6 parameters and options.

This command displays the static entries in the routing table.

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

Parameters

ip-prefix/prefix-length

Displays only those entries matching the specified IP prefix and prefix length.

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)
ipv4-prefix-length —	0 to 32
ipv6-prefix-length —	0 to 128 (7210 SAS-D, 7210 SAS-K 2F1C2T, 7210 SAS-K 3SFP+ 8C)
	0 to 64 (7210 SAS-Dxp)

preference

Displays only static routes with the specified route preference.

Values 0 to 65535

ip-address

Displays only static routes with the specified next hop IP address.

Values

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

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 32: 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 32: 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. 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



Note:

7210 SAS-K 2F6C4T does not support the IPv6 parameters and options.

Description

This command displays the router status.

Output

The following output is an example of router status information, and [Table 33: Output fields: router status](#) describes the output fields.

Sample output

```
A:DUT-B>show>router# show router status
=====
Router Status (Router: Base)
=====
-----
Admin State      Oper State
-----
Router           Up           Up
Max Routes       10000
Total IPv4 Routes 5
ECMP Max Routes  1
=====
A:DUT-B>show>router#
```

Table 33: Output fields: router status

Label	Description
Router	The administrative and operational states for the router
Max Routes	The maximum number of routes configured for the system
Total Routes	The total number of routes in the route table

tunnel-table

Syntax

tunnel-table summary [ipv4 | ipv6]

tunnel-table [protocol *protocol*] [ipv4 | ipv6]

tunnel-table [*ip-prefix* [/mask] [alternative] [ipv4 | ipv6] detail

tunnel-table [*ip-prefix* [/mask] alternative

tunnel-table [*ip-prefix* [/mask] protocol *protocol*

tunnel-table [*ip-prefix* [/mask] sdp *sdp-id*

Context

show>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays tunnel table information.

Parameters

ip-prefix [mask]

Displays only those entries matching the specified IP prefix and mask.

summary ipv4 | ipv6

Displays IPv4 or IPv6 summary tunnel table information.

protocol *protocol* ipv4 | ipv6

Displays IPv4 or IPv6 protocol information for the specified protocol.

protocol

Displays information for the specified protocol.

Values local | static | ospf | isis | aggregate | bgp | bgp-vpn

alternative ipv4 | ipv6

Displays IPv4 or IPv6 LFA and backup tunnel table. information.

alternative

Displays LFA and backup tunnel table information.

detail

Displays detailed tunnel table information.

sdp *sdp-id*

Displays information for the specified SDP.

Values 1 to 17407

Output

The following output is an example of router tunnel table information, and [Table 34: Output fields: tunnel table](#) describes the output fields.

Sample output

```
*A:SAS-1# show router tunnel-table
=====
Tunnel Table (Router: Base)
=====
Destination      Owner Encap TunnelId  Pref  Nexthop      Metric
-----
10.0.0.1/32      sdp   GRE    10       5     10.0.0.1     0
10.0.0.1/32      sdp   GRE    21       5     10.0.0.1     0
10.0.0.1/32      sdp   GRE    31       5     10.0.0.1     0
10.0.0.1/32      sdp   GRE    41       5     10.0.0.1     0
=====
*A:SAS-1#
```

Table 34: Output fields: 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(s)
Nexthop	The next hop for the route destination
Metric	The route metric value for the route

2.6.2.3 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 information for the specified router instance.

Parameters

router-instance

Specifies the router name or service ID.

Values *service-id*: 1 to 2147483647

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 option(s) 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.

dhcp

Syntax

dhcp

Context

clear>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context clear and reset DHCP entities.

local-dhcp-server

Syntax

local-dhcp-server *server-name*

Context

clear>router>dhcp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command clears DHCP server data.

Parameters

server-name

Clears data for the specified local DHCP server.

declined-addresses

Syntax

declined-addresses *ip-address[/mask]*

declined-addresses pool *pool-name*

Context

clear>router>dhcp>local-dhcp-server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command clears declined DHCP addresses.

Parameters

pool-name

Specifies the declined pool name.

ip-address[/mask]

Specifies the declined IP address and mask.

leases

Syntax

leases *ip-address[/mask]* [*state*]

leases all [*state*]

Context

clear>router>dhcp>local-dhcp-server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command clears DHCP leases.

Parameters

ip-address[/mask]

Clears the specified IP address and mask.

state

Clears the state of the lease to be removed.

Values offered | stable | force-renew-pending | remove-pending | held | internal
| internal-orphan | internal-held | sticky

server-stats

Syntax

server-stats

Context

clear>router>dhcp>local-dhcp-server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command clears all server statistics.

statistics

Syntax

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

Context

clear>router>dhcp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command clears DHCP statistics.

Parameters

ip-int-name

Clears DHCP statistics for the specified interface name.

ip-address

Clears DHCP statistics for the specified IP address.

neighbor

Syntax

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

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

Context

clear>router

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

This command clears IPv6 neighbor information.

Parameters

all

Clears IPv6 neighbors.

ip-int-name

Clears the specified neighbor interface information, up to 32 characters.

ip-address

Clears the specified IPv6 neighbors.

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 (hexadecimal)

d - 0 to 255 (decimal)

router-advertisement

Syntax

router-advertisement all

router-advertisement [**interface** *interface-name*]

Context

clear>router

Platforms

7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

This command clears all IPV6 router advertisement counters.

Parameters

all

Clears all router advertisement counters for all interfaces.

interface interface-name

Clear router advertisement counters for the specified interface.

2.6.2.4 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 *service-id*: 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, except 7210 SAS-K 2F1C2T

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

Only displays the interface information associated with the specified IP address.

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

ip-int-name

Only 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.

The **no** form of this command disables debugging for IP packets. If a protocol-id was previously specified, it will be removed from the criteria.

Parameters

ip-int-name

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

ip-address

Displays only 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 only information associated with the packet header.

protocol-id

Specifies the value representing the IP protocol to debug. Well known protocol numbers include ICMP(1), TCP(6), UDP(17). The values can be expressed in decimal, hexadecimal, or binary.

Values 0 to 255

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

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 Filter policies

This chapter provides information about filter policies and management.

3.1 Filter policy configuration overview

Filter policies, also referred to as Access Control Lists (ACLs), are templates applied to services or access uplink ports to control network traffic into (ingress) or out of (egress) a service access port (SAP) or access uplink 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 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.

In 7210 SAS platforms, the available ingress and egress (egress CAM resources allocation is supported only on 7210 SAS-D and 7210 SAS-Dxp) CAM hardware resources can be allocated as per user needs for use with different filter criteria. By default on the 7210 SAS-D, 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 following configuration notes section 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. For the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, both IPv4 and IPv6 ingress and egress filter policy can be used simultaneously with a Layer 2 SAP. For the 7210 SAS-D and 7210 SAS-Dxp, both IPv4 and IPv6 filter policies can be used simultaneously on ingress only; either IPv4 or IPv6 filter policies can be used on egress.

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: IES IP interface in access-uplink mode on the 7210 SAS-D) for which IPv6 addressing is supported. Network filter policies control the forwarding and dropping of packets based on IP match criteria.

**Note:**

Non-IP packets are not hitting the IP filter policy, so the default action in the filter policy will not apply to these packets.

3.1.1 Service-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.

3.1.2 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 the following:

- scope
- default action
- description

Each filter entry contains the following:

- match criteria
- an action

3.1.2.1 Applying filter policies

Filter policies can be applied to specific service types:

- **Epipe**
Both MAC and IP filters are supported on an Epipe SAP.
- **IES**
Only IP filters are supported on IES SAP
- **VPLS**
Both MAC and IP filters are supported on a VPLS SAP.
- **VPRN**
Only IP filters are supported on VPRN SAP.

The following tables describe the support of filter policies on 7210 SAS platforms.

Table 35: Applying filter policies for 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

Service	IPv4 filter	IPv6 filter	MAC filter
Epipe	Epipe access SAP (ingress and egress), Epipe access-uplink SAP (ingress and egress)	Epipe access SAP (ingress and egress), Epipe access-uplink SAP (ingress and egress)	Epipe access SAP (ingress and egress), Epipe access-uplink SAP (ingress and egress)
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)	VPLS access SAP (ingress and egress), VPLS access-uplink SAP (ingress and egress)	Not supported	Not supported
RVPLS (RVPLS IES IP Interface)	Ingress Override filters (ingress)	Not supported	Not supported
IES	IES access SAP (ingress and egress), IES access-uplink SAP (ingress and egress)	IES access-uplink SAP (ingress and egress)	Not supported

Table 36: Applying filter policies for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Service	IPv4 filter	IPv6 filter	MAC filter
Epipe	Epipe access SAP (ingress and egress), Epipe access-uplink SAP (ingress and egress)	Epipe access SAP (ingress and egress), Epipe access-uplink SAP (ingress and egress)	Epipe access SAP (ingress and egress), Epipe access-uplink SAP (ingress and egress)
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)	VPLS access SAP (ingress and egress), VPLS	Not supported	Not supported

Service	IPv4 filter	IPv6 filter	MAC filter
	access-uplink SAP (ingress and egress)		
RVPLS (RVPLS IES IP Interface)	Ingress Override filters (ingress)	Not supported	Not supported
IES	IES access SAP (ingress and egress), IES access-uplink SAP (ingress and egress)	Not supported	Not supported
VPRN	VPRN interface SAP (ingress and egress)	Not supported	Not supported
Network port IP interface	Network port IP interface (ingress and egress)	Not supported	Not supported

3.1.2.2 ACL on range SAPs

The ACLs on VLAN range SAPs are supported only on ingress (for Epipe and VPLS services). The following table lists ACL support on Epipe and VPLAS services.

Table 37: Applying ACLs support on Epipe and VPLS services on 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C variants when using range SAPs

Types of filters	Epipe	VPLS
Ingress IP or IPv6	Yes	Yes
Ingress MAC	Yes	Yes
Egress IP	Yes	Yes
Egress MAC	Yes	Yes

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 ingress policies can be applied on SAP created on access ports or access uplink ports.

- **SAP egress**

Filter policies applied on 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. SAP egress policies can be applied on both access ports and access uplink ports.

- **IES IP interfaces**

IP filter policies are applied to IES SAPs.

- **network ingress**

IP filter policies are applied to network ingress IP interfaces. This is supported only on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

- **network egress**

IP filter policies are applied to network egress IP interfaces. This is supported only on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

The following table lists the Packet Fields available for match in QoS classification policy and ACL policy for different SAPs.

Table 38: Packet fields for match in QoS classification policy and ACL policy

Ingress SAP type	Packet contents (only Ethernet-II frames)	MAC address match	Inner VLAN ID and Dot1p match ⁹	Outer VLAN ID and Dot1p match ⁹	Etype match	IPv4/IPv6 criteria match
NULL SAP	Null tag	Yes	No	No	Yes	Yes
	Priority tag (both VID and Dot1p)	Yes	No	Yes	Yes	Yes
	Single tag	Yes	No	Yes	Yes	Yes
	Two tags	Yes	Yes	Yes	Yes	Yes
	Three or more tags	Yes	Yes	Yes	No	No
Dot1q SAP (includes Dot1q explicit null SAP and Dot1q Default SAP)	Null tag	Yes	No	No	Yes	Yes
	Priority tag (both VID and Dot1p)	Yes	No	Yes	Yes	Yes
	Single tag	Yes	No	Yes	Yes	Yes
	Two tags	Yes	Yes	Yes	Yes	Yes
	Three or more tags	Yes	Yes	Yes	No	No
Dot1q SAP (includes Dot1q	Null tag	Invalid	Invalid	Invalid	Invalid	Invalid

⁹ VLAN tag matching is supported only on 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

Ingress SAP type	Packet contents (only Ethernet-II frames)	MAC address match	Inner VLAN ID and Dot1p match ⁹	Outer VLAN ID and Dot1p match ⁹	Etype match	IPv4/IPv6 criteria match
SAP, Dot1q range SAP)	Priority tag (both VID and Dot1p)	Invalid	Invalid	Invalid	Invalid	Invalid
	Single tag	Yes	No	Yes	Yes	Yes
	Two tags	Yes	Yes	Yes	Yes	Yes
	Three or more tags	Yes	Yes	Yes	No	No
QinQ SAP - 0.* SAP (matches only null and priority tag packets)	Null tag	Yes	No	No	Yes	Yes
	Priority tag (both VID and Dot1p)	Yes	No	Yes	Yes	Yes
	Single tag	Invalid	Invalid	Invalid	Invalid	Invalid
	Two tags	Invalid	Invalid	Invalid	Invalid	Invalid
	Three or more tags	Invalid	Invalid	Invalid	Invalid	Invalid
QinQ SAP (*.*) Default QinQ SAP)	Null tag	Yes	No	No	Yes	Yes
	Priority tag (both VID and Dot1p)	Yes	No	Yes	Yes	Yes
	Single tag	Yes	No	Yes	Yes	Yes
	Two tags	Yes	Yes	Yes	Yes	Yes
	Three or more tags	Yes	Yes	Yes	No	No
QinQ SAP (includes Q1.* SAP)	Null tag	Invalid	Invalid	Invalid	Invalid	Invalid
	Priority tag (both VID and Dot1p)	Invalid	Invalid	Invalid	Invalid	Invalid
	Single tag	Yes	No	Yes	Yes	Yes
	Two tags	Yes	Yes	Yes	Yes	Yes

⁹ VLAN tag matching is supported only on 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.


Ingress SAP type	Packet contents (only Ethernet-II frames)	MAC address match	Inner VLAN ID and Dot1p match ⁹	Outer VLAN ID and Dot1p match ⁹	Etype match	IPv4/IPv6 criteria match
	Three or more tags	Yes	Yes	Yes	No	No
QinQ SAP (includes Q1.0 SAP)	Null tag	Invalid	Invalid	Invalid	Invalid	Invalid
	Priority tag (both VID and Dot1p)	Invalid	Invalid	Invalid	Invalid	Invalid
	Single tag	Yes	No	Yes	Yes	Yes
	Two tags (inner tag is a priority tag)	Yes	Yes	Yes	Yes	Yes
	Two tags (inner tag is not a priority tag)	Invalid	Invalid	Invalid	Invalid	Invalid
	Three or more tags	Yes	Yes	Yes	No	No
QinQ SAP (includes Q1.Q2 SAP)	Null tag	Invalid	Invalid	Invalid	Invalid	Invalid
	Priority tag (both VID and Dot1p)	Invalid	Invalid	Invalid	Invalid	Invalid
	Single tag	Invalid	Invalid	Invalid	Invalid	Invalid
	Two tags (inner tag is a priority tag)	Invalid	Invalid	Invalid	Invalid	Invalid
	Two tags (inner tag is not a priority tag)	Yes	Yes	Yes	Yes	Yes
	Three or more tags	Yes	Yes	Yes	No	No

⁹ VLAN tag matching is supported only on 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

3.2 Creating and applying filter policies

The following figure shows the steps for creating and applying filter policies.

Figure 4: Creating and applying filter policies



3.2.1 Packet matching criteria

As few or as many match parameters can be specified as required, but all conditions must be met in order 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 IPv4 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).

The IPv6 addressing scheme consists of 128 bits expressed in compressed representation of IPv6 addresses (RFC 1924, *A Compact Representation of IPv6 Addresses*).

- 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C, 7210 SAS-D, and 7210 SAS-Dxp support the use of either IPv6 64-bit address match or IPv6 128-bit address match. Use of IPv6 64-bit address in the match criteria provides better scale but provides lesser IPv6 header fields for match criteria. Use of a IPv6 128-bit address in the match criteria provides lesser scale but more IPv6 header fields for match criteria.

- **Destination IP address and mask**

Destination IP address and mask values can be entered as search criteria. A choice similar to that available for source IPv6 addresses is also available for destination IPv6 addresses.

- **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.
- **DSCP marking**
Entering a DSCP marking enables the filter to search for the DSCP marking specified in this field. See [Table 39: DSCP name to DSCP value table](#).
- **ICMP code**
Entering an ICMP code allows the filter to search for matching ICMP codes in the ICMP header.
- **ICMP type**
Entering an ICMP type allows the filter to search for matching ICMP types in the ICMP header.
- **Extension header present**
Enabling this match criterion allows matching of IPv6 packets that have any of the well-known extension headers in the IPv6 header. This match criterion is not supported for IPv6 filters on 7210 SAS-Dxp.
- IPv4 filters created in the mode to use IPv6 resources cannot be applied at the egress SAP. Similarly, IPv4 filters created in the mode to use IPv6 resources will fail to match fragment options.
- **Fragmentation**
Enabling fragmentation allows matches to occur if packets have either the more fragment (MF) 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 EOL 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 and/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 and/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. This is not supported on 7210 SAS-K devices.

- **Ethertype**

Entering an Ethernet type II Ether type 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 Ether type accepts decimal, hex, or binary in the range of 1536 to 65535.

- **Outer Dot1p (Only on 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C)**

Entering the Outer Dot1p value or range (using the mask) allows the filter to search for frames whose outermost Dot1p (that is, the Dot1p in the outermost VLAN tag of the packet) matches the Dot1p value configured. The Dot1p value and mask accepts decimal values in the range 0 to 7.

- **Inner Outer Dot1p (Only on 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C)**

Entering the Inner Dot1p value or range (using the mask) allows the filter to search for frames whose inner Dot1p (that is, the Dot1p in the VLAN tag immediately following the outermost VLAN tag of the packet) matches the Dot1p value configured. The Dot1p value and mask accepts decimal values in the range 0 to 7.

3.2.1.1 DSCP values

The following table describes DSCP names and associated DSCP values.

Table 39: 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	*	

DSCP name	Decimal DSCP value	Hexadecimal DSCP value	Binary DSCP value
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	*	
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		

DSCP name	Decimal DSCP value	Hexadecimal DSCP value	Binary DSCP value
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		
cp54	54		
cp55	55		
cp56	56		
cp57	57		
nc2	58	*	(cs7)
cp60	60		
cp61	61		
cp62	62		

3.2.2 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's specified action is performed (drop or forward).
- If a packet does not match all of the entry criteria, the policy's 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's 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 5: Filtering process example



3.2.3 Applying filters

This section provides information about applying filters.

3.2.3.1 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.

3.2.3.2 Applying a filter to an IES interface

An IP filter can be applied to an IES SAP. 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.

3.2.3.3 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.

3.3 Configuration notes

**Note:**

See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Services Guide* for service specific ACL support and restrictions.

The following information describes filter implementation caveats:

- 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.
- On the 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, IP filters applied on an IES SAP cannot match against IP packets containing IP options.
- 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.
- On the 7210 SAS-D and 7210 SAS-Dxp, ingress filter CAM resources used to match packet fields are shared with other features such as SAP ingress QoS, CFM UP MEP, and G8032. By default software

assigns a fixed amount of resources for use by ingress ACLs. User has an option to either increase this by taking away resources from other features or decrease by taking away resources from ingress ACLs. The number of ACLs that can be supported is directly dependent on the amount of resources allocated toward ingress ACLs.

- On the 7210 SAS-D and 7210 SAS-Dxp when a filter policy is created with the option **ipv6-64bit-address**, the entries can only use only the IPv6 src-ip and IPv6 dst-ip fields in the match criteria.
- On the 7210 SAS-D and 7210 SAS-Dxp when a filter policy is created with the option **ipv6-128bit-address**, the entries can use the IPv6 **src-ip**, IPv6 **dst-ip**, IPv6 DSCP, TCP/UDP port numbers (source and destination port), ICMP code and type, and TCP flags fields in the match criteria.
- On the 7210 SAS-D and 7210 SAS-Dxp 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.
- On the 7210 SAS-D and 7210 SAS-Dxp, the available ingress CAM hardware resources can be allocated as per user needs for use with different filter criteria using the commands under the **configure> system> resource-profile> ingress-internal-tcam> acl-sap-ingress** context. 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).
- On the 7210 SAS-D and 7210 SAS-Dxp, the available egress CAM hardware resources can be allocated as per user needs for use with different filter criteria using the commands under the **configure> system> resource-profile> egress-internal-tcam> acl-sap-egress** context. 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).
- On the 7210 SAS-D and 7210 SAS-Dxp IPv6 ACLs and MAC QoS policies cannot coexist on the SAP.
- On the 7210 SAS-D and 7210 SAS-Dxp 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.
- Only the 7210 SAS-K allows for use of outer VLAN ID and inner VLAN ID for match in MAC criteria with both ingress and egress ACLs. Other 7210 SAS platforms do not support use of outer and inner VLAN ID field for match in the MAC criteria.

3.3.1 MAC filters

The following are configuration notes for 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, routable VPLS or IES services.
- 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. In the 7210 SAS, the default frame-format is "Ethernet-II"

Table 40: MAC match criteria exclusivity rules

Frame format	Etype
Ethernet – II	Yes
802.3	No
802.3 – snap	No
802.3-llc	No

3.3.2 IP filters

The following are configuration notes for IP 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 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.

3.3.3 IPv6 filters

The following are configuration notes for 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.

3.3.3.1 Resource usage for ingress filter policies for 7210 SAS-D and 7210 SAS-Dxp

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 (for example, 256 entries per chunk on 7210 SAS-D).



Note:

The number of entries for each chunk or slice is different for both **ingress-internal-tcam** resource pool and **egress-internal-tcam** resource pool for different platforms.

The usage of these entries by different type of match criteria follows. In the following examples, it is assumed that a chunk/slice has 256 entries considering 7210 SAS-D. The example and the computation needs to be modified suitably for other platforms with different number of entries per chunk/slice.

- **mac-criteria**

User needs to allocate resources for **mac-criteria** from the filter resource pool by using the command **configure>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 **configure>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 ip(v4)-criteria from the filter resource pool by using the command **configure>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**.

- **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 **configure>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 following CLI description for filter for more information. The usage is same as the ipv4 and **mac-criteria**. An IPv6 128 bit address uses 2 entries from the chunk for every match entry configured in filter policy, whereas, an IP filter uses only one entry from the chunk for every entry configured.

- **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 **configure>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 (2) 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 **configure>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. When a chunk is allocated to IPv6 criteria, the software automatically adjusts the number of available entries in that chunk to 128, instead of 256, because 2 entries are needed to match IPv6 fields.

The users can use **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.

3.3.3.2 Resource usage for egress filter policies (supported only for 7210 SAS-D and 7210 SAS-Dxp)

When the user allocates resources for use by filter policies using the **configure system resource-profile egress-internal-tcam** CLI commands, the system allocates resources in chunks of 128 entries from the egress internal tcam pool in hardware. The usage of these entries by different type of match criteria is as follows:

- **mac-criteria**

The user needs to allocate resources for using **mac-criteria** using the **configure system resource-profile egress-internal-tcam acl-sap-egress mac-match-enable 2** or **configure system resource-profile egress-internal-tcam acl-sap-egress mac-ipv4-match-enable 2** command or the **configure system resource-profile egress-internal-tcam acl-sap-egress mac-ipv6-64bit-match-enable 2** command. 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 the **configure system resource-profile egress-internal-tcam acl-sap-egress mac-match-enable 2** command, the user configures two chunks for use by **mac-criteria**, allowing a total of 256 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 250 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, 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 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 4 SAPs using filter policies, such that 2 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 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 **configure system resource-profile egress-internal-tcam acl-sap-egress mac-ipv4-match-enable** command. The resource usage explanation precedes.

- **ipv6-criteria using ipv6-64-bit addresses**

The user need to allocate resources using the **configure system resource-profile egress-internal-tcam acl-sap-egress mac-ipv6-64bit-match-enable** command. The resource usage explanation precedes.

- **ipv6-criteria using ipv6-128-bit addresses**

The user need to allocate resources using the **configure system resource-profile egress-internal-tcam acl-sap-egress ipv6-128bit-match-enable** command. 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 the **configure system resource-profile egress-internal-tcam acl-sap-egress ipv6-128bit-match-enable 2** command, to configure two chunks for use by **ipv6-128bit-criteria**. This allows for a total of 128 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 **tools dump system-resources** command to know the current usage and availability.

3.3.4 Ingress filter policy resource usage: 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

When the user allocates resources from the ingress CAM resource pool for use by filter policies using the **configure system resource-profile ingress-internal-tcam acl-sap-ingress** command, the system allocates resources in chunks of fixed-size entries (512 entries per chunk on 7210 SAS-K). Resources must be allocated using these commands before associating a filter policy with the SAP, otherwise the command returns an error. The usage of these entries by different types of match criteria follow:

- **mac-criteria, ipv4-criteria and ipv6-criteria with 64-bit-address**

User needs to allocate resources, in terms of number of slices, for filter policies that use mac criteria, ipv4 criteria and ipv6 64-bit criteria from the ingress internal tcam resource pool using the **configure system resource-profile ingress-internal-tcam acl-sap-ingress** command. The entries allocated are shared by filter policies that use any of these criteria. Each filter entry configured in the policy takes away a single resource from the pool allocated for filter policies.

- **ipv6-criteria with 128-bit address**

User needs to allocate resources, in terms of number of slices, for filter policies that use ipv6 128-bit criteria from the ingress internal tcam resource pool using the **configure system resource-profile ingress-internal-tcam acl-sap-ingress mac-ipv4-ipv6-128-match-enable** command. User can allocate all the slices allocated for the filter policies (using the **configure system resource-profile ingress-internal-tcam acl-sap-ingress** command) for use by ipv6 criteria with 128-bit addresses or allocation only a portion of it. The entries allocated are used by filter policies that use ipv6 criteria with 128-bit addresses. Each filter entry configured in the policy takes away two (2) resources from the pool. Software uses these resources also for mac criteria, ipv4 criteria, and ipv6 criteria with 64-bit address. Irrespective of the criteria, two (2) resources are taken for each entry configured on the filter policy.

Use the **tools dump system-resources** command to know the current usage and availability.

3.4 Configuring filter policies with CLI

This section provides information to configure filter policies using the CLI.

3.5 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

Example: Configuration output for ingress policy

The following is a sample configuration output of allocation of ingress internal CAM resources for ingress policy for 7210 SAS-D.

```
*A:SASD>config>system>res-prof>ing-internal-tcam# info detail
-----
      acl-sap-ingress 2
        ipv4-match-enable max
        no ipv6-64-only-match-enable
        no ipv4-ipv6-128-match-enable
        mac-match-enable 2
      exit
      no eth-cfm
-----
*A:SASD>config>system>res-prof>ing-internal-tcam# acl-sap-ingress
```

Example: Configuration output for egress policy

The following is a sample configuration output of allocation of egress internal CAM resources for egress policy for 7210 SAS-D.

```
A:SASD>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:SASD>config>system>res-prof>egr-internal-tcam# acl-sap-egress
```


Example: Configuration output of an IP filter policy

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.

```
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 6: Applying an IP filter to an ingress interface



3.6 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.

3.6.1 Allocating resources for filter policies (ingress and egress)

The following provides an example of allocation of CAM hardware resources for use with filter policies that use IPv4 and MAC criteria:

3.6.2 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
- 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

3.6.2.1 IP filter policy

Example: 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#
```

3.6.2.2 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, as follows:

- 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
```

Example: 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#
```

3.6.2.3 IP entry matching criteria

Use the following syntax to configure IP filter matching criteria:

Example: 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#
```

3.6.3 Creating an IPv6 filter policy (applicable only for 7210 SAS-D and 7210 SAS-Dxp)

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

3.6.3.1 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, as follows:

- Enter an IPv6 filter entry ID. The system does not dynamically assign a value.
- Assign an action, either drop or forward.

- Specify matching criteria.

Example: 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 1::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#
```

3.6.4 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

3.6.4.1 MAC filter policy

Example: 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#
```

3.6.4.2 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, as follows:

- Enter a filter entry ID. The system does not dynamically assign a value.
- Assign an action, either drop or forward.
- Specify matching criteria.

Example: AC filter entry configuration output

```
A:siml>config>filter# info
-----
      mac-filter 90 create
        entry 1 create
          description "allow-104"
          match
          exit
          action
            drop
        exit
      exit
A:siml>config>filter#
```

3.6.4.3 MAC entry matching criteria

Example: 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
-----
```

3.6.4.4 Apply IP and MAC filter policies

The following example shows an example of applying 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}
```

Example

The following is a sample output for IP and MAC filters assigned to an ingress and egress SAP.

```
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#
```

3.6.4.5 Apply filter policies to an IES interface

IP filter policies can be applied to an IP interface created in an IES service. These filter policies apply to the routed management traffic.

```
config>service>ies# interface ip-int-name
  address ip-address
  sap sap-id
  ingress
    filter ip ip-filter-id
```

Example

The following is a sample output for an IP filter applied to an IES sap at ingress.

```
A:ALA-48>config>service>ies# info
-----
      interface "to-104" create
        address 10.1.2.1/24
        sap lag-2:0.* create
          ingress
            filter ip 10
        exit
      exit
      ...
-----
A:ALA-48>config>service>ies#
```

3.7 Filter management tasks

This section discusses the filter policy management tasks.

3.7.1 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-number new-entry-number
  mac-filter filter-id
  renum old-entry-number new-entry-number
```

Example: Command usage to renumber filter entries

```
config>filter>ip-filter# renum 10 15
config>filter>ip-filter# renum 20 10
config>filter>ip-filter# renum 40 1
```

Example: Reordered filter entries

The following is a sample original filter entry order on the left side and the reordered filter entries on the right side.

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
```

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
```

<pre> 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# </pre>	<pre> 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# </pre>
--	---

3.7.2 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: Command usage to modify an IP filter policy

```

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#

```

Example: 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#
```

3.7.3 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: Command usage to modify a MAC filter policy

```
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
```

Example: 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#
```

3.7.4 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.

3.7.4.1 From an ingress SAP

Use the following syntax to remove a filter from an ingress SAP.

```
config>service# [epipe | ies | 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
```

3.7.4.2 From an egress SAP

Use the following syntax to remove a filter from an egress SAP.

```
config>service# [epipe | ies | 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
```

3.7.4.3 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
```

3.7.5 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]
```

Example: Command usage

The following shows command usage to copy an existing IP filter ("11") to create a new filter policy ("12").

```
config>filter# copy ip-filter 11 to 12
```

Example: Configuration output

```
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#
```

3.8 Filter command reference

3.8.1 Command hierarchies

- [Configuration commands](#)
 - [IP filter policy commands](#)
 - [IPv6 filter policy commands for 7210 SAS-D and 7210 SAS-Dxp](#)
 - [IPv6 filter policy commands for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C](#)
 - [MAC filter policy commands for 7210 SAS-D and 7210 SAS-Dxp](#)
 - [MAC filter policy commands for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP + 8C](#)
 - [Generic filter commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Monitor commands](#)

3.8.1.1 Configuration commands

3.8.1.1.1 IP filter policy commands

```

config
- filter
- ip-filter filter-id [use-ipv6-resource] [create]
- no ip-filter filter-id
- default-action {drop | forward}
- description description-string
- no description
- filter-name filter-name
- no filter-name
- renum 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{drop}
- action forward
- no action
- 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
- 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}
- no src-port
- tcp-ack {true | false}
- no tcp-ack
- tcp-syn {true | false}
- no tcp-syn

```

3.8.1.1.2 IPv6 filter policy commands for 7210 SAS-D and 7210 SAS-Dxp

```

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 [drop]
  - action forward
  - no action
  - 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
- renam old-entry-id new-entry-id
- scope {exclusive | template}
- no scope

```

3.8.1.1.3 IPv6 filter policy commands for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

```

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 [drop]
      - action forward
      - no action
      - 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
- fragment {true | false | first-only | non-first-only}
- no fragment
- eh-present {true | false}
- no eh-present
- 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

```

3.8.1.1.4 MAC filter policy commands for 7210 SAS-D and 7210 SAS-Dxp

```

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

```

3.8.1.1.5 MAC filter policy commands for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

```

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
- dst-mac ieee-address [ieee-address-mask]
- no dst-mac
- etype 0x0600..0xffff
- no etype
- inner-dot1p dot1p-value [dot1p-mask]
- no inner-dot1p
- inner-tag value [vid-mask]
- no inner-tag
- outer-dot1p dot1p-value [dot1p-mask]
- no outer-dot1p
- no outer-tag
- outer-tag value [vid-mask]
- 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

```

3.8.1.1.6 Generic filter commands

```

config
- filter
- copy ip-filter | mac-filter src-filter-id [src-entry src-entry-id] to dst-filter-id [dst-entry dst-entry-id] [overwrite]

```

3.8.1.2 Show commands

```

show
- filter
- 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]}

```

3.8.1.3 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]
```

3.8.1.4 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]
  - mac mac-filter-id entry entry-id [interval seconds] [repeat repeat] [absolute | rate]
```

3.8.2 Command descriptions

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

3.8.2.1 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 commands](#)
- [MAC filter match criteria commands](#)
- [Policy and entry maintenance commands](#)

3.8.2.1.1 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 stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context 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, and so on), the entire string must be enclosed within double quotes.

3.8.2.1.2 Global filter commands

ip-filter

Syntax

```
[no] ip-filter filter-id [use-ipv6-resource] [create]
```

Context

```
config>filter
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an IP filter policy.

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 as long as the scope of the policy is template.

Any changes made to the existing policy, using any of the sub-commands, 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 then 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 dump 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 where it is applied.

Parameters

filter-id

Specifies the IP filter policy ID number.

Values 1 to 65535

create

Specifies the keyword required when first creating the configuration context. After the context is created, one can navigate into the context without the **create** keyword.

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.

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 IPv6 filter creation, 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, the 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, 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 matching when using this option.

create

Specifies the keyword required when first creating the configuration context. After the context is created, one can navigate into the context without the **create** keyword.

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:

A MAC filter policy cannot be applied to network ports on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

Any changes made to the existing policy, using any of the sub-commands, 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.

3.8.2.1.3 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 all packets will be dropped unless there is a specific filter entry which causes the packet to be forwarded.

forward

Specifies 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). 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.

3.8.2.1.4 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 enables creation or editing of 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 given entry IDs in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

Values 1 to 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.

3.8.2.1.5 IP filter entry commands

action

Syntax

action [drop]

action forward

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 specifies to match packets with a specific IP option or a range of IP options in the first option of the IP header as an IP filter match criterion. The **action** keyword must be entered and a keyword specified in order for the entry to be active.

Multiple action statements entered will overwrite previous actions parameters when defined.

The **no** form of this command removes the specified **action** statement. The filter entry is considered incomplete and therefore rendered inactive without the **action** keyword.

Parameters

drop

Specifies packets matching the entry criteria will be dropped.

forward

Specifies packets matching the entry criteria will be forwarded.

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, 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**

Specifies an IP protocol to be used as an IP filter match criterion. The protocol type, such as TCP or UDP, is identified by its respective protocol number.

protocol-id

Specifies the decimal value representing the IP protocol to be used as an IP filter match criterion. Common protocol numbers include ICMP(1), TCP(6), and UDP(17) (see the following table). The value can be expressed in decimal, hexadecimal, or binary.

Values 0 to 255

Table 41: IP protocol IDs and descriptions

Protocol ID	Protocol	Description
1	icmp	Internet Control Message
2	igmp	Internet Group Management
4	ip	IP in IP (encapsulation)
6	tcp	Transmission Control
8	egp	Exterior Gateway Protocol
9	igp	Any private interior gateway
17	udp	User Datagram
27	rdp	Reliable Data Protocol

Protocol ID	Protocol	Description
45	idrp	Inter-Domain Routing Protocol
46	rsvp	Reservation Protocol
80	iso-ip	ISO Internet Protocol
88	eigrp	EIGRP
89	ospf-igp	OSPF/IGP
97	ether-ip	Ethernet-within-IP Encapsulation
98	encap	Encapsulation Header
102	pnni	PNNI over IP
103	pim	Protocol Independent Multicast
112	vrrp	Virtual Router Redundancy Protocol
115	l2tp	Layer Two Tunneling Protocol
118	stp	Schedule Transfer Protocol
123	ptp	Performance Transparency Protocol
124	isis	ISIS over IPv4
126	crtp	Combat Radio Transport Protocol
127	crudp	Combat Radio User Datagram

3.8.2.1.6 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** keyword.

Parameters

drop

Specifies packets matching the entry criteria will be dropped.

forward

Specifies packets matching the entry criteria will be forwarded.

match

Syntax

match

no match

Context

config>filter>mac-filter>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the context for entering or editing 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.

ethernet_II

Specifies the frame type is Ethernet Type II.

3.8.2.1.7 IP filter match criteria commands

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

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 IP address range to be used as an IP filter match criterion.

To match on the destination IP 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 destination IPv4 address match criterion.

Default

none

Parameters

ip-address

Specifies the IP 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*

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.



Note:

An entry containing L4 match criteria will not match non-initial (2nd, 3rd, and so on) fragments of a fragmented packet because only the first fragment contains the L4 information.

The **no** form of this command removes the destination port match criterion.

Default

none

Parameters

dst-port-number

Specifies the destination port number to be used as a match criteria expressed as a decimal integer.

Values 1 to 65535

eh-present

Syntax

eh-present {true | false}

no eh-present

Context

config>filter>ipv6-filter>entry>match

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

This command allows the user to specify if the presence of the IPv6 extension header should be used to match an IPv6 packet.

The **no** form of this command removes the match criterion.

Default

no eh-present

Parameters

true

Specifies to match an IPv6 packet with an extension header.

false

Specifies to match an IPv6 packet without an extension header.

fragment

Syntax

fragment {true | false}

no fragment

Context

config>filter>ip-filter>entry>match

Platforms

7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

This command configures fragmented or non-fragmented IPv4 packets as IP filter match criteria.



Note:

An entry containing L4 match criteria will not match non-initial (2nd, 3rd, and so on) fragments of a fragmented packet because only the first fragment contains the L4 information.

The **no** form of this command removes the match criterion.

Default

no fragment

Parameters

true

Specifies to match on all fragmented IPv4 packets. A match will occur for all packets that have either the more fragment (MF) bit set or have the Fragment Offset field of the IPv4 header set to a non-zero value.

false

Specifies to match on all non-fragmented IPv4 packets. Non-fragmented IPv4 packets are packets that have the MF bit set to zero and have the Fragment Offset field also set to zero.

fragment

Syntax

fragment {true | false | first-only | non-first-only}

no fragment

Context

config>filter>ipv6-filter>entry>match

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

This command configures fragmented or non-fragmented IPv6 packets as IP filter match criteria.



Note:

An entry containing L4 match criteria will not match non-initial (2nd, 3rd, and so on) fragments of a fragmented packet because only the first fragment contains the L4 information.

The **no** form of this command removes the match criterion.

Default

no fragment

Parameters

true

Specifies to match on all fragmented IPv6 packets. A match will occur for all packets that have either the more fragment (MF) bit set or have the Fragment Offset field of the IPv6 header set to a non-zero value.

false

Specifies to match on all non-fragmented IPv6 packets. Non-fragmented IPv6 packets are packets that have the MF bit set to zero and have the Fragment Offset field also set to zero.

first-only

Specifies to match if a packet is an initial fragment of a fragmented IPv6 packet.

non-first-only

Specifies to match if a packet is a non-initial fragment of a fragmented IPv6 packet.

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 the ICMP code field in the ICMP header of an IP packet as a filter match criterion.



Note:

An entry containing L4 match criteria will not match non-initial (2nd, 3rd, and so on) fragments of a fragmented packet because only the first fragment contains the L4 information.

For an IPv4 filter, this command applies only if the protocol match criterion specifies ICMP (1).

For an IPv6 filter, this command applies only if the next header match criterion specifies **ipv6-icmp** (58).

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 *icmp-code-number* or *icmp-code-keyword*

icmp-code-number

Specifies the ICMP code number in decimal, hexadecimal, or binary, to be used as a match criterion.

Values 0 to 255 (decimal)
0x0 to 0xFF (hexadecimal)
0b0 to 0b11111111 (binary)

icmp-code-keyword

Specifies the ICMP code keyword to be used as a match criterion.

Values none | no-route-to-destination | comm-with-dest-admin-prohibited |
beyond-scope-scr-addr | address-unreachable | port-unreachable

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 packet as a filter match criterion.



Note:

An entry containing L4 match criteria will not match non-initial (2nd, 3rd, and so on) fragments of a fragmented packet because only the first fragment contains the L4 information.

For an IPv4 filter, this command applies only if the protocol match criterion specifies ICMP (1).

For an IPv6 filter, this command applies only if the next header match criterion specifies **ipv6-icmp** (58).
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 *icmp-type-number* or *icmp-type-keyword*

icmp-type-number

Specifies the ICMP type number in decimal, hexadecimal, or binary, to be used as a match criterion.

Values 0 to 255 (decimal)
0x0 to 0xFF (hexadecimal)
0b0 to 0b11111111 (binary)

icmp-type-keyword

Specifies the ICMP type keyword to be used as a match criterion.

Values none | dest-unreachable | packet-too-big | time-exceeded, parameter-problem | echo-request | echo-reply | multicast-listen-query | multicast-listen-report | multicast-listen-done | router-solicitation | router-advt | neighbor-solicitation | neighbor-advertisement | redirect-message | router-renumbering | icmp-node-info-query | icmp-node-info-req | inv-nd-solicitation | inv-nd-adv-message

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 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.

false

Specifies matching on IP packets that do not have any option field present in the IP header.

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*

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.



Note:

An entry containing L4 match criteria will not match non-initial (2nd, 3rd, and so on) fragments of a fragmented packet because only the first fragment contains the L4 information.

The **no** form of this command removes the source port match criterion.

Default

no src-port

Parameters

src-port-number

Specifies the source port number to be used as a match criteria, expressed as a decimal integer.

Values 0 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.



Note:

An entry containing L4 match criteria will not match non-initial (2nd, 3rd, and so on) fragments of a fragmented packet because 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.

The SYN bit is normally set when the source of the packet needs to initiate a TCP session with the specified destination IP address.

**Note:**

An entry containing L4 match criteria will not match non-initial (2nd, 3rd, and so on) fragments of a fragmented packet because only the first fragment contains the L4 information.

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.

3.8.2.1.8 MAC filter match criteria commands

```
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.

Egress Dot1p values used for matching will correspond to the Dot1p values used for remarking.

Default

no dot1p

Parameters

ip-value

Specifies the IEEE 802.1p value in decimal.

Values 0 to 7

mask

Specifies a 3-bit mask that can be configured using the following formats:

Table 42: 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 (decimal)

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

This command 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 43: 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 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.

For the 7210 SAS-D, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C platforms, the dataplane 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

inner-dot1p

Syntax

inner-dot1p *value* [*vid-mask*]

no inner-dot1p

Context

config>filter>mac-filter>entry>match

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

This command configures the Dot1p value to be used to match against the Dot1p value in the inner tag (the one that follows the outermost tag in the packet) of the received packet.

The **no** form of this command removes the previously entered Dot1p value as the match criteria.

Default

no inner-dot1p

Parameters

dot1p-value

Specifies the Dot1p value to match.

Values 0 to 7

dot1p-mask

Specifies the mask value to match a range of Dot1p values. The value can be expressed in decimal or binary.

Values 0 to 7

inner-tag

Syntax

inner-tag *value* [*vid-mask*]

no inner-tag

Context

config>filter>mac-filter>entry>match

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

This command configures the VLAN value to be used to match against the VLAN value in the inner tag (the one that follows the outermost tag in the packet) of the received packet.

The optional *vid_mask* is defaulted to 4095 (exact match) but may be specified to allow pattern matching. The masking operation is $((value \& vid_mask) = (tag \& vid_mask))$. A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6.

The **no** form of this command removes the previously entered VLAN tag value as the match criteria.

Default

no inner-tag

Parameters

value

Specifies the VLAN value to use for the match

Values 0 to 4095 (decimal) or 0x0 to 0xFFF (hexadecimal)

vid-mask

Specifies the mask value to match a range of VLAN values.

Values 0 to 4095 (decimal) or 0x0 to 0xFFF (hexadecimal)

outer-dot1p

Syntax

outer-tag *value* [*vid-mask*]

no outer-tag

Context

config>filter>mac-filter>entry>match

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

The command configures the Dot1p value to be used to match against the Dot1p value in the outermost tag of the received packet.

The **no** form of this command removes the previously entered Dot1p value as the match criteria.

Default

no outer-dot1p

Parameters

dot1p-value

Specifies the Dot1p value to match.

Values 0 to 7

dot1p-mask

Specifies the mask value to match a range of Dot1p values. The value can be expressed in decimal or hexadecimal.

Values 0 to 7

outer-tag

Syntax

outer-tag *value* [*vid-mask*]

no outer-tag

Context

config>filter>mac-filter>entry>match

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Description

This command configures the VLAN value to be used to match against the VLAN value in the inner tag (the one that follows the outermost tag in the packet) of the received packet.

The optional *vid_mask* is defaulted to 4095 (exact match) but may be specified to allow pattern matching. The masking operation is $((value \& vid_mask) = (tag \& vid_mask))$. A value of 6 and a mask of 7 would match all VLANs with the lower 3 bits set to 6.

The **no** form of this command removes the previously entered VLAN tag value as the match criteria.

Default

no outer-tag

Parameters

value

Specifies the VLAN value to use for the match

Values 0 to 4095 (decimal) or 0x0 to 0xFFFF (hexadecimal)

vid-mask

Specifies the mask value to match a range of VLAN values.

Values 0 to 4095 (decimal) or 0x0 to 0xFFFF (hexadecimal)

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 that can be configured using:

Table 44: 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 then the entry should be specified as: 003FA000000 0xFFFFFFFF000000

Values 0x0000000000000000 to 0xFFFFFFFFFFFFFFF (hexadecimal)

Default 0xFFFFFFFFFFFFFFF

3.8.2.1.9 Policy and entry maintenance commands

copy

Syntax

copy {**ip-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 used to create new filters using existing filters. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

If **overwrite** is not specified, an error will occur if the destination policy ID exists.

Parameters

ip-filter

Specifies that the *source-filter-id* and the *dest-filter-id* are IP 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 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 to which the copy command will attempt 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, everything in the existing destination filter ID will be completely overwritten with the contents of the source filter ID. If the destination filter ID exists, either **overwrite** must be specified or an error message will be returned. If **overwrite** is specified, the function of copying from source to destination occurs in a 'break before make' manner and therefore should be handled with care.

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 the **filter-name** attribute of a specific filter. When configured, **filter-name** 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 because 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

3.8.2.2 Show commands

```
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 as to 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.

- [Sample output](#) , [Table 45: Output fields: filter IP](#)
- [Sample output with IP filter ID specified](#), [Table 46: Output fields: filter IP with filter ID specified](#)
- [Sample output with time-range specified](#)
- [Sample output: associations](#), [Table 47: Output fields: filter IP associations](#)
- [Sample output for IP filter counters](#), [Table 48: Output fields: filter IP 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 45: 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 with IP filter ID specified

```

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#

```

Table 46: Output fields: filter IP 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.

Label	Description
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.
	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. Port	The source TCP or UDP port number.
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 — Specifies not to search for packets that contain the option field or have an option field of zero.

Label	Description
	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. as part of the match criteria.
Egr. Matches	The number of egress filter matches or hits for the filter entry.

Sample output with time-range specified

```
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                     Cur. Status   : Active
Src. IP       : 0.0.0.0/0                 Src. Port     : None
Dest. IP      : 10.10.1.1/16             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
=====
A:ALA-49#
```

Sample output: associations

```
A:ALA-49# show filter ip 1 associations
=====
IP Filter
=====
Filter Id       : 1                      Applied       : Yes
```

```

Scope      : Template                      Def. Action : Drop
Entries    : 1
-----
Filter Association : IP
-----
Service Id : 1001                          Type        : VPLS
- SAP      1/1/1:1001 (Ingress)
Service Id : 2000                          Type        :
- SAP      1/1/1:2000 (Ingress)
=====
A:ALA-49#

A:ALA-49# show filter ip 160 associations
=====
IP Filter
=====
Filter Id   : 160                          Applied      : No
Scope      : Template                      Def. Action  : Drop
Entries    : 0
-----
Filter Association : IP
-----
Tod-suite "english_suite"
- ingress, time-range "day" (priority 5)
=====
A:ALA-49#

```

Table 47: 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.
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.
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.

Label	Description
(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 IP filter counters

Table 48: Output fields: filter IP 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. The ingress counters count the packets with Layer 2 encapsulation.
Egr. Matches	The number of egress filter matches or hits for the filter entry. The 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 as to 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 output are examples of IPv6 filter information, and the associated tables describe the output fields.

- [Sample output, Table 49: Output fields: filter IPv6](#)
- [Sample output for IPv6 with a filter ID specified, Table 50: Output fields: filter IPv6 with filter ID specified](#)
- [Sample output for IPv6 filter associations, Table 51: Output fields: filter IPv6 associations](#)
- [Sample output for IPv6 filter counters, Table 52: Output fields: filter IPv6 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 49: Output fields: filter IPv6

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 with a 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       : 1::1/128                Src. Port     : None
Dest. IP      : ::/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                    Src. Port     : None
Dest. IP     : 1:2::1AFC/128           Dest. Port    : None
Next Header  : Undefined               Dscp          : Undefined
ICMP Type    : Undefined               ICMP Code     : Undefined
TCP-syn      : Off                     TCP-ack       : Off
Match action : Drop
Ing. Matches : 819 pkts
Egr. Matches : 0 pkts

=====
*A:7210SAS>show>filter#
```

Table 50: Output fields: filter IPv6 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.
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

Label	Description
	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 51: Output fields: filter IPv6 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 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.

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 52: Output fields: filter IPv6 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. 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. When no parameters are specified, a brief listing of IP filters is produced.

Parameters

mac-filter-id

Displays detailed information for the specified filter ID and its filter entries.

Values 1 to 65535

associations

Displays information as to 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. The associated tables describe the output fields.

- [Sample Detailed Output, Sample output for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, Table 53: Output fields: MAC filter](#)
- [Sample output for MAC filter counters, Table 54: Output fields: filter MAC counters](#)
- [Sample output for MAC filter associations, Table 55: Output fields: filter MAC associations](#)

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 : 200FrameType : 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: 0Egr. 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
=====
```

Sample output for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

```
=====
Mac Filter
=====
Filter Id      : 1                               Applied       : No
Scope         : Template                       Def. Action   : Drop
Entries       : 1                               Type          : unknown
Description    : (Not Specified)
-----
Filter Match Criteria : Mac
-----
Entry         : 1 (Inactive)
Description   : (Not Specified)
Src Mac       :
Dest Mac      :
Outer Dot1p* : none                           Outer Dot1p Mask: none
Inner Dot1p* : none                           Inner Dot1p Mask: none
Outer TagVal  : none                           Outer TagMask   : none
Inner TagVal  : none                           Inner TagMask   : none
Ethertype     : Undefined
Match action  : Drop
Ing. Matches  : 0 pkts
Egr. Matches  : 0 pkts
=====
```

Table 53: Output fields: MAC filter

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 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.

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 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.
Outer Dot1p	The IEEE 802.1p value for the match criteria used to match the Dot1p in the outermost VLAN tag. Undefined indicates no value is specified.
Inner Dot1p	The IEEE 802.1p value for the match criteria used to match the Dot1p in the inner VLAN tag. Undefined indicates no value is specified.
Outer TagVal	The VLAN ID value for the match criteria used to match the VLAN ID in the outermost VLAN tag. Undefined indicates no value is specified.
Inner TagVal	The IEEE 802.1p value for the match criteria used to match the Dot1p in the inner VLAN tag. 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 will be dropped.
	Forward — Packets matching the filter entry criteria is forwarded.

Label	Description
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 counters

```
A:ALA-49# show filter mac 8 counters
=====
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
```

Table 54: Output fields: filter MAC counters

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 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.
	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 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.

Label	Description
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: 3Applied: Yes
Scope: TemplateDef. Action: Drop
Entries: 1
```

```
-----
Filter Association : Mac
-----
```

```
Service Id: 1001Type: VPLS
- SAP 1/1/1:1001(Egress)
```

```
=====
A:ALA-49#
```

Table 55: 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.

3.8.2.3 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.

3.8.2.4 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

repeat repeat

Specifies how many times the command is repeated.

Values 1 to 999

Default 10

absolute

Displays the 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

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 the 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-filter-id

Specifies 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

interval

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 the 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.

4 Common CLI command descriptions

4.1 In this chapter

This section provides information about common Command Line Interface (CLI) syntax and command usage.

Topics in this chapter include:

- [SAP commands](#)

4.2 Common service commands

This section provides information about common service commands.

4.2.1 SAP commands

4.2.1.1 SAP command description

```
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 56: 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>]: <i>qtag1</i>	<i>port-id</i> : <i>qtag1</i> : 1/1/3:100 <i>lag-id</i> :lag-1:102
qinq	[<i>port-id</i> <i>lag-id</i>]: <i>qtag1.qtag2</i>	<i>port-id</i> : <i>qtag1.qtag2</i> : 1/1/ 3:100.10 <i>lag-id</i> : <i>qtag1.qtag2</i> :lag-10:

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 depends on the encapsulation type configured for the interface. The following table describes the allowed values for the port and encapsulation types.

Table 57: 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.
Ethernet	QinQ	qtag1: 0 to 4094 qtag2: 0 to 4094	The SAP is identified by two 802.1Q tags on the port. Note that a 0 qtag1 value also accepts untagged packets on the Dot1q port.

5 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) means 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/S-1/10GE-VC indicates 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE platforms in standalone-VC mode.
- 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.
- 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 is 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.

5.1 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

5.2 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

5.3 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.

5.4 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

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

RFC 7490, Remote Loop-Free Alternate (LFA) Fast Reroute (FRR) is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12



Note:
With Segment Routing.

5.5 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-vrrp-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

5.6 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 are 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.

5.7 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.

5.8 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

5.9 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.

5.10 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, 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, R6, and R12

5.11 Management

draft-ietf-snmpv3-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

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 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

5.12 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

5.13 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

5.14 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.

5.15 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

5.16 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

5.17 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

5.18 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, 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, R6, and R12

5.19 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

5.20 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

5.21 RIP

RFC 1058, Routing Information Protocol is supported on Mxp

RFC 2082, RIP-2 MD5 Authentication is supported on Mxp

RFC 2453, RIP Version 2 is supported on Mxp

5.22 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:

Only on 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:

Only on 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:

Only on Dxp-12p ETR. Dxp-ETR and 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:

Only on 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:

Only on 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:

Only on 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:

Only on 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:

Only on 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:

Only on 7210 SAS-Sx 10/100GE QSFP28 variant and Dxp-12p ETR.

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

5.23 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](#)