



7210 Service Access System

Release 25.9.R1

7210 SAS-K 2F6C4T, K 3SFP+ 8C Routing Protocols Guide

3HE 21181 AAAB TQZZA 01
Edition: 01
September 2025

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

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.

© 2025 Nokia.

Table of contents

List of tables.....	13
List of figures.....	17
1 Getting started.....	19
1.1 About this guide.....	19
1.1.1 Document structure and content.....	19
1.2 7210 SAS modes of operation.....	20
1.3 7210 SAS port modes.....	22
1.4 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C routing configuration process.....	24
1.5 Conventions.....	25
1.5.1 Precautionary and information messages.....	25
1.5.2 Options or substeps in procedures and sequential workflows.....	25
2 Multicast.....	27
2.1 Overview of multicast.....	27
2.1.1 Multicast models (SSM).....	27
2.1.1.1 SSM.....	28
2.2 Multicast features.....	28
2.2.1 IGMP.....	28
2.2.1.1 IGMP versions and interoperability requirements.....	29
2.2.1.2 IGMP version transition.....	29
2.2.1.3 SSM groups.....	29
2.2.2 PIM-SM.....	30
2.2.2.1 PIM-SM functions.....	30
2.2.2.2 Encapsulating data packets in the register tunnel.....	32
2.2.2.3 PIM bootstrap router mechanism.....	32
2.2.2.4 PIM-SM routing policies.....	33
2.2.2.5 Reverse path forwarding checks.....	34
2.2.2.6 Distributing PIM joins over multiple ECMP paths.....	36
2.2.3 Multicast debugging tools.....	39
2.2.3.1 Mtrace.....	39
2.2.3.2 Mrinfo.....	41
2.2.4 Configuration guidelines for 7210 SAS.....	41

2.3	Configuring multicast parameters with CLI.....	41
2.3.1	Multicast configuration overview.....	41
2.3.2	Basic configuration.....	42
2.3.3	Common configuration tasks.....	44
2.3.4	Configuring IGMP parameters.....	44
2.3.4.1	Enabling IGMP.....	45
2.3.4.2	Configuring an IGMP interface.....	45
2.3.4.3	Configuring static parameters.....	46
2.3.4.4	Configuring SSM translation.....	47
2.3.5	Configuring PIM parameters.....	48
2.3.5.1	Enabling PIM.....	48
2.3.5.2	Configuring PIM interface parameters.....	49
2.3.5.3	Importing PIM join or register policies.....	53
2.3.6	Disabling IGMP or PIM.....	54
2.4	Multicast command reference.....	55
2.4.1	Command hierarchies.....	55
2.4.1.1	Configuration commands.....	56
2.4.1.2	IGMP commands.....	56
2.4.1.3	PIM commands.....	57
2.4.1.4	Operational commands.....	58
2.4.1.5	Show commands.....	58
2.4.1.6	Clear commands.....	58
2.4.1.7	Debug commands.....	59
2.4.2	Command descriptions.....	60
2.4.2.1	Configuration commands.....	60
2.4.2.2	Show commands.....	105
2.4.2.3	Show router PIM commands.....	116
2.4.2.4	Clear commands.....	134
2.4.2.5	Debug commands.....	141
3	OSPF.....	154
3.1	Configuring OSPF.....	154
3.1.1	OSPF areas.....	155
3.1.1.1	Backbone area.....	155
3.1.1.2	Stub area.....	156
3.1.1.3	Not-So-Stubby Area.....	156

3.1.2	OSPFv3 authentication.....	160
3.1.3	OSPFv3 graceful restart helper.....	161
3.1.4	Virtual links.....	162
3.1.5	Neighbors and adjacencies.....	162
3.1.6	Link-state advertisements.....	162
3.1.7	Metrics.....	163
3.1.8	Authentication.....	163
3.1.9	Multiple OSPF instances.....	163
3.1.9.1	Route export policies for OSPF.....	163
3.1.9.2	Preventing route redistribution loops.....	164
3.1.10	IP subnets.....	164
3.1.11	Preconfiguration recommendations.....	164
3.2	IP Fast-Reroute (IP FRR) for OSPF and IS-IS prefixes.....	165
3.2.1	IP FRR/LFA configuration.....	165
3.2.1.1	Reducing the scope of the LFA calculation by SPF.....	166
3.2.2	ECMP considerations.....	166
3.2.3	IP FRR and RSVP shortcut (IGP shortcut).....	166
3.2.4	IP FRR and BGP next-hop resolution.....	166
3.2.5	OSPF and IS-IS support for Loop-Free Alternate calculation.....	167
3.2.5.1	Loop-Free Alternate calculation in the presence of IGP shortcuts.....	169
3.2.5.2	Loop-Free Alternate calculation for inter-area/inter-level prefixes.....	170
3.3	Loop-Free Alternate Shortest Path First (LFA SPF) policies.....	170
3.3.1	Configuration of route next-hop policy template.....	171
3.3.1.1	Configuring affinity or admin group constraint in route next-hop policy.....	171
3.3.1.2	Configuring SRLG group constraint in route next-hop policy.....	172
3.3.1.3	Interaction of IP and MPLS admin group and SRLG.....	173
3.3.1.4	Configuring protection type and next-hop type preference in route next-hop policy template.....	173
3.3.2	Application of route next-hop policy template to an interface.....	174
3.3.3	Excluding prefixes from LFA SPF.....	174
3.3.4	Modification to LFA next-hop selection algorithm.....	175
3.4	Segment routing in shortest path forwarding.....	176
3.4.1	LFA protection using segment routing backup node SID.....	176
3.4.1.1	Detailed operation of LFA protection using backup node SID.....	177
3.4.1.2	Duplicate SID handling.....	179
3.4.1.3	OSPF control plane extensions.....	180

3.5	OSPF configuration process overview.....	182
3.6	Configuration notes.....	182
3.6.1	General.....	182
3.6.1.1	OSPF defaults.....	183
3.7	Configuring OSPF with CLI.....	183
3.8	OSPF configuration guidelines.....	183
3.9	Basic OSPF configuration.....	183
3.9.1	Configuring the router ID.....	184
3.10	Configuring OSPF components.....	185
3.10.1	Configuring OSPF parameters.....	185
3.10.2	Configuring OSPFv3 parameters.....	186
3.10.3	Configuring an OSPF and OSPFv3 area.....	186
3.10.4	Configuring an OSPF and OSPFv3 stub area.....	187
3.10.5	Configuring a Not-So-Stubby Area.....	188
3.10.6	Configuring a virtual link.....	189
3.10.7	Configuring an interface.....	190
3.10.8	Configuring authentication.....	191
3.10.9	Assigning a designated router.....	194
3.10.10	Configuring route summaries.....	195
3.10.11	Configuring route preferences.....	197
3.11	OSPF configuration management tasks.....	198
3.11.1	Modifying a router ID.....	198
3.11.2	Deleting a router ID.....	199
3.11.3	Modifying OSPF parameters.....	199
3.12	OSPF command reference.....	200
3.12.1	Command hierarchies.....	200
3.12.1.1	Configuration commands for OSPF.....	201
3.12.1.2	Show commands.....	204
3.12.1.3	Clear commands.....	205
3.12.1.4	Debug commands.....	205
3.12.2	Command descriptions.....	206
3.12.2.1	Configuration commands.....	206
3.12.2.2	Show commands.....	257
3.12.2.3	Clear commands.....	299
3.12.2.4	Debug commands.....	301

4	IS-IS.....	310
4.1	Configuring IS-IS.....	310
4.1.1	Routing.....	311
4.1.2	IS-IS frequently used terms.....	312
4.1.3	ISO network addressing.....	313
4.1.4	IS-IS PDU configuration.....	314
4.1.5	IS-IS operations.....	315
4.1.6	IS-IS route summarization.....	315
4.1.7	IS-IS multi-topology for IPv6.....	315
4.1.8	IS-IS administrative tags.....	316
4.1.8.1	Setting route tags.....	316
4.1.8.2	Using route tags.....	316
4.1.9	Segment routing in shortest path forwarding.....	316
4.1.9.1	Segment routing operational procedures.....	317
4.1.9.2	Segment routing tunnel management.....	321
4.1.9.3	Remote LFA with segment routing.....	323
4.1.9.4	Data path support.....	327
4.1.9.5	Control protocol changes.....	328
4.1.9.6	BGP label route resolution using segment routing tunnels.....	332
4.1.9.7	Service packet forwarding with segment routing.....	333
4.1.9.8	Mirror services.....	333
4.1.10	IGP-LDP synchronization.....	334
4.2	IS-IS configuration process overview.....	334
4.3	Configuration notes.....	335
4.3.1	General.....	335
4.4	Configuring IS-IS with CLI.....	335
4.5	IS-IS configuration overview.....	336
4.5.1	Router levels.....	336
4.5.2	Area address attributes.....	336
4.5.3	Interface level capability.....	337
4.5.4	Route leaking.....	337
4.6	Basic IS-IS configuration.....	338
4.7	Common configuration tasks.....	340
4.8	Configuring IS-IS components.....	340
4.8.1	Enabling IS-IS.....	340

4.8.2	Modifying router-level parameters.....	341
4.8.3	Configuring ISO area addresses.....	342
4.8.4	Configuring global IS-IS parameters.....	342
4.8.5	Migration to IS-IS multi-topology.....	343
4.8.6	Configuring interface parameters.....	346
4.8.6.1	Example: configuring a Level 1 area.....	348
4.8.6.2	Example: modifying a router's level capability.....	349
4.9	IS-IS configuration management tasks.....	350
4.9.1	Disabling IS-IS.....	350
4.9.2	Removing IS-IS.....	350
4.9.3	Modifying global IS-IS parameters.....	351
4.9.4	Modifying IS-IS interface parameters.....	352
4.9.5	Configuring leaking.....	353
4.9.6	Redistributing external IS-IS routers.....	355
4.10	IS-IS command reference.....	356
4.10.1	Command hierarchies.....	356
4.10.1.1	Configuration commands.....	356
4.10.1.2	Global commands.....	356
4.10.1.3	Interface command.....	357
4.10.1.4	Show commands.....	358
4.10.1.5	Clear commands.....	359
4.10.1.6	Debug commands.....	359
4.10.2	Command descriptions.....	360
4.10.2.1	IS-IS configuration commands.....	360
4.10.2.2	Show commands.....	413
4.10.2.3	Clear commands.....	443
4.10.2.4	Debug commands.....	445
5	BGP.....	453
5.1	BGP overview.....	453
5.2	BGP communication.....	453
5.2.1	Message types.....	454
5.3	Group configuration and peers.....	455
5.4	Hierarchical levels.....	456
5.5	Route reflection.....	456
5.6	Fast external failover.....	458

5.7	Sending of BGP communities.....	459
5.8	ECMP and BGP route tunnels.....	459
5.9	Next-hop resolution of BGP labeled routes to tunnels.....	459
5.9.1	VPN-IPv4 and VPN-IPv6 route resolution.....	460
5.10	Route selection criteria.....	461
5.11	BGP path attributes.....	462
5.11.1	NEXT_HOP attribute.....	462
5.11.1.1	Next-hop indirection.....	462
5.12	BGP Routing Information Base.....	463
5.12.1	LOC-RIB features.....	463
5.12.2	BGP fast reroute.....	464
5.12.2.1	Calculating backup paths.....	464
5.12.2.2	Failure detection and switchover to the backup path.....	464
5.12.3	RIB-OUT features.....	465
5.12.3.1	BGP export policies.....	465
5.12.3.2	Outbound Route Filtering.....	466
5.12.3.3	RT constrained route distribution.....	467
5.12.3.4	Minimum Route Advertisement Interval.....	468
5.12.3.5	Advertise-inactive.....	469
5.12.3.6	Split-horizon.....	469
5.13	Add-path.....	470
5.13.1	Receiving multiple paths per prefix from a BGP peer.....	470
5.13.2	Path selection with add-paths.....	471
5.13.3	BGP decision process with add-path.....	471
5.13.4	Advertising multiple paths using add-path.....	472
5.13.5	Limiting the number of paths per prefix.....	473
5.14	AIGP metric.....	473
5.15	Command interactions and dependencies.....	474
5.15.1	Changing the ASN.....	474
5.15.2	Changing the local ASN.....	474
5.15.3	Changing the router ID at the configuration level.....	474
5.15.4	Hold time and keep alive timer dependencies.....	474
5.15.5	Import and export route policies.....	475
5.15.6	Route damping and route policies.....	475
5.15.7	AS Override.....	475
5.16	Configuration guideline for BGP.....	476

5.17	BGP configuration process overview.....	476
5.18	Configuration notes.....	476
5.18.1	General.....	477
5.18.1.1	BGP defaults.....	477
5.18.1.2	BGP MIB notes.....	477
5.19	Configuring BGP with CLI.....	479
5.20	BGP configuration overview.....	479
5.20.1	Preconfiguration requirements.....	479
5.20.2	BGP hierarchy.....	479
5.20.3	Internal and external BGP configurations.....	479
5.21	Basic BGP configuration.....	480
5.22	Common configuration tasks.....	481
5.22.1	Creating an autonomous system.....	482
5.22.2	Configuring a router ID.....	483
5.22.3	BGP components.....	484
5.22.4	Configuring BGP.....	484
5.22.5	Configuring group attributes.....	484
5.22.6	Configuring neighbor attributes.....	485
5.22.7	Configuring AIGP.....	486
5.23	BGP configuration management tasks.....	487
5.23.1	Modifying an ASN.....	487
5.23.2	Modifying the BGP router ID.....	488
5.23.3	Modifying the router-level router ID.....	489
5.23.4	Deleting a neighbor.....	489
5.23.5	Deleting groups.....	490
5.23.6	Editing BGP parameters.....	491
5.24	BGP command reference.....	491
5.24.1	Command hierarchies.....	491
5.24.1.1	Configuration commands.....	491
5.24.1.2	Show commands.....	495
5.24.1.3	Clear commands.....	496
5.24.1.4	Debug commands.....	496
5.24.2	Command descriptions.....	497
5.24.2.1	Configuration commands.....	497
5.24.2.2	Other BGP-related commands.....	545
5.24.2.3	Show commands.....	546

5.24.2.4	Clear commands.....	600
5.24.2.5	Debug commands.....	603
6	Route policies.....	610
6.1	Configuring route policies.....	610
6.1.1	Policy statements.....	610
6.1.1.1	Default action behavior.....	611
6.1.1.2	Denied IP prefixes.....	611
6.1.1.3	Controlling route flapping.....	611
6.2	Regular expressions.....	612
6.2.1	BGP and OSPF route policy support.....	616
6.2.1.1	BGP route policies.....	617
6.2.1.2	Re-advertised route policies.....	618
6.2.2	When to use route policies.....	618
6.3	Route policy configuration process overview.....	618
6.4	Configuration notes.....	619
6.4.1	General.....	619
6.5	Configuring route policies with CLI.....	619
6.6	Route policy configuration overview.....	619
6.6.1	When to create routing policies.....	620
6.6.2	Default route policy actions.....	620
6.6.3	Policy evaluation.....	621
6.6.4	Damping.....	622
6.7	Basic configurations.....	623
6.8	Configuring route policy components.....	624
6.8.1	Beginning the policy statement.....	624
6.8.2	Creating a route policy.....	625
6.8.3	Configuring a default action.....	626
6.8.4	Configuring an entry.....	626
6.8.5	Configuring damping.....	627
6.8.5.1	Configuring a prefix list.....	627
6.9	Route policy configuration management tasks.....	628
6.9.1	Editing policy statements and parameters.....	628
6.9.2	Deleting an entry.....	629
6.9.3	Deleting a policy statement.....	629
6.9.4	Use of route policies for IGMP filtering.....	630

6.10	Route policy command reference.....	631
6.10.1	Command hierarchies.....	631
6.10.1.1	Route policy configuration commands.....	631
6.10.1.2	Show commands.....	633
6.10.2	Command descriptions.....	633
6.10.2.1	Route policy command reference.....	633
6.10.2.2	Show commands.....	667
7	Standards and protocol support.....	673
7.1	BGP.....	673
7.2	Ethernet.....	675
7.3	EVPN.....	676
7.4	Fast Reroute.....	676
7.5	Internet Protocol (IP) — General.....	677
7.6	IP — Multicast.....	679
7.7	IP — Version 4.....	680
7.8	IP — Version 6.....	681
7.9	IPsec.....	682
7.10	IS-IS.....	683
7.11	Management.....	684
7.12	MPLS — General.....	687
7.13	MPLS — GMPLS.....	688
7.14	MPLS — LDP.....	688
7.15	MPLS — MPLS-TP.....	688
7.16	MPLS — OAM.....	689
7.17	MPLS — RSVP-TE.....	689
7.18	OSPF.....	690
7.19	Pseudowire.....	691
7.20	Quality of Service.....	692
7.21	RIP.....	692
7.22	Timing.....	692
7.23	VPLS.....	694

List of tables

Table 1: Supported modes of operation and configuration methods.....	21
Table 2: Supported port modes by mode of operation.....	23
Table 3: 7210 SAS platforms supporting port modes.....	23
Table 4: Configuration process.....	24
Table 5: Join filter policy match conditions.....	33
Table 6: Register filter policy match conditions.....	34
Table 7: Output fields: mrinfo.....	102
Table 8: Output fields: mtrace.....	104
Table 9: Output fields: IGMP group.....	107
Table 10: Output fields: IGMP interface.....	108
Table 11: Output fields: IGMP SSM translate.....	111
Table 12: Output fields: IGMP static.....	112
Table 13: Output fields: IGMP statistics.....	114
Table 14: Output fields: IGMP status.....	115
Table 15: Output fields: PIM anycast.....	117
Table 16: Output fields: PIM CRP.....	118
Table 17: Output fields: PIM group.....	120
Table 18: Output fields: PIM interface.....	123
Table 19: Output fields: PIM neighbor.....	125
Table 20: Output fields: PIM RP.....	127
Table 21: Output fields: PIM RP hash.....	128

Table 22: Output fields: PIM statistics.....	129
Table 23: Output fields: PIM status.....	133
Table 24: Handling of duplicate SIDs.....	179
Table 25: OSPF control plane extension fields.....	180
Table 26: OSPF control plane extension flags.....	181
Table 27: Route preference defaults by route type.....	197
Table 28: Route preference defaults by route type.....	215
Table 29: Output fields: OSPF area.....	259
Table 30: Output fields: OSPF database.....	263
Table 31: Output fields: OSPF interface.....	266
Table 32: Output fields: OSPF interface detail.....	268
Table 33: Output fields: OSPF neighbor.....	273
Table 34: Output fields: OSPF neighbor detail.....	275
Table 35: Output fields: prefix SIDs.....	279
Table 36: Output fields: OSPF range.....	281
Table 37: Output fields: OSPF SFP.....	285
Table 38: Output fields: OSPF statistics.....	287
Table 39: Output fields: OSPF status.....	291
Table 40: Output fields: OSPF virtual link.....	293
Table 41: Output fields: OSPF virtual neighbor.....	297
Table 42: Data path support.....	327
Table 43: Potential adjacency capabilities.....	337
Table 44: Default route preferences.....	373

Table 45: Potential adjacency capabilities.....	386
Table 46: Default preferences.....	401
Table 47: Output fields: ISIS adjacency.....	417
Table 48: Output fields: IS-IS capabilities.....	420
Table 49: Output fields: router IS-IS database.....	422
Table 50: Output fields: router IS-IS hostname.....	425
Table 51: Output fields: IS-IS interface.....	427
Table 52: Output fields: LFA coverage.....	428
Table 53: Output fields: link group member status.....	429
Table 54: Output fields: prefix SIDs.....	432
Table 55: Output fields: IS-IS routes.....	435
Table 56: Output fields: IS-IS SPF log.....	436
Table 57: Output fields: IS-IS statistics.....	437
Table 58: Output fields: IS-IS status.....	439
Table 59: Output fields: IS-IS summary address.....	441
Table 60: Output fields: IS-IS topology.....	442
Table 61: BGP fast reroute scenarios (base router context).....	464
Table 62: 7210 SAS and IETF MIB variations.....	477
Table 63: MIB variable with SNMP.....	478
Table 64: Output fields: BGP auth-keychain.....	549
Table 65: Output fields: router BGP damping.....	554
Table 66: Output fields: router BGP group.....	558
Table 67: Output fields: router BGP neighbor.....	568

Table 68: Output fields: router BGP neighbor received-routes.....

572

Table 69: Output fields: show neighbor add-path.....

574

Table 70: Output fields: router BGP next-hop.....

580

Table 71: Output fields: router BGP paths.....

581

Table 72: Output fields: router BGP routes.....

590

Table 73: Output fields: BGP routes IPv4.....

592

Table 74: Output fields: BGP EVPN routes.....

595

Table 75: Output fields: router BGP summary.....

598

Table 76: Regular expression operators.....

613

Table 77: AS path and community regular expression examples.....

614

Table 78: Default route policy actions.....

620

Table 79: Output fields: router policy.....

672

List of figures

Figure 1: Anycast RP for PIM-SM implementation example.....	35
Figure 2: Backbone area.....	156
Figure 3: PEs connected to an MPLS-VPN super backbone.....	158
Figure 4: Sham links.....	159
Figure 5: Topology example with primary and LFA routes.....	167
Figure 6: Topology example with broadcast interfaces.....	168
Figure 7: Label stack for remote LFA in ring topology.....	177
Figure 8: Backup ABR node SID.....	178
Figure 9: OSPF configuration and implementation flow.....	182
Figure 10: Checking corresponding bit.....	210
Figure 11: IS-IS routing domain.....	311
Figure 12: Using area addresses to form adjacencies.....	314
Figure 13: Handling of the same prefix and SID in different IS-IS instances.....	320
Figure 14: Example topology remote LFA algorithm.....	324
Figure 15: Remote LFA next-hop in segment routing.....	326
Figure 16: IS-IS configuration and implementation flow.....	335
Figure 17: Configuring a Level 1 area.....	348
Figure 18: Configuring a Level 1/2 area.....	350
Figure 19: BGP configuration.....	455
Figure 20: Fully meshed BGP configuration.....	457
Figure 21: BGP configuration with route reflectors.....	458

Figure 22: BGP Update message with path identifier for IPv4 NLRI.....	470
Figure 23: BGP configuration and implementation flow.....	476
Figure 24: BGP route policy diagram.....	617
Figure 25: OSPF route policy diagram.....	617
Figure 26: Route policy configuration and implementation flow.....	619
Figure 27: Route policy process example.....	622
Figure 28: Damping example.....	623

1 Getting started

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

1.1 About this guide



Note:

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

This guide describes system concepts and provides configuration examples to configure the boot option file (BOF) on the following 7210 SAS platforms, operating in one of the modes described in [Table 1: Supported modes of operation and configuration methods](#). If multiple modes of operation apply, they are explicitly noted in the topic.

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

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



Note:

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

- access-uplink mode of operation
7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C
- network mode of operation
7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 25.x.Rx content and may include some content that will be released in later maintenance loads. See the *7210 SAS Software Release Notes 25.x.Rx*, part number 3HE 21188 000x TQZZA, for information about features supported in each load of the Release 25.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.
- 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 25.x.Rx*, part number 3HE 21188 000x TQZZA, and to the appropriate 7210 SAS software user guide for information about features and capabilities supported by a 7210 SAS platform when operating in a specific mode.

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

- **access-uplink**

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

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

- **network**

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

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

- **satellite**

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

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

- **standalone**

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

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

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

- **standalone-VC**

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

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

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

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

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

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

Table 1: Supported modes of operation and configuration methods

7210 SAS platform	Mode of operation and configuration method				
	Network	Access-uplink	Standalone	Standalone-VC	Satellite
7210 SAS-D		Implicit	Implicit		
7210 SAS-Dxp		Implicit	Implicit		
7210 SAS-K 2F1C2T		Implicit	Implicit		
7210 SAS-K 2F6C4T ¹	Port Mode Configuration ²	Port Mode Configuration ²	Implicit		
7210 SAS-K 3SFP+ 8C ¹	Port Mode Configuration ²	Port Mode Configuration ²	Implicit		
7210 SAS-Mxp	Implicit ³		Explicit BOF Configuration		Explicit BOF Configuration
7210 SAS-R6 ⁴	Implicit		Implicit		

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

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

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

⁴ Supports MPLS uplinks only and implicitly operates in network mode



Note:

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

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

Table 2: Supported port modes by mode of operation

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

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

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

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

Platform	Port mode			
	Access	Network	Hybrid	Access-uplink
7210 SAS-R6 IMM-c 100GE (IMM-c 1QSFP28)	Yes	Yes	Yes	No
7210 SAS-R12 IMM-b	Yes	Yes	Yes	No
7210 SAS-R12 IMM-c 100GE (IMM-c 1QSFP28)	Yes	Yes	Yes	No
7210 SAS-Sx/S 1/10GE	Yes	Yes	Yes	No
7210 SAS-Sx 10/100GE	Yes	Yes	Yes	No
7210 SAS-T	Yes	Yes ⁶	Yes ⁷	Yes ⁸

1.4 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C routing configuration process

The following table lists the tasks necessary to configure OSPF, IS-IS, BGP, and route policies. 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.

Table 4: Configuration process

Area	Task	Chapter
Protocol configuration	Configure routing protocols:	
	• Multicast	Multicast
	• OSPF	OSPF
	• IS-IS	IS-IS
	• BGP	BGP
Policy configuration	• Configure route policies	Route policies
Reference	List of IEEE, IETF, and other proprietary entities	Standards and protocol support

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.

1.5 Conventions

This section describes the general conventions used in this guide.

1.5.1 Precautionary and information messages

The following information symbols are used in the documentation.



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



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



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



Note: Note provides additional operational information.



Tip: Tip provides suggestions for use or best practices.

1.5.2 Options or substeps in procedures and sequential workflows

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

Example: Options in a procedure

1. User must perform this step.
2. This step offers three options. User must perform one option to complete this step:
 - This is one option.
 - This is another option.
 - This is yet another option.

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

Example: Substeps in a procedure

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

- b.** This is another substep.

2 Multicast

This chapter provides information about Protocol Independent Multicast (PIM).

2.1 Overview of multicast

IP multicast is a method of many-to-many communication that simplifies the delivery of unicast datagrams. In the case of unicast delivery, IP packets are sent from a single source to a single recipient. The source inserts the address of the target host in the IP header destination field of an IP datagram, and intermediate routers (if present) forward the datagram toward the target in accordance with their respective routing tables.

However, some applications, such as audio or video streaming broadcasts, require the delivery of individual IP packets to multiple destinations. In such applications, multicast is used to distribute datagrams sourced from one or more hosts to a set of receivers that may be distributed over different (sub) networks. The delivery of multicast datagrams is significantly more complex.

Multicast sources can send a single copy of data using a single address for the entire group of recipients. The routers between the source and recipients route the data using the group address route. Multicast packets are delivered to a multicast group. A multicast group specifies a set of recipients who are interested in a particular data stream and is represented by an IP address from a specified range. Data addressed to the IP address is forwarded to the members of the group. A source host sends data to a multicast group by specifying the multicast group address in the datagram destination IP address. A source does not have to register to send data to a group, nor does it need to be a member of the group.

Routers and Layer 3 (L3) switches use the Internet Group Management Protocol (IGMP) to manage membership for a multicast session. When a host needs to receive one or more multicast sessions, it signals its local router by sending a join message to each multicast group it needs to join. When a host needs to leave a multicast group, it sends a leave message.

To extend multicast to the Internet, the multicast backbone (Mbone) is used. The Mbone is layered on top of portions of the Internet. These portions, or islands, are interconnected using tunnels. The tunnels allow multicast traffic to pass between the multicast-capable portions of the Internet. As more and more routers in the Internet are multicast-capable (and scalable), the unicast and multicast routing table will converge.

The original Mbone was based on the Distance Vector Multicast Routing Protocol (DVMRP) and was very limited. The Mbone is, however, converging around the following protocol set:

- IGMP
- Protocol Independent Multicast (Sparse Mode) (PIM-SM)

2.1.1 Multicast models (SSM)

This section provides information about the Source-Specific Multicast (SSM) model.

2.1.1.1 SSM

The SSM service model defines a channel identified by an (S,G) pair, where S is a source address and G is an SSM destination address. In contrast to the ASM model, SSM only provides network-layer support for one-to-many delivery.

The SSM service model attempts to alleviate the following deployment problems:

- **address allocation**

SSM defines channels on a per-source basis. For example, the channel (S1,G) is distinct from the channel (S2,G), where S1 and S2 are source addresses, and G is an SSM destination address.

This averts the problem of global allocation of SSM destination addresses and makes each source independently responsible for resolving address collisions for the various channels it creates.

- **access control**

SSM provides an efficient solution to the access control problem. When a receiver subscribes to an (S,G) channel, it receives data sent only by the source S. In contrast, any host can transmit to an ASM host group. At the same time, when a sender picks a channel (S,G) to transmit on, it is automatically ensured that no other sender will be transmitting on the same channel (except in the case of malicious acts such as address spoofing). This makes it harder to spam an SSM channel than an ASM multicast group.

- **handling of well-known sources**

SSM requires only source-based forwarding trees. This eliminates the need for a shared tree infrastructure. In terms of the IGMP and PIM-SM, this implies that neither the RP-based shared tree infrastructure of PIM-SM nor the MSDP protocol is required. Therefore, the complexity of the multicast routing infrastructure for SSM is low, making it viable for immediate deployment.

- **handling point-to-point applications**

Anticipating that point-to-multipoint applications such as Internet TV will be significant in the future; the SSM model is better suited for such applications.

2.2 Multicast features

This section contains information about the multicast protocols required to support a Nokia router in the network.

2.2.1 IGMP

IGMP is used by IPv4 hosts and routers to report their IP multicast group memberships to neighboring multicast routers. A multicast router keeps a list of multicast group memberships for each attached network, and a timer for each membership.

Multicast group memberships include at least one member of a multicast group on an attached network. In each of its attached networks, a multicast router can assume one of two roles: querier or non-querier. There is typically only one querier per physical network.

The querier issues two types of queries: a general query and a group-specific query. General queries are issued to solicit membership information with regard to any multicast group. Group-specific queries are

issued when a router receives a leave message from the node it perceives as being the last remaining group member on that network segment.

If the host needs to receive a multicast session issue and a multicast group membership report, the reports must be sent to all multicast-enabled routers.

2.2.1.1 IGMP versions and interoperability requirements

If routers run different versions of IGMP, they negotiate the lowest common version of IGMP that is supported on their subnet and operate in that version. Three versions of IGMP are supported:

- **version 1**

Specified in RFC 1112, *Host extensions for IP Multicasting* was the first widely deployed version and the first version to become an Internet standard.

- **version 2**

Specified in RFC 2236, *Internet Group Management Protocol* added support for “low leave latency”, that is, a reduction in the time it takes for a multicast router to learn that there are no longer any members of a particular group present on an attached network.

- **version 3**

Specified in RFC 3376, *Internet Group Management Protocol* added support for source filtering, that is, the ability for a system to report interest in receiving packets only from specific source addresses, as required to support [SSM](#), or from all but specific source addresses, sent to a particular multicast address.

IGMPv3 must keep track of the state of each group for each attached network. The group state consists of a filter-mode, a list of sources, and various timers. For each attached network running IGMP, a multicast router records the desired reception state for that network.

2.2.1.2 IGMP version transition

Nokia routers are capable of interoperating with routers and hosts running IGMPv1, IGMPv2, and/or IGMPv3. RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3)/Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction* explores the interoperability issues and how they affect the routing protocols.

IGMPv3 specifies that if a router receives an earlier version query message on an interface, it must immediately switch to a mode that is compatible with the earlier version. Because the previous versions of IGMP are not source-aware, should this occur and the interface switches to version 1 or 2 compatibility mode, any previously learned group memberships with specific sources (learned via the IGMPv3-specific INCLUDE or EXCLUDE mechanisms) must be converted to non-source specific group memberships. The routing protocol will then treat the query as if there is no EXCLUDE definition present.

2.2.1.3 SSM groups

IGMPv3 permits a receiver to join a group and specify that it only needs to receive traffic for a group if that traffic comes from a particular source. If a receiver does this, and no other receiver on the LAN requires all the traffic for the group, the designated router (DR) can omit performing a (*,G) join to set up the shared tree, and instead issue a source-specific (S,G) join only.

The range of multicast addresses from 232.0.0.0 to 232.255.255.255 is currently set aside for source-specific multicast in IPv4. For groups in this range, receivers should only issue source-specific IGMPv3 joins. If a PIM router receives a non-source-specific join for a group in this range, it should ignore it.

A Nokia PIM router must silently ignore a received (*,G) PIM join message where G is a multicast group address from the multicast address group range that has been explicitly configured for SSM. This occurrence should generate an event. If configured, the IGMPv2 request can be translated into IGMPv3. The router allows for the conversion of an IGMPv2 (*,G) request into a IGMPv3 (S,G) request based on manual entries. A maximum of 32 SSM ranges is supported.

IGMPv3 also permits a receiver to join a group and specify that it only needs to receive traffic for a group if that traffic does not come from a specific source or sources. In this case, the DR will perform a (*,G) join as normal, but can combine this with a prune for each of the sources the receiver does not wish to receive.

2.2.2 PIM-SM

PIM-SM leverages the unicast routing protocols that are used to create the unicast routing table, OSPF, IS-IS, BGP, and static routes. Because PIM uses this unicast routing information to perform the multicast forwarding function, it is effectively IP protocol independent. Unlike DVMRP, PIM does not send multicast routing table updates to its neighbors.

PIM-SM uses the unicast routing table to perform the Reverse Path Forwarding (RPF) check function instead of building up a completely independent multicast routing table.

PIM-SM only forwards data to network segments with active receivers that have explicitly requested the multicast group. PIM-SM in the ASM model initially uses a shared tree to distribute information about active sources. Depending on the configuration options, the traffic can remain on the shared tree or switch over to an optimized source distribution tree. As multicast traffic starts to flow down the shared tree, routers along the path determine whether there is a better path to the source. If a more direct path exists, the router closest to the receiver sends a join message toward the source and reroutes the traffic along this path.

PIM-SM relies on an underlying topology-gathering protocol to populate a routing table with routes. The routing table is called the Multicast Routing Information Base (MRIB). The routes in this table can be taken directly from the unicast routing table, or they can be different and provided by a separate routing protocol such as MBGP. Regardless of how it is created, the primary role of the MRIB in the PIM-SM protocol is to provide the next hop router along a multicast-capable path to each destination subnet. The MRIB is used to determine the next hop neighbor to whom any PIM join/prune message is sent. Data flows along the reverse path of the join messages. In contrast to the unicast RIB that specifies the next hop that a data packet would take to get to a subnet, the MRIB gives reverse-path information and indicates the path that a multicast data packet would take from its origin subnet to the router that has the MRIB.

2.2.2.1 PIM-SM functions

This section provides information about the three phases of PIM-SM functions.

2.2.2.1.1 Phase one

In this phase, a multicast receiver expresses its interest in receiving traffic destined for a multicast group. Typically, it does this using IGMP or MLD, but other mechanisms can also serve this purpose. One of the receiver's local routers is elected as the DR for that subnet. When the expression of interest is received, the DR sends a PIM join message toward the RP for that multicast group. This join message is known

as a (*,G) join because it joins group G for all sources to that group. The (*,G) join travels hop-by-hop toward the RP for the group, and in each router it passes through the multicast tree state for group G is instantiated. Eventually, the (*,G) join either reaches the RP or reaches a router that already has the (*,G) join state for that group.

When many receivers join the group, their join messages converge on the RP and form a distribution tree for group G that is rooted at the RP. This is known as the RP tree and is also known as the shared tree because it is shared by all sources sending to that group. Join messages are resent periodically as long as the receiver remains in the group. When all receivers on a leaf-network leave the group, the DR sends a PIM (*,G) prune message toward the RP for that multicast group. However, if the prune message is not sent for any reason, the state will eventually time out.

A multicast data sender starts sending data destined for a multicast group. The sender's local router (the DR) takes those data packets, unicast-encapsulates them, and sends them directly to the RP. The RP receives these encapsulated data packets, de-encapsulates them, and forwards them to the shared tree. The packets then follow the (*,G) multicast tree state in the routers on the RP tree, being replicated wherever the RP tree branches, and eventually reaching all the receivers for that multicast group. The process of encapsulating data packets to the RP is called registering, and the encapsulation packets are known as PIM register packets.

At the end of phase one, multicast traffic flows encapsulated to the RP, and then natively over the RP tree to the multicast receivers.

2.2.2.1.2 Phase two

In this phase, register-encapsulation of data packets is performed. However, register-encapsulation of data packets is unsuitable for the following reasons:

- Encapsulation and de-encapsulation can be resource intensive operations for a router to perform depending on whether or not the router has appropriate hardware for the tasks.
- Traveling to the RP and then back down the shared tree can cause the packets to travel a relatively long distance to reach receivers that are close to the sender. For some applications, increased latency is unwanted.

Although register-encapsulation can continue indefinitely, for these reasons, the RP will switch to native forwarding. To do this, when the RP receives a register-encapsulated data packet from source S on group G, it will initiate an (S,G) source-specific join toward S. This join message travels hop-by-hop toward S, instantiating the (S,G) multicast tree state in the routers along the path. The (S,G) multicast tree state is used only to forward packets for group G if those packets come from source S. Eventually the join message reaches S's subnet or a router that already has the (S,G) multicast tree state, and packets from S start to flow following the (S,G) tree state toward the RP. These data packets can also reach routers with the (*,G) state along the path toward the RP, and if this occurs, they can take a shortcut to the RP tree at this point.

While the RP is in the process of joining the source-specific tree for S, the data packets will continue being encapsulated to the RP. When packets from S also start to arrive natively at the RP, the RP receives two copies of each of these packets. At this point, the RP starts to discard the encapsulated copy of these packets and sends a register-stop message back to S's DR to prevent the DR unnecessarily encapsulating the packets. At the end of phase 2, traffic will be flowing natively from S along a source-specific tree to the RP and from there along the shared tree to the receivers. Where the two trees intersect, traffic can transfer from the shared RP tree to the shorter source tree.



Note:

A sender can start sending before or after a receiver joins the group, and therefore, phase two may occur before the shared tree to the receiver is built.

2.2.2.1.3 Phase three

In this phase, the RP joins back toward the source using the shortest path tree. Although having the RP join back toward the source removes the encapsulation overhead, it does not completely optimize the forwarding paths. For many receivers, the route via the RP can involve a significant detour when compared with the shortest path from the source to the receiver.

To obtain lower latencies, a router on the receiver's LAN, typically the DR, may optionally initiate a transfer from the shared tree to a source-specific shortest-path tree (SPT). To do this, it issues an (S,G) Join toward S. This instantiates the state in the routers along the path to S. Eventually, this join either reaches S's subnet or reaches a router that already has the (S,G) state. When this happens, data packets from S start to flow following the (S,G) state until they reach the receiver.

At this point, the receiver (or a router upstream of the receiver) receives two copies of the data — one from the SPT and one from the RPT. When the first traffic starts to arrive from the SPT, the DR or upstream router starts to drop the packets for G from S that arrive via the RP tree. In addition, it sends an (S,G) prune message toward the RP. The prune message travels hop-by-hop instantiating the state along the path toward the RP indicating that traffic from S for G should not be forwarded in this direction. The prune message is propagated until it reaches the RP or a router that still needs the traffic from S for other receivers.

By now, the receiver is receiving traffic from S along the SPT between the receiver and S. In addition, the RP is receiving the traffic from S, but this traffic is no longer reaching the receiver along the RP tree. As far as the receiver is concerned, this is the final distribution tree.

2.2.2.2 Encapsulating data packets in the register tunnel

Conceptually, the register tunnel is an interface with a smaller MTU than the underlying IP interface toward the RP. IP fragmentation on packets forwarded on the register tunnel is performed based on this smaller MTU. The encapsulating DR can perform path-MTU discovery to the RP to determine the effective MTU of the tunnel. This smaller MTU takes both the outer IP header and the PIM register header overhead into consideration.

2.2.2.3 PIM bootstrap router mechanism

For proper operation, every PIM-SM router within a PIM domain must be able to map a particular global-scope multicast group address to the same RP. If this is not possible, black holes can appear (this is where some receivers in the domain cannot receive some groups). A domain in this context is a contiguous set of routers that all implement PIM and are configured to operate within a common boundary.

The bootstrap router (BSR) mechanism provides a way in which viable group-to-RP mappings can be created and distributed to all the PIM-SM routers in a domain. Each candidate BSR originates bootstrap messages (BSMs). Each BSM contains a BSR priority field. Routers within the domain flood the BSMs throughout the domain. A candidate BSR that hears about a higher-priority candidate BSR suppresses its sending of further BSMs for a period of time. The single remaining candidate BSR becomes the elected BSR and its BSMs inform the other routers in the domain that it is the elected BSR.

The PIM bootstrap routing mechanism is adaptive, meaning that if an RP becomes unreachable, it will be detected and the mapping tables will be modified so that the unreachable RP is no longer used and the new tables will be rapidly distributed throughout the domain.

2.2.2.4 PIM-SM routing policies

Multicast traffic can be restricted from certain source addresses by creating routing policies. Join messages can be filtered using import filters. PIM join policies can be used to reduce denial of service attacks and subsequent PIM state explosion in the router and to remove unwanted multicast streams at the edge of the network before it is carried across the core. Route policies are created in the **config>router>policy-options** context. Join and register route policy match criteria for PIM-SM can specify the following:

- router interface or interfaces specified by name or IP address
- neighbor address (the source address in the IP header of the join and prune message)
- multicast group address embedded in the join and prune message
- multicast source address embedded in the join and prune message

Join policies can be used to filter PIM join messages so that no *,G or S,G state is created on the router. The following table describes the match conditions.

Table 5: Join filter policy match conditions

Match condition	Matches
Interface	The RTR interface by name
Neighbor	The neighbors source address in the IP header
Group Address	The multicast group address in the join/prune message
Source Address	The source address in the join/prune message

PIM register messages are sent by the first hop designated router that has a direct connection to the source. This serves a dual purpose:

- It notifies the RP that a source has active data for the group.
- It delivers the multicast stream in register encapsulation to the RP and its potential receivers.
- If no one has joined the group at the RP, the RP will ignore the registers.

In an environment where the sources to particular multicast groups are always known, it is possible to apply register filters at the RP to prevent any unwanted sources from transmitting a multicast stream. You can apply these filters at the edge so that register data does not travel unnecessarily over the network toward the RP.

The following table describes the match conditions.

Table 6: Register filter policy match conditions

Match condition	Matches
Interface	The RTR interface by name
Group Address	The multicast group address in the join/prune message
Source Address	The source address in the join/prune message

2.2.2.5 Reverse path forwarding checks

Multicast implements a reverse path forwarding check (RPF). An RPF checks the path that multicast packets take between their sources and the destinations to prevent loops. Multicast requires that an incoming interface be the outgoing interface used by unicast routing to reach the source of the multicast packet. RPF forwards a multicast packet only if it is received on an interface that is used by the router to route to the source.

If the forwarding paths are modified due to routing topology changes, any dynamic filters that may have been applied must be reevaluated. If filters are removed, the associated alarms are also cleared.

2.2.2.5.1 Anycast RP for PIM-SM

The implementation of anycast RP for PIM-SM environments enables fast convergence when a PIM rendezvous point (RP) router fails by allowing receivers and sources to rendezvous at the closest RP. It allows an arbitrary number of RPs per group in a single shared-tree protocol Independent Multicast-Sparse Mode (PIM-SM) domain. This is particularly important for triple play configurations that choose to distribute multicast traffic using PIM-SM, not SSM. In this case, RP convergence must be fast enough to avoid the loss of multicast streams, which could cause loss-of-TV delivery to the end customer.

Anycast RP for PIM-SM environments are supported in the base routing/PIM-SM instance of the service router. This feature is supported in Layer 3-VRPN instances that are configured with PIM.

2.2.2.5.1.1 Implementation

The Anycast RP for PIM-SM implementation is defined in RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*, and is similar to that described in RFC 3446, *Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*. The implementation extends the register mechanism in PIM so that anycast RP functionality can be retained without using Multicast Source Discovery Protocol (MSDP).


The mechanism works as follows:

- An IP address is chosen as the RP address. This address is statically configured, or distributed using a dynamic protocol, to all PIM routers throughout the domain.
- A set of routers in the domain are chosen to act as RPs for this RP address. These routers are called the anycast-RP set.
- Each router in the anycast-RP set is configured with a loopback interface using the RP address.

- Each router in the anycast-RP set also needs a separate IP address to be used for communication between the RPs.
- The RP address, or a prefix that covers the RP address, is injected into the unicast routing system inside of the domain.
- Each router in the anycast-RP set is configured with the addresses of all other routers in the anycast-RP set. This must be consistently configured in all RPs in the set.

The following figure shows a scenario where all routers are connected, and where R1A, R1B, and R2 are receivers for a group, and S1 and S2 send to that group. Assume RP1, RP2, and RP3 are all assigned the same IP address that is used as the anycast-RP address (for example, the IP address is RPA).

Figure 1: Anycast RP for PIM-SM implementation example



OSSG271



Note:

The address used for the RP address in the domain (the anycast-RP address) must be different from the addresses used by the anycast-RP routers to communicate with each other.

The following procedure is used when S1 starts sourcing traffic:

1. S1 sends a multicast packet.
2. The DR directly attached to S1 forms a PIM register message to send to the anycast-RP address (RPA). The unicast routing system delivers the PIM register message to the nearest RP, in this case RP1.
3. RP1 receives the PIM register message, de-encapsulates it, and sends the packet down the shared tree to get the packet to receivers R1A and R1B.
4. RP1 is configured with RP2 and RP3's IP address. Because the register message did not come from one of the RPs in the anycast-RP set, RP1 assumes the packet came from a DR. If the register message is not addressed to the anycast-RP address, an error has occurred and it should be rate-limited logged.
5. RP1 sends a copy of the register message from S1's DR to both RP2 and RP3. RP1 uses its own IP address as the source address for the PIM register message.
6. RP1 may join back to the source tree by triggering a (S1,G) Join message toward S1; however, RP1 must create the (S1,G) state.
7. RP2 receives the register message from RP1, de-encapsulates it, and also sends the packet down the shared tree to get the packet to receiver R2.
8. RP2 sends a register-stop message back to the RP1. RP2 may wait to send the register-stop message if it decides to join the source tree. RP2 should wait until it has received data from the source on the source tree before sending the register-stop message. If RP2 decides to wait, the register-stop message will be sent when the next register is received. If RP2 decides not to wait, the register-stop message is sent now.

9. RP2 may join back to the source tree by triggering a (S1,G) Join message toward S1; however, RP2 must create the (S1,G) state.
10. RP3 receives the register message from RP1 and de-encapsulates it, but, since there are no receivers joined for the group, it can discard the packet.
11. RP3 sends a register-stop message back to RP1.
12. RP3 creates a (S1,G) state so that when a receiver joins after S1 starts sending, RP3 can join quickly to the source tree for S1.
13. RP1 processes the register-stop message from RP2 and RP3. RP1 may cache on a per-RP/per-(S,G) basis the receipt of register-stop messages from the RPs in the anycast-RP set. This option is performed to increase the reliability of register message delivery to each RP. When this option is used, subsequent register messages received by RP1 are sent only to the RPs in the anycast-RP set that have not previously sent register-stop messages for the (S,G) entry.
14. RP1 sends a register-stop message back to the DR the next time a register message is received from the DR and, if all RPs in the anycast-RP set have returned register-stop messages for a particular (S,G) route when RP1 caches on a per-RP/per-(S,G) basis the receipt of register-stop messages from the RPs in the anycast-RP set.

The procedure for S2 sending follows the same preceding steps, but it is RP3 that sends a copy of the register originated by S2's DR to RP1 and RP2. This example shows how sources anywhere in the domain, associated with different RPs, can reach all receivers, also associated with different RPs, in the same domain.

2.2.2.6 Distributing PIM joins over multiple ECMP paths

The per bandwidth/round robin method is commonly used in multicast load balancing. However, the interface in an ECMP set can also be used for a channel to be predictable without any knowledge of the other channels using the ECMP set.

The **mc-ecmp-hashing-enabled** command enables PIM joins to be distributed over multiple ECMP paths based on a hash of S and G. When a link in the ECMP set is removed, the multicast streams using the link are redistributed over the remaining ECMP links using the same hash algorithm. When a link is added to the ECMP set, new joins may be allocated to the new link based on the hash algorithm, but existing multicast streams using the other ECMP links stay on those links until they are pruned.

The default is **no mc-ecmp-hashing-enabled**, which means that the use of multiple ECMP paths is controlled by the existing implementation and CLI commands, that is, **mc-ecmp-balance**.

The **mc-ecmp-hashing-enabled** command and the **mc-ecmp-balance** command are mutually exclusive in the same context.

The following procedure is used to achieve distribution of streams across the ECMP links:

1. For a specific S, G get all possible nHops.
2. Sort these nHops based on nHop addresses.
3. xor S and G addresses.
4. Hash the xor address over a number of PIM next hops.
5. Use the hash value obtained in step 4, and get that element, in the sorted list obtained in step 2, as the preferred nHop.
6. If this element is not available or it is not a PIM nHop (PIM neighbor), the next available next hop is chosen.

Example: PIM status

The following is a sample PIM status indicating ECMP hashing is disabled.

```
*B:BB# show router pim status

=====
PIM Status ipv4
=====
Admin State                : Up
Oper State                 : Up

IPv4 Admin State           : Up
IPv4 Oper State            : Up

BSR State                  : Accept Any

Elected BSR
  Address                  : None
  Expiry Time              : N/A
  Priority                  : N/A
  Hash Mask Length         : 30
  Up Time                  : N/A
  RPF Intf toward E-BSR   : N/A

Candidate BSR
  Admin State              : Down
  Oper State               : Down
  Address                  : None
  Priority                  : 0
  Hash Mask Length         : 30

Candidate RP
  Admin State              : Down
  Oper State               : Down
  Address                  : 0.0.0.0
  Priority                  : 192
  Holdtime                 : 150

SSM-Default-Range          : Enabled
SSM-Group-Range
  None

MC-ECMP-Hashing            : Disabled

Policy                     : None

RPF Table                  : rtable-u

Non-DR-Attract-Traffic    : Disabled
=====

-----
*B:BB>config>service>vprn>pim# no mc-ecmp-balance mc-ecmp-balance mc-ecmp-balance
-hold
*B:BB>config>service>vprn>pim# no mc-ecmp-balance
*B:BB>config>service>vprn>pim# mc-ecmp-mc-ecmp-balance mc-ecmp-balance-hold mc-ecmp
-hashing-enabled
*B:BB>config>service>vprn>pim# mc-ecmp-hashing-enabled
*B:BB>config>service>vprn>pim# info
-----

          apply-to all
          rp
          static
```

```

        address 10.3.3.3
        group-prefix 224.0.0.0/4
        exit
    exit
    bsr-candidate
        shutdown
    exit
    rp-candidate
        shutdown
    exit
exit
no mc-ecmp-balance
mc-ecmp-hashing-enabled
-----
*B:BB>config>service>vprn>pim#
apply-to      - Create/remove interfaces in PIM
[no] import    - Configure import policies
[no] interface  + Configure PIM interface
[no] mc-ecmp-balance - Enable/
Disable multicast balancing of traffic over ECMP links
[no] mc-ecmp-balanc* - Configure hold time for multicast balancing over ECMP links
[no] mc-ecmp-hashin* - Enable/
Disable hash based multicast balancing of traffic over ECMP links
[no] non-dr-attract* - Enable/disable attracting traffic when not DR
rp            + Configure the router as static or Candidate-RP
[no] shutdown  - Administratively enable or disable the operation of PIM
[no] spt-switchover* - Configure shortest path tree (spt tree) switchover
threshold for a group prefix
[no] ssm-default-ra* - Enable the disabling of SSM Default Range
[no] ssm-groups    + Configure the SSM group ranges

```

Example: PIM joins over multiple ECMP paths

The following is a sample distribution of PIM joins over multiple ECMP paths.

```

*A:BA# show router pim group
=====
PIM Groups ipv4
=====

```

Group Address Source Address	Type RP	Spt Bit Inc Intf	No.0ifs
239.1.1.1	(S,G)	spt to_C0	1
172.0.100.33	10.20.1.6		
239.1.1.2	(S,G)	spt to_C3	1
172.0.100.33	10.20.1.6		
239.1.1.3	(S,G)	spt to_C2	1
172.0.100.33	10.20.1.6		
239.1.1.4	(S,G)	spt to_C1	1
172.0.100.33	10.20.1.6		
239.1.1.5	(S,G)	spt to_C0	1
172.0.100.33	10.20.1.6		
239.1.1.6	(S,G)	spt to_C3	1
172.0.100.33	10.20.1.6		
239.2.1.1	(S,G)	spt to_C0	1
172.0.100.33	10.20.1.6		
239.2.1.2	(S,G)	spt to_C3	1
172.0.100.33	10.20.1.6		
239.2.1.3	(S,G)	spt to_C2	1
172.0.100.33	10.20.1.6		
239.2.1.4	(S,G)	spt to_C1	1

172.0.100.33	10.20.1.6		
239.2.1.5	(S,G) spt	to_C0	1
172.0.100.33	10.20.1.6		
239.2.1.6	(S,G) spt	to_C3	1
172.0.100.33	10.20.1.6		
239.3.1.1	(S,G) spt	to_C0	1
172.0.100.33	10.20.1.6		
239.3.1.2	(S,G) spt	to_C3	1
172.0.100.33	10.20.1.6		
239.3.1.3	(S,G) spt	to_C2	1
172.0.100.33	10.20.1.6		
239.3.1.4	(S,G) spt	to_C1	1
172.0.100.33	10.20.1.6		
239.3.1.5	(S,G) spt	to_C0	1
172.0.100.33	10.20.1.6		
239.3.1.6	(S,G) spt	to_C3	1
172.0.100.33	10.20.1.6		
239.4.1.1	(S,G) spt	to_C0	1
172.0.100.33	10.20.1.6		
239.4.1.2	(S,G) spt	to_C3	1
172.0.100.33	10.20.1.6		
239.4.1.3	(S,G) spt	to_C2	1
172.0.100.33	10.20.1.6		
239.4.1.4	(S,G) spt	to_C1	1
172.0.100.33	10.20.1.6		
239.4.1.5	(S,G) spt	to_C0	1
172.0.100.33	10.20.1.6		
239.4.1.6	(S,G) spt	to_C3	1
172.0.100.33	10.20.1.6		

Groups : 24			
=====			

2.2.3 Multicast debugging tools

This section describes multicast debugging tools for the 7210 SAS.

The debugging tools for multicast consist of two elements: mtrace and mrimf.

2.2.3.1 Mtrace

Assessing problems in the distribution of IP multicast traffic can be difficult. The **mtrace** feature uses a tracing feature implemented in multicast routers that is accessed via an extension to the IGMP protocol. The **mtrace** feature is used to print the path from the source to a receiver; it does this by passing a trace query hop-by-hop along the reverse path from the receiver to the source. At each hop, information such as the hop address, routing error conditions, and packet statistics are gathered and returned to the requester.

Data added by each hop includes:

- query arrival time
- incoming interface
- outgoing interface
- previous hop router address
- input packet count

- output packet count
- total packets for this source/group
- Routing protocol
- TTL threshold
- Forwarding/error code

The information enables the network administrator to determine the following:

- where multicast flows stop
- the flow of the multicast stream

When the trace response packet reaches the first-hop router (the router that is directly connected to the source's network interface), that router sends the completed response to the response destination (receiver) address specified in the trace query.

If a multicast router along the path does not implement the traceroute feature or if there is an outage, no response is returned. To solve this problem, the trace query includes a maximum hop count field to limit the number of hops traced before the response is returned. This allows a partial path to be traced.

The reports inserted by each router contain not only the address of the hop, but also the TTL required to forward, flags to indicate routing errors, and counts of the total number of packets on the incoming and outgoing interfaces and those forwarded for the specified group. Examining the differences in these counts for two separate traces and comparing the output packet counts from one hop with the input packet counts of the next hop allows the calculation of packet rate and packet loss statistics for each hop to isolate congestion problems.

2.2.3.1.1 Finding the last hop router

The trace query must be sent to the multicast router, which is the last hop on the path from the source to the receiver. If the receiver is on the local subnet (as determined using the subnet mask), the default method is to multicast the trace query to all-routers.mcast.net (224.0.0.2) with a TTL of 1. Otherwise, the trace query is sent to the group address because the last-hop router will be a member of that group if the receiver is. Therefore, it is necessary to specify a group that the intended receiver has joined. This multicast is sent with a default TTL of 64, which may not be sufficient for all cases.

When tracing from a multihomed host or router, the default receiver address may not be the desired interface for the path from the source. In such cases, the desired interface should be specified explicitly as the receiver.

2.2.3.1.2 Directing the response

Unless the number of hops to trace is explicitly set with the hop option, mtrace first attempts to trace the full reverse path by default. If there is no response within a 3 second timeout interval, a "*" is displayed and the probing switches to hop-by-hop mode. Trace queries are issued starting with a maximum hop count of one and increasing by one until the full path is traced or no response is received. At each hop, multiple probes are sent. The first attempt is made with the unicast address of the host running mtrace as the destination for the response. Since the unicast route may be blocked, the remainder of attempts request that the response be multicast to mtrace.mcast.net (224.0.1.32) with the TTL set to 32 more than what is needed to pass the thresholds seen so far along the path to the receiver. For the final attempts, the TTL is increased by another 32.

Alternatively, the TTL may be set explicitly with the TTL option.

The output of mtrace is a short listing of the hops in the order they are queried, that is, in the reverse of the order from the source to the receiver. For each hop, a line is displayed showing:

- the hop number (counted negatively to indicate that this is the reverse path)
- the multicast routing protocol
- the threshold required to forward data (to the previous hop in the listing as indicated by the up-arrow character)
- the cumulative delay for the query to reach that hop (valid only if the clocks are synchronized)

The response ends with a line showing the round-trip time, which measures the interval from the time the query is issued until the response is received, both derived from the local system clock.

Mtrace packets use special IGMP packets with IGMP type codes of 0x1E and 0x1F.

2.2.3.2 Mrinfo

The **mrinfo** feature is a simple mechanism to display configuration information from the target multicast router. The type of information displayed includes the multicast capabilities of the router, code version, metrics, TTL thresholds, protocols, and status. This information can be used by network operators to verify if bidirectional adjacencies exist. When the specified multicast router responds, the configuration is displayed.

2.2.4 Configuration guidelines for 7210 SAS

The following are the configuration guidelines for the 7210 SAS:

- 7210 SAS platforms can be used as RPs.
- Static RP configuration using PIM BSR messages is supported.
- It is possible to configure the 7210 SAS as a First Hop Multicast router (FHR) from the source in a PIM-SM network.
- 7210 SAS devices provide an option to either switch over to the SPT or continue to use the shared tree. However, the traffic rate threshold cannot be configured to trigger the switch over.
- RFP checks are performed using the unicast routing table. Multicast BGP and multicast routing table are not supported.

2.3 Configuring multicast parameters with CLI

This section provides information to configure multicast, IGMP, and PIM.

2.3.1 Multicast configuration overview

7210 SAS routers use IGMP to manage membership for a specific multicast session. IGMP is not enabled by default. The IGMP context is not operational until at least one IGMP interface is specified in the context, at which time the interface is enabled for IGMP.

Traffic can only flow away from the router to an IGMP interface, and to and from a PIM interface. A router directly connected to a source must have PIM enabled on the interface to that source. In a network, traffic travels from PIM interface to PIM interface, and arrives on an IGMP-enabled interface.

The IGMP CLI context allows you to specify an existing IP interface and modify the interface-specific parameters. Static IGMP group memberships can be configured to test multicast forwarding without a receiver host. When IGMP static group membership is enabled, data is forwarded to an interface without receiving membership reports from host members.

When static IGMP group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static IGMP group entries do not generate join messages toward the RP. When a host needs to receive multicast sessions, it sends a join message for each multicast group it needs to join. Then, a leave message may be sent for each multicast group it no longer wishes to participate with.

A multicast router keeps a list of multicast group memberships for each attached network, and an interval timer for each membership. Hosts issue a Multicast Group Membership Report when they want to receive a multicast session. The reports are sent to all multicast routers.

PIM is not enabled by default. Because it is an interface function, PIM is not operational until at least one interface is specified in the PIM context, at which time the interface is enabled for PIM. When PIM is enabled, data is forwarded to network segments with active receivers that have explicitly requested the multicast group.

2.3.2 Basic configuration

Perform the following basic multicast configuration tasks.

For IGMP:

- enable IGMP (required)
- configure IGMP interfaces (required)
- specify the IGMP version on the interface (optional)
- configure static (S,G)/(*,G) (optional)
- configure SSM translation (optional)

For PIM:

- enable PIM (required)
- add interfaces so the protocol establishes adjacencies with the neighboring routers (required)
- configure a way to calculate group-to-RP mapping (required) by either:
 - using static group-to-RP mapping
 - enabling the candidate RP/bootstrap mechanism on some routers.
- enable unicast routing protocols to learn routes toward the RP/source for reverse path forwarding (required)
- add SSM ranges (optional)
- enable Candidate BSR (optional)
- enable Candidate RP (optional)
- change the hello interval (optional)
- configure route policies (bootstrap-export, bootstrap-import, import join and register)

Example: Enabled IGMP and PIM configuration output

```
A:LAX>config>router>igmp# info
-----
interface "lax-vls"
exit
interface "pl-ix"
exit
-----
*A:Dut-B>config>router>igmp# info detail
-----
        interface "C_Rx"
            no import
            version 3
            subnet-check
            no max-groups
            no max-sources
            no max-grp-sources
            no disable-router-alert-check
            no query-interval
            no query-last-listener-interval
            no query-response-interval
            no shutdown
        exit
        no grp-if-query-src-ip
        query-interval 125
        query-last-member-interval 1
        query-response-interval 10
        robust-count 2
        no shutdown
-----
A:7210SAS>config>router>igmp# exit
A:7210SAS>config>router# pim
A:7210SAS>config>router>pim# info
-----
        interface "lax-vls"
        exit
        interface "lax-vls"
        exit
        interface "lax-sjc"
        exit
        interface "pl-ix"
        exit
        rp
            static
                address 239.22.187.237
                group-prefix 239.24.24.24/32
            exit
        exit
        shutdown
bsr-candidate
        exit
        rp-candidate
        shutdown
        exit
    exit
-----
A:7210SAS>config>router>pim# info detail
-----
        no import join-policy
        no import register-policy
        interface "system"
            priority 1
            hello-interval 30
```

```
        multicast-senders auto
        no tracking-support
        no shutdown
    exit
    interface "lax-vls"
        priority 1
        hello-interval 30
        multicast-senders auto
        no tracking-support
        no shutdown
    exit
    interface "lax-sjc"
        priority 1
        hello-interval 30
        multicast-senders auto
        no tracking-support
        no shutdown
    exit
    interface "pl-ix"
        priority 1
        hello-interval 30
        multicast-senders auto
        no tracking-support
        no shutdown
    exit
    apply-to none
    rp
        no bootstrap-import
        no bootstrap-export
        static
            address 239.22.187.237
            no override
            group-prefix 239.24.24.24/32
        exit
    exit
        shutdown
        priority 0
        hash-mask-len 30
        no address
    exit
    rp-candidate
        shutdown
bsr-candidate
        no address
        holdtime 150
        priority 192
    exit
    exit
    no shutdown
-----
A:7210SAS>config>router>pim#
```

2.3.3 Common configuration tasks

The following sections describe basic multicast configuration tasks.

2.3.4 Configuring IGMP parameters

This section provides information to configure IGMP parameters.

2.3.4.1 Enabling IGMP

Use the following syntax to enable IGMP.

```
config>router# igmp
```

Example: Enabled IGMP detailed output

```
A:7210SAS>>config>router# info detail
...
#-----
echo "IGMP Configuration"
#-----
      igmp
        query-interval 125
        query-last-member-interval 1
        query-response-interval 10
        robust-count 2
        no shutdown
      exit
#-----
A:7210SAS>>config>system#
```

2.3.4.2 Configuring an IGMP interface

Use the following syntax to configure an IGMP interface.

```
config>router# igmp
  interface ip-int-name
  max-groups value
  import policy-name
  version version
  no shutdown
```

Example: IGMP interface configuration command usage

```
config>router#
  config>router>igmp# interface "lax-vls"
  config>router>igmp>if? no shutdown
  config>router>igmp>if# exit
  config>router>igmp# interface "pl-ix"
  config>router>igmp>if? no shutdown
  config>router>igmp>if# exit
  config>router>igmp# interface "lax-sjc"
  config>router>igmp>if? no shutdown
  config>router>igmp>if# exit
```

Example: IGMP configuration output

```
A:7210SAS>config>router>igmp# info
-----
      interface "lax-sjc"
      exit
      interface "lax-vls"
      exit
      interface "pl-ix"
```

```
exit
-----
A:7210SAS>config>router>igmp# exit
```

2.3.4.3 Configuring static parameters

Use the following syntax to add an IGMP static multicast source.

```
config>router# igmp
interface ip-int-name
no shutdown
static
    group grp-ip-address
    source ip-address
```

Example: Command usage

Use the following syntax to configure static group addresses and source addresses for the SSM translate group ranges.

```
config>router>igmp# interface lax-vls
config>router>igmp>if# static
config>router>igmp>if>static# group 239.255.0.2
config>router>igmp>if>static>group# source 172.22.184.197
config>router>igmp>if>static>group# exit
config>router>igmp>if>static# exit
config>router>igmp>if# exit
```

Example: Configuration output

```
A:LAX>config>router>igmp# info
-----
    interface "lax-sjc"
    exit
    interface "lax-vls"
        static
            group 239.255.0.2
            source 172.22.184.197
        exit
    exit
    interface "pl-ix"
    exit
-----
A:LAX>config>router>igmp#
```

Use the following syntax to add an IGMP static starg entry.

```
config>router# igmp
interface ip-int-name
no shutdown
static
    group grp-ip-address
    starg
```

Example: Command usage

Use the following syntax to configure static group addresses and add a static (*,G) entry.

```
config>router>igmp# interface lax-sjc
config>router>igmp>if# static
config>router>igmp>if>static# group 239.1.1.1
config>router>igmp>if>static>group# starg
config>router>igmp>if>static>group# exit
config>router>igmp>if>static# exit
config>router>igmp>if# exit
config>router>igmp#
```

Example: Configuration output

```
A:LAX>config>router>igmp# info
-----
      interface "lax-sjc"
        static
          group 239.1.1.1
          starg
        exit
      exit
    exit
  interface "lax-vls"
    static
      group 239.255.0.2
      source 172.22.184.197
    exit
  exit
exit
interface "pl-ix"
exit
-----
A:LAX>config>router>igmp#
```

2.3.4.4 Configuring SSM translation

Use the following CLI syntax to configure IGMP parameters.

```
config>router# igmp
  ssm-translate
  grp-range start end
  source ip-address
```

Example: Command usage to configure IGMP parameters

```
config>router# igmp
config>router>igmp# ssm-translate
config>router>igmp>ssm# grp-range 239.255.0.1 231.2.2.2
config>router>igmp>ssm>grp-range# source 10.1.1.1
```

Example: Configuration output

```
A:LAX>config>router>igmp# info
-----
      ssm-translate
      grp-range 239.255.0.1 239.2.2.2
-----
```

```

        source 10.1.1.1
    exit
exit
interface "lax-sjc"
    static
        group 239.1.1.1
        starg
    exit
exit
exit
interface "lax-vls"
    static
        group 239.255.0.2
        source 172.22.184.197
    exit
exit
exit
interface "pl-ix"
exit
-----
A:LAX>config>router>igmp# exit

```

2.3.5 Configuring PIM parameters

The following section describes the syntax used to configure the PIM parameters.

2.3.5.1 Enabling PIM

PIM must be enabled on all interfaces for the routing instance; failure to do so might result in multicast routing errors.

Use the following syntax to enable PIM.

```
config>router# pim
```

Example: Detailed output of an enabled PIM

```

A:LAX>>config>router# info detail
...
#-----
echo "PIM Configuration"
#-----
    pim
        no import join-policy
        no import register-policy
        apply-to none
        rp
            no bootstrap-import
            no bootstrap-export
            static
            exit
            shutdown
            priority 0
            hash-mask-len 30
            no address
        exit
        rp-candidate
        shutdown

```



```

                no address
                holdtime 150
                priority 192
            exit
        exit
        no shutdown
    exit
#-----
...
A:LAX>>config>system#

```

2.3.5.2 Configuring PIM interface parameters

The following examples show the command usage to configure PIM interface parameters and the resulting configuration outputs.

Example: Command usage 1

```

A:LAX>config>router# pim
A:LAX>config>router>pim# interface "system"
A:LAX>config>router>pim>if# exit
A:LAX>config>router>pim# interface "lax-vls"
A:LAX>config>router>pim>if# exit
A:LAX>config>router>pim# interface "lax-sjc"
A:LAX>config>router>pim>if# exit
A:LAX>config>router>pim# interface "pl-ix"
A:LAX>config>router>pim>if# exit
A:LAX>config>router>pim# rp
A:LAX>config>router>pim>rp# static
A:LAX>config>router>pim>rp>static# address 239.22.187.237
A:LAX>config>router>.>address# group-prefix 239.24.24.24/32
A:LAX>config>router>pim>rp>static>address# exit
A:LAX>config>router>pim>rp>static# exit
A:LAX>config>router>pim>rp# exit

```

Example: Configuration output 1

```

A:LAX>config>router>pim# info
-----
        interface "system"
        exit
        interface "lax-vls"
        exit
        interface "lax-sjc"
        exit
        interface "pl-ix"
        exit
        rp
            static
                address 239.22.187.237
                group-prefix 239.24.24.24/32
                exit
                address 10.10.10.10
                exit
            exit
            shutdown
bsr-candidate
        exit
        rp-candidate

```

```
                shutdown
            exit
        exit
    -----
A:LAX>config>router>pim#
```

Example: Command usage 2

```
A:SJC>config>router# pim
A:SJC>config>router>pim# interface "system"
A:SJC>config>router>pim>if# exit
A:SJC>config>router>pim# interface "sjc-lax"
A:SJC>config>router>pim>if# exit
A:SJC>config>router>pim# interface "sjc-nyc"
A:SJC>config>router>pim>if# exit
A:SJC>config>router>pim# interface "sjc-sfo"
A:SJC>config>router>pim>if# exit
A:SJC>config>router>pim# rp
A:SJC>config>router>pim>rp# static
A:SJC>config>router>pim>rp>static# address 239.22.187.237
A:SJC>config>router>pim>rp>static>address# group-prefix 239.24.24.24/32
A:SJC>config>router>pim>rp>static>address# exit
A:SJC>config>router>pim>rp>static# exit
A:SJC>config>router>pim>rp# exit
A:SJC>config>router>pim#
```

Example: Configuration output 2

```
A:SJC>config>router>pim# info
-----
        interface "system"
        exit
        interface "sjc-lax"
        exit
        interface "sjc-nyc"
        exit
        interface "sjc-sfo"
        exit
        rp
            static
                address 239.22.187.237
                group-prefix 239.24.24.24/32
            exit
        exit
        shutdown
bsr-candidate
        exit
        rp-candidate
        shutdown
        exit
    exit
    -----
A:SJC>config>router>pim#
```

Example: Command usage 3

```
A:MV>config>router# pim
A:MV>config>router>pim# interface "system"
A:MV>config>router>pim>if# exit
```

```
A:MV>config>router>pim# interface "mv-sfo"
A:MV>config>router>pim>if# exit
A:MV>config>router>pim# interface "mv-vlc"
A:MV>config>router>pim>if# exit
A:MV>config>router>pim# interface "p3-ix"
A:MV>config>router>pim>if# exit
A:MV>config>router>pim# rp
A:MV>config>router>pim>rp# static
A:MV>config>router>pim>rp>static# address 239.22.187.237
A:MV>config>router>pim>rp>static>address# group-prefix 239.24.24.24/32
A:MV>config>router>pim>rp>static>address# exit
A:MV>config>router>pim>rp>static#
A:MV>config>router>pim>rp# exit
A:MV>config>router>pim#
```

Example: Configuration output 3

```
A:MV>config>router>pim# info
-----
      interface "system"
      exit
      interface "mv-sfo"
      exit
      interface "mv-vlc"
      exit
      interface "p3-ix"
      exit
      rp
        static
          address 239.22.187.237
          group-prefix 239.24.24.24/32
          exit
        exit
          address 239.22.187.236
          no shutdown
        exit
        rp-candidate
          address 239.22.187.236
          no shutdown
      bsr-candidate
        exit
      exit
-----
A:MV>config>router>pim#
```

Example: Command usage 4

```
A:SF0>config>router# pim
A:SF0>config>router>pim# interface "system"
A:SF0>config>router>pim>if# exit
A:SF0>config>router>pim# interface "sfo-sfc"
A:SF0>config>router>pim>if# exit
A:SF0>config>router>pim# interface "sfo-was"
A:SF0>config>router>pim>if# exit
A:SF0>config>router>pim# interface "sfo-mv"
A:SF0>config>router>pim>if# exit
A:SF0>config>router>pim# rp
A:SF0>config>router>pim>rp# static
A:SF0>config>router>pim>rp>static# address 239.22.187.237
A:SF0>config>router>pim>rp>static>address# group-prefix 239.24.24.24/32
A:SF0>config>router>pim>rp>static>address# exit
```

```
A:SF0>config>router>pim>rp>static# exit
A:SF0>config>router>pim>rp # exit
A:SF0>config>router>pim#
```

Example: Configuration output 4

```
A:SF0>config>router>pim# info
-----
      interface "system"
      exit
      interface "sfo-sjc"
      exit
      interface "sfo-was"
      exit
      interface "sfo-mv"
      exit
      rp
        static
          address 239.22.187.237
          group-prefix 239.24.24.24/32
          exit
        exit
        address 239.22.187.239
        no shutdown
      exit
      rp-candidate
        address 239.22.187.239
        no shutdown
      bsr-candidate
        exit
      exit
-----
A:SF0>config>router>pim#
```

Example: Command usage 5

```
A:WAS>config>router# pim
A:WAS>config>router>pim# interface "system"
A:WAS>config>router>pim>if# exit
A:WAS>config>router>pim# interface "was-sfo"
A:WAS>config>router>pim>if# exit
A:WAS>config>router>pim# interface "was-vlc"
A:WAS>config>router>pim>if# exit
A:WAS>config>router>pim# interface "p4-ix"
A:WAS>config>router>pim>if# exit
A:WAS>config>router>pim# rp
A:WAS>config>router>pim>rp# static
A:WAS>config>router>pim>rp>static# address 239.22.187.237
A:WAS>config>router>pim>rp>static>address# group-prefix 239.24.24.24/32
A:WAS>config>router>pim>rp>static>address# exit
A:WAS>config>router>pim>rp>static# exit
A:WAS>config>router>pim>rp# exit
A:WAS>config>router>pim#
```

Example: Configuration output 5

```
A:WAS>config>router>pim# info
-----
      interface "system"
      exit
-----
```

```
interface "was-sfo"
exit
interface "was-vlc"
exit
interface "p4-ix"
exit
rp
  static
    address 239.22.187.237
    group-prefix 239.24.24.24/32
  exit
exit
  address 239.22.187.240
  no shutdown
exit
rp-candidate
  address 239.22.187.240
  no shutdown
bsr-candidate
  exit
exit
-----
A:WAS>config>router>pim#
```

2.3.5.3 Importing PIM join or register policies

The **import** command provides a mechanism to control the (*,G) and (S,G) state that is created on a router. Import policies are defined in the **config>router>policy-options** context.



Note:

In the import policy, if a policy **action** is not specified in the **entry**, the **default-action** takes precedence. In the same way, if there are no entry matches, the **default-action** takes precedence. If no **default-action** is specified, the default **default-action** is executed.

Use the following syntax to configure PIM parameters.

```
config>router# pim
  import {join-policy|register-policy} [policy-name]
[.. policy-name]
```

Output example: Applying the policy statement

The following example shows the command usage to apply the policy statement, which does not allow join messages for group 229.50.50.208/32 and source 192.168.0.0/16, but allows join messages for 192.168.0.0/16, 229.50.50.208 (see [Configuring route policy components](#)).

```
config>router# pim
config>router>pim# import join-policy "foo"
config>router>pim# no shutdown
```

Example: PIM configuration output

```
A:LAX>config>router>pim# info
-----
import join-policy "foo"
interface "system"
exit
```

```
interface "lax-vls"
exit
interface "lax-sjc"
exit
interface "pl-ix"
exit
rp
  static
    address 239.22.187.237
    group-prefix 239.24.24.24/3
    exit
    address 10.10.10.10
    exit
  exit
  shutdown
exit
rp-candidate
  shutdown
exit
exit
-----
A:LAX>config>router>pim#
```

2.3.6 Disabling IGMP or PIM

Use the following syntax to disable IGMP and PIM.

```
config>router#
igmp
shutdown
pim
shutdown
```

Example: Command usage to disable multicast

```
config>router# igmp
config>router>igmp# shutdown
config>router>igmp# exit
config>router#
config>router# pim
config>router>pim# shutdown
config>router>pim# exit
```

Example: Configuration output

```
A:LAX>config>router# info
-----
...
#-----
echo "IGMP Configuration"
#-----
    igmp
      shutdown
      ssm-translate
      grp-range 239.255.0.1 231.2.2.2
      source 10.1.1.1
    exit
  exit
  interface "lax-sjc"
```

```
        static
        group 239.1.1.1
        starg
        exit
    exit
exit
interface "lax-vls"
    static
    group 239.255.0.2
    source 172.22.184.197
    exit
    exit
exit
interface "p1-ix"
exit
exit
#-----
echo "PIM Configuration"
#-----
    pim
    shutdown
    import join-policy "foo"
    interface "system"
    exit
    interface "lax-sjc"
    exit
    interface "lax-vls"
    exit
    interface "p1-ix"
    exit
    rp
        static
        address 239.22.187.237
        group-prefix 239.24.24.24/32
        exit
        address 10.10.10.10
        exit
    exit
    shutdown
    exit
    rp-candidate
    shutdown
bsr-candidate
    exit
    exit
    exit
#-----
....
-----
A:LAX>config>router#
```

2.4 Multicast command reference

2.4.1 Command hierarchies

- [Configuration commands](#)
- [IGMP commands](#)

- [PIM commands](#)
- [Operational commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

2.4.1.1 Configuration commands

```
config
- router
  - mc-maximum-routes number [log-only] [threshold threshold]
  - no mc-maximum-routes
```

2.4.1.2 IGMP commands

```
config
- router
  - [no] igmp
    - [no] interface ip-int-name
      - [no] disable-router-alert-check
      - import policy-name
      - no import
      - max-groups [value]
      - no max-groups
      - max-sources [value]
      - no max-sources
      - query-interval seconds
      - no query-interval
      - query-last-listener-interval seconds
      - no query-last-listener-interval
      - query-response-interval seconds
      - no query-response-interval
      - [no] shutdown
      - ssm-translate
        - [no] grp-range start end
        - [no] source ip-address
      - static
        - [no] group grp-ip-address
        - [no] source ip-address
        - [no] starg
      - [no] subnet-check
      - version version
      - no version
    - query-interval seconds
    - no query-interval
    - query-last-member-interval seconds
    - no query-last-member-interval
    - query-response-interval seconds
    - no query-response-interval
    - robust-count robust-count
    - no robust-count
    - [no] shutdown
    - ssm-translate
      - [no] grp-range start end
      - [no] source ip-address
```


2.4.1.3 PIM commands

```

config
- router
- [no] pim
- [no] enable-mdt-spt
- import {join-policy | register-policy} policy-name [policy-name (up to 5 max)]
- no import {join-policy | register-policy}
- [no] interface ip-int-name
- assert-period assert-period
- no assert-period
- [no] bfd-enable [ipv4]
- [no] bsm-check-rtr-alert
- hello-interval hello-interval
- no hello-interval
- hello-multiplier deci-units
- no hello-multiplier
- [no] improved-assert
- [no] instant-prune-echo
- max-groups value
- no max-groups
- multicast-senders {auto | always | never}
- no multicast-senders
- priority dr-priority
- no priority
- [no] shutdown
- sticky-dr [priority dr-priority]
- no sticky-dr
- three-way-hello [compatibility-mode]
- no three-way-hello
- [no] tracking-support
- [no] mc-ecmp-balance
- mc-ecmp-balance-hold minutes
- no mc-ecmp-balance-hold
- [no] mc-ecmp-hashing-enabled
- [no] non-dr-attract-traffic
- rp
- [no] anycast rp-ip-address
- [no] rp-set-peer ip-address
- bootstrap-export policy-name [.. policy-name ...(up to 5 max)]
- no bootstrap-export
- bootstrap-import policy-name [.. policy-name ...(up to 5 max)]
- no bootstrap-import
- bsr-candidate
- address ip-address
- no address
- hash-mask-len hash-mask-length
- no hash-mask-len
- priority bootstrap-priority
- no priority
- [no] shutdown
- rp-candidate
- address ip-address
- no address
- [no] group-range {grp-ip-address/mask | grp-ip-address netmask}
- holdtime holdtime
- no holdtime
- priority priority
- no priority
- [no] shutdown
- static
- [no] address ip-address

```

```

- [no] group-prefix {grp-ip-address/mask | grp-ip-address netmask}
- [no] override
- [no] rpf-table rtable-u
- [no] shutdown
- spt-switchover-threshold {grp-ipv4-prefix/ipv4-prefix-length | grp-ipv4-
prefix netmask} spt-threshold
- no spt-switchover-threshold {grp-ipv4-prefix/ipv4-prefix-length | grp-ipv4-
prefix netmask}
- ssm-assert-compatible-mode [enable | disable]
- ssm-default-range-disable ipv4
- no ssm-default-range-disable ipv4
- [no] ssm-groups
- [no] group-range {ip-prefix/mask | ip-prefix netmask}

```

2.4.1.4 Operational commands

<GLOBAL>

```

- mrinfo ip-address | dns-name [router router-instance | service-name service-name]
- mtrace source ip-address | dns-name group ip-address | dns-name [destination ip-address
| dns-name] [hop hop] [router router-instance | service-name service-name] [wait-time wait-
time]

```

2.4.1.5 Show commands

```

show
- router
- igmp
- group [grp-ip-address] [host | interface | saps]
- group summary [host | interface | saps]
- interface [ip-int-name | ip-address] [group] [grp-ip-address] [detail]
- ssm-translate interface-name
- static [ip-int-name | ip-addr]
- statistics [ip-int-name | ip-address]
- status

show
- router
- pim
- anycast [family] [detail]
- crp [family | ip-address]
- group [grp-ip-address] [source ip-address] [type {starstarrp | starg | sg}]
[detail] [family]
- interface [ip-int-name | int-ip-address] [group group-ip-address source ip-
address] [type {starstarrp | starg | sg}] [detail] [family]
- mc-ecmp-balance
- neighbor [ip-address | ip-int-name [address neighbor-ip-address]] [detail]
[family]
- rp [family | ip-address]
- rp-hash ip-address
- statistics [ip-int-name | int-ip-address | mpls-if-name] [family]
- status [detail] [family]

```

2.4.1.6 Clear commands

```
clear
```

```
- router
- igmp
- database [group grp-ip-address [source src-ip-address]]
- database interface {ip-int-name | ip-address} [group grp-ip-address [source src-
ip-address]]
- database host ip-address [group grp-ip-address [source src-ip-address]]
- database host all [group grp-ip-address [source src-ip-address]]
- database group-interface all
- statistics group-interface [fwd-service service-id] ip-int-name
- statistics group-interface all
- statistics host ip-address
- statistics host all
- statistics [interface ip-int-name | ip-address]
- version group-interface [fwd-service service-id] ip-int-name
- version group-interface all
- version host ip-address
- version host all
- version [interface ip-int-name | ip-address]
- pim
- database [interface ip-int-name | int-ip-address] [group grp-ip-address [source
ip-address]][family]
- neighbor [interface ip-int-name] [family]
- statistics [[[interface ip-int-name | int-ip-address]]] {[group grp-ip-address
[source ip-address]]}] [family]]
clear
- service
- id
- igmp-snooping
- port-db sap sap-id [group grp-ip-address [source src-ip-address]]
- port-db sdp sdp-id:vc-id [group grp-ip-address [source src-ip-address]]
- querier
- statistics [all | sap sap-id | sdp sdp-id:vc-id]
```

2.4.1.7 Debug commands

```
debug
- router
- igmp
- [no] group-interface [fwd-service service-id] [ip-int-name]
- [no] interface [ip-int-name | ip-address]
- [no] misc
- no packet [query | v1-report | v2-report | v3-report | v2-leave] group-
interface ip-int-name
- no packet [query | v1-report | v2-report | v3-report | v2-leave] host ip-int-name
- packet [query | v1-report | v2-report | v3-report | v2-leave] [ip-int-name | ip-
int-name] [mode {dropped-only | ingr-and-dropped | egr-ingr-and-dropped}]
- packet [query | v1-report | v2-report | v3-report | v2-leave] [mode {dropped-only
| ingr-and-dropped | egr-ingr-and-dropped}] group-interface ip-int-name
- packet [query | v1-report | v2-report | v3-report | v2-leave] host ip-
address [mode {dropped-only | ingr-and-dropped | egr-ingr-and-dropped}]
debug
- router
- pim
- [no] adjacency
- all [group grp-ip-address] [source ip-address] [detail]
- no all
- assert [group grp-ip-address] [source ip-address] [detail]
- no assert
- bgp [source ip-address] [group group-ip-address] [peer peer-ip-address]
- no bgp
- bsr [detail]
```

```
- no bsr
- data [group grp-ip-address] [source ip-address] [detail]
- no data
- db [group grp-ip-address] [source ip-address] [detail]
- no db
- interface [ip-int-name | mt-int-name| ip-address] [detail]
- no interface
- jp [group grp-ip-address] [source ip-address] [detail]
- no jp
- mrib[group grp-ip-address] [source ip-address] [detail]
- no mrib
- msg [detail]
- no msg
- packet [hello | register | register-stop | jp | bsr | assert] [ip-int-name | int-
ip-address]
- no packet
- red [detail]
- no red
- register [group grp-ip-address] [source ip-address] [detail]
- no register
- rtm [detail]
- no rtm
```

2.4.2 Command descriptions

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

2.4.2.1 Configuration commands

2.4.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

config>router>igmp

config>router>igmp>interface

config>router>pim

config>router>pim>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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

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:    config>router>igmp
                config>router>igmp>interface ip-int-name
                config>router>pim
```

2.4.2.1.2 Multicast commands

ssm-translate

Syntax

ssm-translate

Context

```
config>router>igmp>interface
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds or removes SSM translate group ranges.

source

Syntax

[no] **source** *ip-address*

Context

```
config>router>igmp>interface>shutdown>ssm-translate>grp-range
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds or removes source addresses for the SSM translate group range.

Parameters

ip-address

Specifies the unicast source address.

Values a.b.c.d

grp-range

Syntax

[no] **grp-range** *start end*

Context

config>router>igmp>interface>shutdown>ssm-translate

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds or removes SSM translate group range entries.

Parameters

start

Specifies the multicast group range start address.

Values a.b.c.d

end

Specifies the multicast group range end address.

Values a.b.c.d

mc-maximum-routes

Syntax

mc-maximum-routes *number*[log-only][threshold *threshold*]

no mc-maximum-routes

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the maximum number of multicast routes that can be held within a VPN routing/forwarding (VRF) context. When this limit is reached, a log and SNMP trap are sent. If the **log-only** parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, no new joins will be processed.

The **no** form of this command disables the limit of multicast routes within a VRF context. Issue the **no** form of this command only when the VPRN instance is **shutdown**.

Default

no mc-maximum-routes

Parameters

number

Specifies the maximum number of routes to be held in a VRF context.

Values 1 to 2147483647

log-only

Keyword to specify that if the maximum limit is reached, only log the event. This keyword does not disable the learning of new routes.

threshold

Specifies the percentage at which a warning log message and SNMP trap are sent.

Values 0 to 100

2.4.2.1.3 Router IGMP commands

igmp

Syntax

[no] igmp

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the Internet Group Management Protocol (IGMP) context. When the context is created, IGMP is enabled.

IGMP is used by IPv4 systems (hosts and routers) to report their IP multicast group memberships to neighboring multicast routers. An IP multicast router can be a member of one or more multicast groups, in which case it performs both the "multicast router part" of the protocol, which collects the membership information needed by its multicast routing protocol, and the "group member part" of the protocol, which informs it and other neighboring multicast routers of its memberships.

The **no** form of this command disables the IGMP instance. To start or suspend execution of IGMP without affecting the configuration, use the **no shutdown** command.

interface

Syntax

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

Context

config>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure an IGMP interface. The interface is a local identifier of the network interface on which reception of the specified multicast address is to be enabled or disabled.

The **no** form of this command deletes the IGMP interface. The **shutdown** command in the **config>router>igmp>interface** context can be used to disable an interface without removing the configuration for the interface.

Default

no interface

Parameters

ip-int-name

Specifies the IP interface name. Interface names must be unique within the group of defined IP interfaces for the **config>router>interface** and **config>service>ies>interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

If the IP interface name does not exist or does not have an IP address configured, an error message is returned.

If the IP interface exists in a different area, it will be moved to this area.

disable-router-alert-check

Syntax

[no] **disable-router-alert-check**

Context

config>router>igmp>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the router alert checking for IGMP messages received on this interface.

The **no** form of this command disables the IGMP router alert check option.

import

Syntax

import *policy-name*

no import

Context

configure>router>igmp>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command applies the referenced IGMP policy (filter) to an interface subscriber or a group interface. An IGMP filter is also known as an allowlist/denylist and it is defined under the **configure>router>policy-options** context.

The **no** form of this command removes the policy association from the IGMP instance.

Default

no import

Parameters

policy-name

Specifies the route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#,

\$, spaces, and so on.), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policyoptions** context.

max-groups

Syntax

max-groups [*value*]

no max-groups

Context

config>router>igmp>if

config>router>pim>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the maximum number of groups for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed. When the value is 0, there is no limit to the number of groups. This command is applicable for IPv4 only.

Default

max-groups 0

Parameters

value

Specifies the maximum number of groups for this interface.

Values	1 to 900 (for the 7210 SAS-K 2F6C4T)
	1 to 950 (for the 7210 SAS-K 3SFP+ 8C)

max-sources

Syntax

max-sources [*value*]

no max-sources

Context

config>router>igmp>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum number of group sources for this interface

Parameters

value

Specifies the maximum number of group sources that can be configured.

Values 1 to 1000

query-last-listener-interval

Syntax

query-last-listener-interval *seconds*

no query-last-listener-interval

Context

config>router>igmp>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the frequency at which the querier sends group-specific query messages, including messages sent in response to leave-group messages. The lower the interval, the faster the detection of the loss of the last member of a group.

Default

no query-last-listener-interval

Parameters

seconds

Specifies the frequency, in seconds, at which the router transmits group-specific host-query messages.

Values 1 to 1023

static

Syntax

static

Context

config>router>igmp>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command tests multicast forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without receiving membership reports from host members.

group

Syntax

[no] **group** *grp-ip-address*

Context

config>router>igmp>if>static

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds a static multicast group either as a (*,G) or one or more (S,G) records. Use IGMP static group memberships to test multicast forwarding without a receiver host. When IGMP static groups are enabled, data is forwarded to an interface without receiving membership reports from host members.

When static IGMP group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static IGMP group entries do not generate join messages toward the RP.

Parameters

grp-ip-address

Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

Values a.b.c.d

source

Syntax

[no] source *ip-address*

Context

config>router>igmp>if>static>group

config>router>igmp>ssm-translate>grp-range

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies an IPv4 unicast address that sends data on an interface. This enables a multicast receiver host to signal to a router the group from which to receive multicast traffic, and the sources from which the traffic is expected.

The **source** command is mutually exclusive with the specification of individual sources for the same group.

The source command in combination with the group is used to create a specific (S,G) static group entry.

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

Parameters

ip-address

Specifies the IPv4 unicast address.

Values a.b.c.d

starg

Syntax

[no] starg

Context

config>router>igmp>if>static>group

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.

The **no** form of this command is used to remove the starg entry from the configuration.

subnet-check

Syntax

[no] subnet-check

Context

config>router>igmp>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables subnet checking for IGMP messages received on this interface. All IGMP packets with a source address that is not in the local subnet are dropped.

The **no** form of this command disables subnet checking.

Default

subnet-check

version

Syntax

version *version*

no version

Context

config>router>igmp>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the IGMP version. If routers run different versions of IGMP, they will negotiate the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version. For IGMP to function correctly, all routers on a LAN should be configured to run the same version of IGMP on that LAN.

For IGMPv3, a multicast router that is also a group member performs both parts of IGMPv3, receiving and responding to its own IGMP message transmissions as well as those of its neighbors.

Default

version 3

Parameters

version

Specifies the IGMP version number.

Values 1, 2, 3

query-interval

Syntax

query-interval *seconds*

no query-interval

Context

config>router>igmp

config>router>igmp>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the frequency at which the querier router transmits general host-query messages. The host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.

Default

query-interval 125

Parameters

seconds

Specifies the time frequency, in seconds, that the router transmits general host-query messages.

Values 2 to 1024

query-last-member-interval

Syntax

query-last-member-interval *seconds*

Context

config>router>igmp

config>router>igmp>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the frequency at which the querier sends group-specific query messages, including messages sent in response to leave-group messages. The lower the interval, the faster the detection of the loss of the last member of a group.

Default

query-last-member-interval 1

Parameters

seconds

Specifies the frequency, in seconds, at which query messages are sent.

Values 1 to 1024

query-response-interval

Syntax

query-response-interval *seconds*

Context

config>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies how long the querier router waits to receive a response to a host-query message from a host.

Default

query-response-interval 10

Parameters

seconds

Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 to 1023

robust-count

Syntax

robust-count *robust-count*

no robust-count

Context

config>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the robust count. The *robust-count* variable allows tuning for the expected packet loss on a subnet. If a subnet anticipates losses, the robust-count variable can be increased.

Default

robust-count 2

Parameters

robust-count

Specifies the robust count value.

Values 2 to 10

ssm-translate

Syntax

ssm-translate

Context

config>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure group ranges, which are translated to source-specific multicast (SSM) (S,G) entries. If the static entry needs to be created, it has to be translated from an IGMPv1 or IGMPv2 request to an SSM join. An SSM translate source can only be added if the **starg** command is not enabled. An error message is generated if you try to configure the **source** command with the **starg** command enabled.

grp-range

Syntax

[no] **grp-range** *start end*

Context

config>router>igmp>ssm-translate

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command is used to configure group ranges, which are translated to SSM (S,G) entries.

Parameters

start

Specifies an IP address that indicates the start of the group range.

Values a.b.c.d

end

Specifies an IP address that indicates the end of the group range. This value should always be greater than or equal to the value of the *start* value.

Values a.b.c.d

source

Syntax

[no] **source** *ip-address*

Context

config>router>igmp>ssm-translate>grp-range

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the source IP address for the group range. Whenever a (*,G) report is received in the range specified by **grp-range** *start* and *end* parameters, it is translated to an (S,G) report with the value of this object as the source address.

Parameters

ip-address

Specifies the IP address that will be sending data.

Values a.b.c.d

2.4.2.1.4 Router PIM commands

```
pim
```

Syntax

[no] pim

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a protocol independent multicast (PIM) instance.

PIM is used for multicast routing within the network. Devices in the network can receive the multicast feed requested and non-participating routers can be pruned. The router OS supports PIM sparse mode (PIM-SM).

Default

no pim

```
enable-mdt-spt
```

Syntax

[no] enable-mdt-spt

Context

config>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command is used to enable SPT switchover for default MDT.

The **no** form of this command disables SPT switchover for default MDT. If disabled, the PIM instance resets all MDTs and reinitiates setup.

Default

no enable-mdt-spt

import

Syntax

import {**join-policy** | **register-policy**} [*policy-name* [*.. policy-name*]]

no import {**join-policy** | **register-policy**}

Context

config>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the import route policy to be used. Route policies are configured in the **config>router>policy-options** context.

When an import policy is not specified, BGP routes are accepted by default. Up to five import policy names can be specified.

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

Default

no import join-policy

no import register-policy

Parameters

join-policy

Keyword to filter PIM join messages, which prevents unwanted multicast streams from traversing the network.

register-policy

Keyword to filter register messages. PIM register filters prevent register messages from being processed by the RP. This filter can only be defined on an RP. When a match is found, the RP immediately sends back a register-stop message.

policy-name

Specifies the route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

interface

Syntax

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

Context

config>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a logical IP routing interface.

Interface names are case-sensitive and must be unique within the group of IP interfaces defined for **config>router>interface** and **config>service>ies>interface**. Interface names must not be in the dotted-decimal notation of an IP address. For example, the name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed. Show commands for router interfaces use either the interface names or the IP addresses. Ambiguity can exist if an IP address is used as an IP address and an interface name. Duplicate interface names can exist in different router instances, although this is not recommended because it is confusing.

The **no** form of this command removes the IP interface and all the associated configurations.

Parameters

ip-int-name

Specifies the name of the IP interface, up to 32 characters. Interface names must be unique within the group of defined IP interfaces for the **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, and so on.), the entire string must be enclosed within double quotes.

If the *ip-int-name* already exists, the context is changed to maintain that IP interface. If *ip-int-name* does not exist, the interface is created and the context is changed to that interface for further command processing.

assert-period

Syntax

assert-period *assert-period*

no assert-period

Context

config>router>pim>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the period for periodic refreshes of PIM Assert messages on an interface.

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

Default

no assert-period

Parameters

assert-period

Specifies the period for periodic refreshes of PIM Assert messages on an interface.

Values 1 to 300 seconds

bfd-enable

Syntax

[no] bfd-enable [ipv4]

Context

config>router>pim>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the use of IPv4 bidirectional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a specific protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set using the BFD command under the IP interface.

The **no** form of this command removes BFD from the associated IGP protocol adjacency.

For information about the protocols and platforms that support BFD, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*.

Default

no bfd-enable

bsm-check-rtr-alert

Syntax

[no] **bsm-check-rtr-alert**

Context

config>router>pim>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the checking of the router alert option in the bootstrap messages received on this interface.

The **no** form of this command enables accepting of BSM packets without the router alert option.

Default

no bsm-check-rtr-alert

mc-ecmp-balance

Syntax

[no] **mc-ecmp-balance**

Context

configure>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables multicast balancing of traffic over ECMP links. When this command is enabled, each multicast stream that needs to be forwarded over an ECMP link is reevaluated for the total multicast bandwidth utilization. Reevaluation occurs on the ECMP interface in question.

The **no** form of this command disables multicast balancing.

mc-ecmp-balance-hold

Syntax

mc-ecmp-balance-hold *minutes*

no mc-ecmp-balance-hold

Context

configure>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the hold time for multicast balancing over ECMP links.

Parameters

minutes

Specifies the hold time, in minutes, that applies after an interface has been added to the ECMP link.

Values 2 to 600

mc-ecmp-hashing-enabled

Syntax

[no] mc-ecmp-hashing-enabled

Context

configure>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables hash-based multicast balancing of traffic over ECMP links and causes PIM joins to be distributed over the multiple ECMP paths based on a hash of S and G (and possibly next-hop IP). When a link in the ECMP set is removed, the multicast streams that were using that link are redistributed over the remaining ECMP links using the same hash algorithm. When a link is added to the ECMP set, new joins may be allocated to the new link based on the hash algorithm, but existing multicast streams using the other ECMP links stay on those links until they are pruned.

Hash-based multicast balancing is supported for IPv4 only.

This command is mutually exclusive with the **mc-ecmp-balance** command in the same context.

The **no** form of this command disables the hash-based multicast balancing of traffic over ECMP links.

Default

no mc-ecmp-hashing-enabled

hello-interval

Syntax

hello-interval *hello-interval*
no hello-interval

Context

config>router>pim>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the frequency at which PIM hello messages are transmitted on this interface.
The **no** form of this command reverts to the default value.

Default

hello-interval 30

Parameters

hello-interval

Specifies the hello interval in seconds. A 0 (zero) value disables the sending of hello messages (the PIM neighbor will never timeout the adjacency).

Values 0 to 255 seconds

hello-multiplier

Syntax

hello-multiplier *deci-units*
no hello-multiplier

Context

config>router>pim>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the multiplier to determine the hold time for a PIM neighbor on this interface.
The **hello-multiplier** in conjunction with the **hello-interval** determines the hold time for a PIM neighbor.

Parameters

deci-units

Specifies the value, in multiples of 0.1, for the formula used to calculate the hello-hold time based on the hello-multiplier:

$(\text{hello-interval} * \text{hello-multiplier}) / 10$

This allows the PIMv2 default timeout of 3.5 seconds to be supported.

Values 20 to 100

Default 35

improved-assert

Syntax

[no] improved-assert

Context

config>router>pim>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

The PIM assert process establishes a forwarder for a LAN and requires interaction between the control and forwarding planes. The assert process is started when data is received on an outgoing interface, meaning that duplicate traffic is forwarded to the LAN until the forwarder is negotiated among the routers.

When the **improved-assert** command is enabled, the PIM assert process is done entirely in the control plane. The advantages are that it eliminates duplicate traffic forwarding to the LAN. It also improves performance because it removes the required interaction between the control and data planes.



Note:

The **improved-assert** command is still fully interoperable with the draft-ietf-pim-sm-v2-new-xx, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Revised*, and RFC 2362, *Protocol Independent Multicast-Sparse Mode (PIM-SM)*, implementations. However, there may be conformance tests that may fail if the tests expect control-data plane interaction in determining the assert winner. Nokia recommends disabling the **improved-assert** command when performing conformance tests.

Default

enabled

instant-prune-echo

Syntax

[no] **instant-prune-echo**

Context

config>router>pim>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables a PIM router to echo the PIM prune message received from a downstream router. It is typically used in a multi-access broadcast network (For example: Ethernet LAN) to reduce the probability of loss of PIM prune messages.

Default

no instant-prune-echo

multicast-senders

Syntax

multicast-senders {auto | always | never}

no multicast-senders

Context

config>router>pim>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures how traffic from directly-attached multicast sources should be treated on broadcast interfaces. It can also be used to treat all traffic received on an interface as traffic coming from a directly-attached multicast source. This is particularly useful if a multicast source is connected to a point-to-point or unnumbered interface.

Default

auto

Parameters

auto

Specifies that, on broadcast interfaces, the forwarding plane performs subnet-match checks on multicast packets received on the interface to determine whether the packet is from a directly-attached source. On unnumbered/point-to-point interfaces, all traffic is implicitly treated as coming from a remote source.

always

Specifies that all traffic received on the interface be treated as coming from a directly-attached multicast source.

never

Specifies that, on broadcast interfaces, traffic from directly-attached multicast sources is not forwarded. Traffic from a remote source is still forwarded if there is a multicast state for it. On unnumbered/point-to-point interfaces, all traffic received on that interface must not be forwarded.

priority

Syntax

priority *dr-priority*

no priority

Context

config>router>pim>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets the priority value to elect the designated router (DR). The DR election priority is a 32-bit unsigned number and the numerically larger priority is always preferred.

The **no** form of this command restores the default values.

Default

priority 1

Parameters

priority

Specifies the priority to become the designated router. The higher the value, the higher the priority.

Values 1 to 4294967295

sticky-dr

Syntax

sticky-dr [*priority dr-priority*]
no sticky-dr

Context

config>router>pim>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the **sticky-dr** operation on this interface. When the operation is enabled, the priority in PIM hello messages sent on this interface when elected as the designated router (DR) are modified to the value configured in *dr-priority*. This is done to avoid the delays in forwarding caused by DR recovery, when switching back to the old DR on a LAN when it comes back up.

By enabling **sticky-dr** on an interface, it will continue to act as the DR for the LAN even after the old DR comes back up.

The **no** form of this command disables the **sticky-dr** operation on this interface.

Default

no sticky-dr

Parameters

dr-priority

Specifies the DR priority to be sent in PIM Hello messages following the election of that interface as the DR when **sticky-dr** operation is enabled.

Values 1 to 4294967295

three-way-hello

Syntax

three-way-hello [*compatibility-mode*]
no three-way-hello

Context

config>router>pim>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the compatibility mode to enable three-way hello. By default, the value is disabled on all interfaces, which specifies that the standard two-way hello is supported. When enabled, the three-way hello is supported.

Default

no three-way-hello

tracking-support

Syntax

[no] tracking-support

Context

config>router>pim>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets the T bit in the LAN prune delay option of the hello message. This indicates that the router is capable of enabling join message suppression. This capability allows for upstream routers to explicitly track join membership.

Default

no tracking-support

rp

Syntax

rp

Context

config>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure rendezvous point (RP) parameters. The address of the root of the group shared multicast distribution tree is known as its RP. Packets received from a source upstream and join messages from downstream routers rendezvous at this router.

If this command is disabled, the router cannot become the RP.

anycast

Syntax

[no] **anycast** *rp-ip-address*

Context

config>router>pim>rp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a PIM anycast protocol instance for the RP being configured. Anycast enables fast convergence when a PIM RP router fails by allowing receivers and sources to rendezvous at the closest RP.

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

Parameters

rp-ip-address

Specifies the loopback IP address shared by all routes that form the RP set for this anycast instance. Only a single address can be configured. If another anycast command is entered with an address, the old address is replaced with the new address. If no IP address is entered, the command is used to enter the anycast CLI level.

Values a.b.c.d

rp-set-peer

Syntax

[no] **rp-set-peer** *ip-address*

Context

config>router>pim>rp>anycast

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a peer in the anycast RP set. The address identifies the address used by the other node as the RP candidate address for the same multicast group address range as configured on this node.

This is a manual procedure. Caution should be taken to produce a consistent configuration of an RP set for a specific multicast group address range. The priority should be identical on each node and be a higher value than any other configured RP candidate that is not a member of this RP set.

Although there is no set maximum number of addresses that can be configured in an RP set, up to 15 IP addresses is recommended.

The **no** form of this command removes an entry from the list.

Parameters

ip-address

Specifies a peer in the anycast RP set.

Values a.b.c.d

bootstrap-export

Syntax

bootstrap-export *policy-name*[*policy-name*...(up to 5 max)]

no bootstrap-export

Context

config>router>pim>rp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command applies export policies to control the flow of bootstrap messages from the RP and apply them to the PIM configuration. Up to five policy names can be specified.

Default

no bootstrap-export

Parameters

policy-name

Specifies the export policy name, up to 32 characters.

bootstrap-import

Syntax

bootstrap-import *policy-name*[..*policy-name*...(5 maximum)]

no bootstrap-import

Context

config>router>pim>rp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command applies import policies to control the flow of bootstrap messages to the RP, and apply them to the PIM configuration. Up to 5 policy names can be specified.

Default

no bootstrap-import

Parameters

policy-name

Specifies the import policy name, up to 32 characters.

bsr-candidate

Syntax

bsr-candidate

Context

config>router>pim>rp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure Candidate Bootstrap (BSR) parameters.

address

Syntax

address *ip-address*

Context

config>router>pim>rp>bsr-candidate

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the candidate BSR IP address. This address is for bootstrap router election.

Parameters

ip-address

Specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

hash-mask-len

Syntax

hash-mask-len *hash-mask-length*

no hash-mask-len

Context

config>router>pim>rp>bsr-candidate

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command is used to configure the length of a mask that is to be combined with the group address before the hash function is called. All groups with the same hash map to the same RP. For example, if the *hash-mask-length* value is 24, only the first 24 bits of the group addresses matter. This mechanism is used to map one group or multiple groups to an RP.

Default

hash-mask-len 30

Parameters

hash-mask-length

Specifies the hash mask length.

Values 0 to 32 (v4)

priority

Syntax

priority *bootstrap-priority*

no priority

Context

config>router>pim>rp>bsr-candidate

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the bootstrap priority of the router. The RP is sometimes called the bootstrap router. The priority determines if the router is eligible to be a bootstrap router. In the case of a tie, the router with the highest IP address is elected to be the bootstrap router.

Default

priority 0

Parameters

bootstrap-priority

Specifies the priority to become the bootstrap router. The higher the value, the higher the priority. A 0 value means the router is not eligible to be the bootstrap router. A value of 1 means the router is the least likely to become the designated router.

Values 0 to 255

rp-candidate

Syntax

rp-candidate

Context

config>router>pim>rp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure the candidate RP parameters.

Routers use a set of available rendezvous points distributed in bootstrap messages to get the proper group-to-RP mapping. A set of routers within a domain are also configured as candidate RPs (C-RPs); typically, these will be the same routers that are configured as candidate BSRs.

Every multicast group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The rendezvous point (RP) is the root of this shared tree.

Default

rp-candidate shutdown

address

Syntax

[no] **address** *ip-address*

Context

config>router>pim>rp>rp-candidate

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the local RP address. This address is sent in the RP candidate advertisements to the bootstrap router.

Parameters

ip-address

Specifies the IP address.

Values a.b.c.d

group-range

Syntax

[no] **group-range** {*grp-ip-address/mask*| *grp-ip-address netmask*}

Context

config>router>pim>rp>rp-candidate

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the address ranges of the multicast groups for which this router can be an RP.

Parameters

grp-ip-address

Specifies the multicast group IP address expressed in dotted-decimal notation.

Values a.b.c.d (multicast group address)

mask

Specifies the mask associated with the IP prefix expressed as a mask length or in dotted-decimal notation; for example /16 for a sixteen-bit mask. The mask can also be entered in dotted-decimal notation.

Values 4 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)

holdtime

Syntax

holdtime *holdtime*

no holdtime

Context

config>router>pim>rp>rp-candidate

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the length of time, in seconds, that neighbors should consider the sending router to be operationally up. A local RP cannot be configured on a logical router.

Parameters

holdtime

Specifies the hold time, in seconds.

Values 5 to 255

priority

Syntax

priority *priority*
no priority

Context

config>router>pim>rp>rp-candidate

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the candidate-RP priority for becoming a rendezvous point (RP). This value is used to elect an RP for a group range.

Default

priority 192

Parameters

priority

Specifies the priority to become a rendezvous point (RP). A value of 0 is considered as the highest priority.

Values 0 to 255

static

Syntax

static

Context

config>router>pim>rp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure static rendezvous point (RP) addresses for a multicast group range.

Entries can be created or destroyed. If no IP addresses are configured in the

config>router>pim>rp>static>address context, the multicast group-to-RP mapping is derived from the RP-set messages received from the bootstrap router.

address

Syntax

address *ip-address*

no address

Context

config>router>pim>rp>static

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command indicates the RP address that should be used by the router for the range of multicast groups configured by the **group-range** command.

Parameters

ip-address

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

Values a.b.c.d

group-range

Syntax

[no] group-range {*ip-prefix/mask* | *ip-prefix netmask*}

Context

config>router>pim>ssm-groups

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the SSM multicast group address ranges for this router.

Parameters

ip-prefix/mask

Specifies the IP prefix in dotted-decimal notation and the associated mask.

Values

ipv4-prefix: a.b.c.d

ipv4-prefix-le: 0 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)

group-prefix

Syntax

[no] **group-prefix** {*grp-ip-address/mask* | *grp-ip-address netmask*}

Context

config>router>pim>rp>static>address

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the range of multicast group addresses that should be used by the router as the RP. The **config router pim rp static address** command implicitly defaults to deny all for all multicast groups (224.0.0.0/4). A group-prefix must be specified for that static address. This command does not apply to the whole group range.

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

Parameters

grp-ip-address

Specifies the multicast group IP address expressed in dotted-decimal notation.

Values a.b.c.d

mask

Specifies the mask associated with the IP prefix expressed as a mask length or in dotted-decimal notation; for example /16 for a sixteen-bit mask. The mask can also be entered in dotted-decimal notation.

Values 4 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)

override

Syntax

[no] override

Context

config>router>pim>rp>static>address

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command changes the precedence of static RP over dynamically learned RP.

When this command is enabled, the static group-to-RP mappings take precedence over the dynamically learned mappings.

Default

no override

non-dr-attract-traffic

Syntax

[no] non-dr-attract-traffic

Context

config>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies whether the router should ignore the designated router state and attract traffic even when it is not the designated router.

An operator can configure an interface (router, IES, or VPRN interfaces) to IGMP and PIM. The interface state is synchronized to the backup node if it is associated with the redundant peer port. The interface can be configured to use PIM, which causes multicast streams to be sent to the elected DR only. The DR is also the router sending traffic to the DSLAM. Because it may be required to attract traffic to both routers, the **non-dr-attract-traffic** flag can be used in the PIM context to have the router ignore the DR state and attract traffic when not DR. While using this flag, the router may not send the stream down to the DSLAM while not DR.

When this command is enabled, the designated router state is ignored.

The **no** form of this command causes the router to honor the designated router value.

Default

no non-dr-attract-traffic

rpf-table

Syntax

rpf-table {*rtable-u*}

no rpf-table

Context

config>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the sequence of route tables used to find a Reverse Path Forwarding (RPF) interface for a particular multicast route.

By default, only the unicast route table is looked up to calculate the RPF interface toward the source/rendezvous point. However, the operator can specify the use of the unicast route table (*rtable-u*).

Default

rpf-table *rtable-u*

Parameters

rtable-u

Specifies only that the unicast route table will be used by the multicast protocol (PIM) for IPv4 RPF checks. This route table will contain routes submitted by all the unicast routing protocols.

spt-switchover-threshold

Syntax

spt-switchover-threshold {*grp-ipv4-prefix/ipv4-prefix-length* | *grp-ipv4-prefix netmask*} *spt-threshold*

Context

config>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the shortest path tree (SPT) switchover thresholds for group prefixes.

PIM-SM routers with directly connected routers receive multicast traffic initially on a shared tree rooted at the RP. When the traffic arrives on the shared tree and the source of the traffic is known, a switchover to the SPT tree rooted at the source is attempted.

For a group that falls in the range of a prefix configured in the table, the corresponding threshold value determines when the router should switch over from the shared tree to the source-specific tree. The switchover is attempted only if the traffic rate on the shared tree for the group exceeds the configured threshold.



Note:

On the 7210 SAS, this command is used to enable or disable switch over to the SPT tree. To disable switch over to SPT, a threshold value of infinity must be configured (that is, to continue using the shared tree forever, configure the IP multicast prefix with this command and set the threshold to infinity). To use the SPT tree, do not configure the IP multicast address prefix using this command and the default behavior will apply to the multicast group. The default behavior is to switch over to SPT when the first packet is received.

In the absence of any matching prefix in the table, the default behavior is to switch over when the first packet is seen. In the presence of multiple prefixes matching a specific group, the most specific entry is used.

Parameters

grp-ipv4-prefix

Specifies the multicast group IP address expressed in dotted-decimal notation.

Values a.b.c.d (multicast IP address)

ipv4-prefix-length

Specifies the length of the IPv4 prefix.

Values 4 to 32

netmask

Specifies the netmask associated with the IPv4 prefix expressed in dotted-decimal notation.

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

spt-threshold

Specifies the configured threshold in kilobits per second (kbps) for a group prefix. A switchover is attempted only if the traffic rate on the shared tree for the group exceeds this configured threshold.

Values 1, infinity

infinity

Keyword to specify that no switchover will occur at any time, regardless of the traffic level is detected. The threshold value, in kilobits per second (KBPS), is 4294967295.

ssm-assert-compatible-mode

Syntax

ssm-assert-compatible-mode [enable|disable]

Context

config>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

When this command is enabled, packets are treated as if SPT bit was set regardless of whether it is set or not.

Default

ssm-assert-compatible-mode disable

Parameters

enable

Enables SSM assert in compatibility mode for this PIM protocol instance.

disable

Disables SSM assert in compatibility mode for this PIM protocol instance.

ssm-default-range-disable

Syntax

[no] **ssm-default-range-disable** ipv4

Context

config>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command allows the user to disable the reservation and allows PIM to accept and create (*,G) entries for addresses in this range on receiving IGMPv2 reports. PIM SSM has a default range of 232/8 (232.0.0.0 to 232.255.255.255) reserved by IANA. These addresses are not used by PIM ASM.

Default

ssm-default-range-disable ipv4

ssm-groups

Syntax

[no] ssm-groups

Context

config>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure SSM group ranges.

2.4.2.1.5 Operational commands

mrinfo

Syntax

mrinfo *ip-address* | *dns-name* [**router** *router-instance* | **service-name** *service-name*]

Context

<GLOBAL>

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays relevant multicast information from the target multicast router. Information displayed includes adjacency information, protocol, metrics, thresholds, and flags from the target multicast router. This information can be used by network operators to determine whether bidirectional adjacencies exist.

Parameters

ip-address

Specifies the IP address of the multicast capable target router.

Values ip-address ipv4 unicast address (a.b.c.d)

dns-name

Specifies the DNS name, up to 63 characters.

router-instance

Specifies the router instance.

Default router-name - "Base" | "management" Default - Base

service-name

Specifies the service name, up to 64 characters

Output

The following output is an example of multicast information, and [Table 7: Output fields: mrinfo](#) describes the output fields.

Sample output

```
A:dut-f# mrinfo 10.1.1.2

10.1.1.2 [version 3.0,prune,genid,mtrace]:
 10.1.1.2 -> 10.1.1.1 [1/0/pim]
 16.1.1.1 -> 0.0.0.0 [1/0/pim/down/disabled]
 17.1.1.1 -> 0.0.0.0 [1/0/pim/querier/leaf]
 200.200.200.3 -> 200.200.200.5 [1/0/tunnel/pim]...

A:dut-g# mrinfo 1.1.1.1

1.1.1.1 [version 7.0,prune,genid,mtrace]:
? 1.1.1.1 -> ? 0.0.0.0 [1/0/pim/leaf]
? 12.1.1.1 -> ? 12.1.1.2 [1/0/pim]
? 19.1.1.1 -> ? 19.1.1.9 [1/0/pim]
? 11.1.1.1 -> ? 0.0.0.0 [1/0/pim/leaf]
? 17.1.1.1 -> ? 17.1.1.7 [1/0/pim]
? 17.1.2.1 -> ? 17.1.2.7 [1/0/pim]
```

Table 7: Output fields: mrinfo

Label	Description
General flags	
version	Displays the software version on queried router
prune	Indicates that router understands pruning
genid	Indicates that router sends generation IDs
mtrace	Indicates that the router handles mtrace requests
Neighbors flags	
1	Metric
0	Threshold (multicast time-to-live)
pim	PIM enabled on interface

Label	Description
down	Operational status of interface
disabled	Administrative status of interface
leaf	No downstream neighbors on interface
querier	Interface is IGMP querier
tunnel	Neighbor reached via tunnel

mtrace

Syntax

mtrace **source** *ip-address* | *dns-name* [**group** *ip-address* | *dns-name*] [**destination** *ip-address* | *dns-name*] [**hop** *hop*] [**router** *router-instance* | **service-name** *service-name*] [**wait-time** *wait-time*]

Context

<GLOBAL>

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command traces the multicast path from a source to a receiver by passing a trace query hop-by-hop along the reverse path from the receiver to the source. At each hop, information such as the hop address, routing error conditions, and packet statistics are gathered and returned to the requester. A network administrator can determine where multicast flows stop and verify the flow of the multicast stream.

Parameters

source *ip-address*

Specifies the IP address of the multicast-capable source. This is a unicast address of the beginning of the path to be traced.

dns-name

Specifies the DNS name, up to 63 characters.

Values *ip-address* ipv4 unicast address (a.b.c.d)

group *ip-address*

Specifies the multicast address.

destination *ip-address*

Specifies the IP address of the unicast destination. If this parameter is omitted, the IP address of the system where the command is entered is used. The destination parameter can also be used to specify a local interface address as the destination address to send the trace query.

Default The default address for the destination address is the incoming IETF format for that (S,G)

hop

Specifies the maximum number of hops that will be traced from the receiver back toward the source.

Values 1 to 255

Default 32 hops (infinity for the DVMRP routing protocol).

router-instance

Specifies the router name or service ID used to identify the router instance.

Default router-name - "Base" | "management" Default - Base

service-name

Specifies the service name, up to 64 characters.

wait-time

Specifies the number of seconds to wait for the response.

Values 1 to 60

Default 10

Output

The following output is an example of mtrace information, and [Table 8: Output fields: mtrace](#) describes the output fields.

Sample output

```
A:Dut-F# mtrace source 10.10.16.9 group 224.5.6.7

Mtrace from 10.10.16.9 via group 224.5.6.7
Querying full reverse path...

 0 ? (10.10.10.6)
-1 ? (10.10.10.5) PIM thresh^ 1 No Error
-2 ? (10.10.6.4) PIM thresh^ 1 No Error
-3 ? (10.10.4.2) PIM thresh^ 1 Reached RP/Core
-4 ? (10.10.1.1) PIM thresh^ 1 No Error
-5 ? (10.10.2.3) PIM thresh^ 1 No Error
-6 ? (10.10.16.9)
Round trip time 29 ms; total ttl of 5 required.
```

Table 8: Output fields: mtrace

Label	Description
hop	Displays the number of hops from the source to the listed router

Label	Description
router name	Displays the name of the router for this hop. If a DNS name query is not successful a "?" displays
address	Displays the address of the router for this hop
protocol	Displays the protocol used
ttl	Displays the forward TTL threshold. TTL that a packet is required to have before it will be forwarded over the outgoing interface
forwarding code	Displays the forwarding information or error code for this hop

2.4.2.2 Show commands

2.4.2.2.1 IGMP commands

group

Syntax

group [*grp-ip-address*] [**host** | **interface** | **saps**]

group summary [**host** | **interface** | **saps**]

Context

show>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the multicast group and (S,G) addresses. If no *grp-ip-address* parameters are specified, all IGMP group, (*,G) and (S,G) addresses are displayed.

Parameters

grp-ip-address

Displays specific multicast group addresses.

host

Displays hosts for the multicast group addresses.

interface

Displays interfaces for the multicast group addresses.

saps

Displays SAPs for the multicast group addresses.

Output

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

Sample output

```
*B:Dut-C# show router igmp group
=====
IGMP Interface Groups
=====
IGMP Host Groups
=====
(*,225.0.0.1)
  Fwd List : 112.112.1.2          Up Time : 0d 00:00:21
(11.11.0.1,225.0.0.1)
  Fwd List : 112.112.1.1          Up Time : 0d 00:00:30
  Blk List : 112.112.1.2          Up Time : 0d 00:00:21
(11.11.0.2,225.0.0.1)
  Fwd List : 112.112.1.1          Up Time : 0d 00:00:30
(*,225.0.0.2)
  Fwd List : 112.112.1.2          Up Time : 0d 00:00:21
(11.11.0.1,225.0.0.2)
  Blk List : 112.112.1.2          Up Time : 0d 00:00:21
-----
(*,G)/(S,G) Entries : 5
=====
*B:Dut-C#

*B:Dut-C# show router igmp group summary
=====
IGMP Interface Groups
=====
IGMP Host Groups Summary          Nbr Fwd   Nbr Blk
=====
(*,225.0.0.1)                     1         0
(11.11.0.1,225.0.0.1)             1         1
(11.11.0.2,225.0.0.1)             1         0
(*,225.0.0.2)                     1         0
(11.11.0.1,225.0.0.2)             0         1
-----
(*,G)/(S,G) Entries : 5
=====
*B:Dut-C#

A:NYC# show router igmp group 224.24.24.24
=====
IGMP Groups
=====
(*,224.24.24.24)                  Up Time : 0d 05:23:23
  Fwd List : nyc-vlc
-----
(*,G)/(S,G) Entries : 1
=====
A:NYC#
```

Table 9: Output fields: IGMP group

Label	Description
IGMP Groups	Displays the IP multicast sources corresponding to the IP multicast groups which are statically configured
Fwd List	Displays the list of interfaces in the forward list

interface

Syntax

interface [*ip-int-name* | *ip-address*] [*group*] [*grp-address*] [*detail*]

Context

show>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays IGMP interface information.

Parameters

ip-int-name

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

ip-address

Displays the information associated with the specified IP address.

Values a.b.c.d

grp-address

Displays IP multicast group address for which this entry contains information.

Values a.b.c.d, multicast group address or 0

detail

Displays detailed IP interface information along with the source group information learned on that interface.

Output

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

Sample output

```
A:Dut-C# show router igmp interface
=====
IGMP Interfaces
=====
Interface           Adm  Oper Querier           Cfg/Opr Num      Policy
                    Version Groups
-----
C_Rx_net1           Up   Up   10.2.1.3           3/3    900    none
C_Rx_acc1           Up   Up   10.1.1.3           3/3    900    none
C_Rx_acc2           Up   Up   10.1.2.3           3/3    900    none
C_Rx_net2           Up   Up   10.2.2.3           3/3    900    none
-----
Interfaces : 4
=====
```

```
*A:Dut-C# show router igmp interface detail
=====
IGMP Interface C_Rx_net1
=====
Interface           : C_Rx_net1
Admin Status        : Up
Querier             : 10.2.1.3
Querier Expiry Time: N/A
Admin/Oper version  : 3/3
Policy              : none
Max Groups Allowed  : No Limit
Use LAG port weight: no
Router Alert Check  : Enabled
Redundant Multicast: no

Oper Status         : Up
Querier Up Time     : 0d 00:00:55
Time for next query: 0d 00:01:51
Num Groups          : 900
Subnet Check        : Enabled
Max Groups Till Now: 900
Max Sources Allowed: No Limit
Max GrpSrcs Allowed: No Limit
Red. Multicast Fwd  : N/A

IGMP Group
-----
Group Address : 239.225.1.1
Interface     : C_Rx_net1
Last Reporter : 10.2.1.1
V1 Host Timer : Not running
V2 Host Timer : Not running
Up Time       : 0d 00:00:51
Expires       : N/A
Mode          : include
Type          : dynamic
Compat Mode   : IGMP Version 3

Source Address Expires      Type      Fwd/Blk
-----
10.1.1.2       0d 00:04:07  dynamic Fwd
-----
```

Table 10: Output fields: IGMP interface

Label	Description
Interface	Specifies the interfaces that participate in the IGMP protocol
Adm Admin Status	Displays the administrative state for the IGMP protocol on this interface
Oper Oper Status	Displays the current operational state of IGMP protocol on the interface

Label	Description
Querier	Displays the address of the IGMP querier on the IP subnet to which the interface is attached
Querier Up Time	Displays the time since the querier was last elected as querier
Querier Expiry Timer	Displays the time remaining before the querier ages out. If the querier is the local interface address, the value will be zero.
Cfg/Opr Version Admin/Oper version	Cfg — The configured version of IGMP running on this interface. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN. Opr — The operational version of IGMP running on this interface. If the cfg value is 3 but all of the routers in the local subnet of this interface use IGMP version v1 or v2, the operational version will be v1 or v2.
Num Groups	Displays the number of multicast groups which have been learned by the router on the interface
Policy	Displays the policy that is to be applied on the interface
Group Address	Displays the IP multicast group address for which this entry contains information
Up Time	Displays the time since this source group entry got created
Last Reporter	Displays the IP address of the source of the last membership report received for this IP Multicast group address on this interface. If no membership report has been received, this object has the value 0.0.0.0.
Mode	The mode is based on the type of membership report(s) received on the interface for the group. In the 'include' mode, reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the source-list parameter of the IGMP membership report. In 'exclude' mode, reception of packets sent to the specific multicast address is requested from all IP source addresses except those listed in the source-list parameter.
V1 Host Timer	The time remaining until the local router will assume that there are no longer any IGMP version 1 members on the IP subnet attached to this interface. Upon hearing any IGMPv1 Membership Report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 Leave messages for this group that it receives on this interface.
V2 Host Timer	The time remaining until the local router will assume that there are no longer any IGMP version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2

Label	Description
	Membership Report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv3 Leave messages for this group that it receives on this interface.
Type	Indicates how this group entry was learned. If this group entry was learned by IGMP, it will be set to "dynamic". For statically configured groups, the value will be set to 'static'.
Compat Mode	Used in order for routers to be compatible with earlier version routers. IGMPv3 hosts MUST operate in version 1 and version 2 compatibility modes. IGMPv3 hosts MUST keep state per local interface regarding the compatibility mode of each attached network. A host's compatibility mode is determined from the Host Compatibility Mode variable which can be in one of three states: IGMPv1, IGMPv2 or IGMPv3. This variable is kept per interface and is dependent on the version of General Queries heard on that interface as well as the Earlier Version Querier Present timers for the interface.

ssm-translate

Syntax

ssm-translate

ssm-translate interface *interface-name*

Context

show>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays IGMP SSM translate configuration information.

Parameters

interface-name

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

Output

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

Sample output

=====		
IGMP SSM Translate Entries		
=====		
Group Range	Source	Interface

<234.1.1.1 - 234.1.1.2>	10.1.1.1	
<232.1.1.1 - 232.1.1.5>	10.1.1.2	ies-abc

Table 11: Output fields: IGMP SSM translate

Label	Description
Group Range	Displays the address ranges of the multicast groups for which this router can be an RP
Source	Displays the unicast address that sends data on an interface
SSM Translate Entries	Displays the total number of SSM translate entries

static

Syntax

static [*ip-int-name* | *ip-addr*]

Context

show>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays static IGMP, (*,G), and (S,G) information.

Parameters

ip-int-name

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

ip-addr

Displays the information associated with the specified IP address.

Values a.b.c.d

Output

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

Sample output

```
*A:Dut-C# show router igmp static
=====
IGMP Static Group Source
=====
Source          Group          Interface
-----
*               239.1.2.1      C_Rx_acc1
10.2.1.1        239.12.1.1     C_Rx_acc1
*               239.1.1.1      C_Rx_net2
-----
Static (*,G)/(S,G) Entries : 3
=====
```

Table 12: Output fields: IGMP static

Label	Description
Source	Displays entries which represent a source address from which receivers are interested/not interested in receiving multicast traffic
Group	Displays the IP multicast group address for which this entry contains information
Interface	Displays the interface name

statistics

Syntax

```
statistics [ip-int-name | ip-address]
statistics group-interface [fwd-service service-id] [ip-int-name]
statistics host [ip-address]
```

Context

```
show>router>igmp
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays IGMP statistics information.

Parameters

- ip-int-name**

Displays the information associated with the specified IP interface name up to 32 characters.
- ip-address**

Displays the information associated with the specified IP address.

Values a.b.c.d
- service-id**

Displays the information associated with the specified service ID.

Values 1 to 2147483647 | 64 char max

Output

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

Sample output

```
*A:dut-e>show>router# igmp statistics

=====
IGMP Interface Statistics
=====
Message Type      Received      Transmitted
-----
Queries           0             57
Report V1         0             0
Report V2         0             0
Report V3         0             0
Leaves            0             0
-----

Global General Statistics
-----
Bad Length       : 0
Bad Checksum     : 0
Unknown Type     : 0
Drops           : 0
Rx Non Local     : 0
Rx Wrong Version : 0
Policy Drops     : 0
No Router Alert  : 0
Rx Bad Encodings : 0
Local Scope Pkts : 0
Resvd Scope Pkts : 0
-----

Global Source Group Statistics
-----
(S,G)           : 0
(*,G)           : 75
=====

*A:dut-e>show>router#
```

Table 13: Output fields: IGMP statistics

Label	Description
IGMP Interface Statistics	Displays the IGMP statistics for a particular interface
Message Type	<p>Queries — The number of IGMP general queries transmitted or received on this interface</p> <p>Report — The total number of IGMP V1, V2, or V3 reports transmitted or received on this interface</p> <p>Leaves — The total number of IGMP leaves transmitted on this interface</p>
Received	Displays the total number of IGMP packets received on this interface
Transmitted	Column that displays the total number of IGMP packets transmitted from this interface
General Interface Statistics	Displays the general IGMP statistics
Bad Length	Displays the total number of IGMP packets with bad length received on this interface
Bad Checksum	Displays the total number of IGMP packets with bad checksum received on this interface
Unknown Type	Displays the total number of IGMP packets with unknown type received on this interface
Bad Receive If	Displays the total number of IGMP packets incorrectly received on this interface
Rx Non Local	Displays the total number of IGMP packets received from a non-local sender
Rx Wrong Version	Displays the total number of IGMP packets with wrong versions received on this interface
Policy Drops	Displays the total number of times IGMP protocol instance matched the host IP address or group/source addresses specified in the import policy
No Router Alert	Displays the total number of IGMPv3 packets received on this interface which did not have the router alert flag set

status

Syntax
status

Context
show>router>igmp

Platforms
Supported on all 7210 SAS platforms as described in this document

Description
This command displays IGMP status information.
If IGMP is not enabled, the following message appears:

```
A:NYC# show router igmp status
MINOR: CLI IGMP is not configured.
A:NYC#
```

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

Sample output

```
*A:ALA-BA# show>router# igmp status

=====
IGMP Status
=====
Admin State           : Up
Oper State            : Up
Query Interval        : 125
Last Member Query Interval : 1
Query Response Interval : 10
Robust Count          : 2
=====
*A:ALA-BA#
```

Table 14: Output fields: IGMP status

Label	Description
Admin State	Displays the administrative status of IGMP
Oper State	Displays the current operating state of this IGMP protocol instance on this router

Label	Description
Query Interval	Displays the frequency at which IGMP query packets are transmitted
Last Member Query Interval	Displays the maximum response time inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages
Query Response Interval	Displays the maximum query response time advertised in IGMPv2 queries
Robust Count	Displays the number of times the router will retry a query

2.4.2.3 Show router PIM commands

anycast

Syntax

anycast [**detail**] [**family**]

Context

show>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays PIM anycast RP-set information.

Parameters

detail

Displays detailed information.

family

Displays IPv4 information.

Output

The following output displays an example of a PIM anycast information, and [Table 15: Output fields: PIM anycast](#) describes the output fields.

Sample output

```
A:7210SAS# show router pim anycast
=====
PIM Anycast RP Entries
```

```
=====
Anycast RP          Anycast RP Peer
-----
100.100.100.1       10.102.1.1
                    10.103.1.1
                    10.104.1.1
-----
PIM Anycast RP Entries : 3
=====
```

Table 15: Output fields: PIM anycast

Label	Description
Anycast Address	Displays the candidate anycast address
Anycast RP Peer	Displays the candidate anycast RP peer address

crp

Syntax

crp [family|ip-address]

Context

show>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays PIM candidate RP (CRP) information received at the elected bootstrap router (BSR).

Parameters

ip-address

Specifies the candidate RP IP address.

family

Displays IPv4 information.

Output

The following output is an example of a PIM CRP configuration, and [Table 16: Output fields: PIM CRP](#) describes the output fields.

Sample output

```
A:7210SAS# show router pim crp
=====
PIM Candidate RPs
```

```
=====
RP Address      Group Address      Priority    Holdtime  Expiry Time
-----
239.22.187.236  239.0.0.0/4        192        150      0d 00:02:19
239.22.187.239  239.0.0.0/4        192        150      0d 00:02:19
239.22.187.240  239.0.0.0/4        192        150      0d 00:02:09
-----
Candidate RPs : 3
=====
A:7210SAS#
```

Table 16: Output fields: PIM CRP

Label	Description
RP Address	Displays the Candidate RP address
Group Address	Displays the range of multicast group addresses for which the CRP is the Candidate RP
Priority	Displays the candidate RP priority for becoming a rendezvous point (RP). This value is used to elect RP for a group range. A value of 0 is considered as the highest priority.
Holdtime	Displays the hold time of the candidate RP. It is used by the Bootstrap router to time out the RP entries if it does not listen to another CRP advertisement within the hold time period.
Expiry	Displays the minimum time remaining before the CRP will be declared down. If the local router is not the BSR, this value is 0.
Candidate RPs	Displays the number of CRP entries

group

Syntax

group [*group-ip-address*] [**source** *ip-address*] [**type** {*starstarrp*|*starg*|*sg*}] [**detail**] [**family**]

Context

show>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays PIM source group database information.

Parameters

- group-ip-address**
Specifies the IP multicast group address for which this entry contains information.
- source ip-address**
Specifies the source address for which this entry contains information.
- type starstarrp**
Specifies that only (*, *, rp) entries be displayed.
- type starg**
Specifies that only (*,G) entries be displayed.
- type sg**
Specifies that only (S,G) entries be displayed.
- detail**
Displays detailed group information.
- family**
Displays IPv4 information.

Output

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

Sample output

```
*A:Dut-C# show router pim group
=====
Legend:  A = Active   S = Standby
=====
PIM Groups ipv4
=====
Group Address      Type      Spt Bit  Inc Intf  No.0ifs
Source Address      RP        State    Inc Intf(S)
-----
239.225.1.1        (S,G)     spt      C_A      4
10.1.1.2           10.4.4.4
239.225.1.2        (S,G)     spt      C_A      4
10.1.1.2           10.4.4.4
239.225.1.3        (S,G)     spt      C_A      4
10.1.1.2           10.4.4.4
239.225.1.4        (S,G)     spt      C_A      4
10.1.1.2           10.4.4.4
239.225.1.5        (S,G)     spt      C_A      4
10.1.1.2           10.4.4.4
239.225.1.6        (S,G)     spt      C_A      4
10.1.1.2           10.4.4.4
239.225.1.7        (S,G)     spt      C_A      4
10.1.1.2           10.4.4.4
-----
Groups : 7
=====

*A:Dut-C# show router pim group detail
=====
PIM Source Group ipv4
```

```

=====
Group Address      : 239.225.1.1
Source Address     : 10.1.1.2
RP Address         : 10.4.4.4
Advt Router       : 10.1.1.1
Flags             : spt                      Type           : (S,G)
MRIB Next Hop     : 13.1.1.1
MRIB Src Flags    : remote
Keepalive Timer Exp: 0d 00:03:16
Up Time           : 0d 00:21:44             Resolved By      : rtable-u

Up JP State       : Joined                   Up JP Expiry      : 0d 00:00:43
Up JP Rpt        : Not Joined StarG         Up JP Rpt Override: 0d 00:00:00

Register State    : No Info
Reg From Anycast RP: No

Rpf Neighbor      : 13.1.1.1
Incoming Intf     : C_A
Outgoing Intf List: C_Rx_net1, C_Rx_acc1, C_Rx_acc2, C_Rx_net2

Spt threshold     : 0 kbps                   ECMP opt threshold: 7
=====

```

Table 17: Output fields: PIM group

Label	Description
Group Address	Displays the IP multicast group address for which this entry contains information
Source Address	Displays the source address of the multicast sender. It will be 0 if the type is configured as starg. It will be the address of the Rendezvous Point (RP) if the type is configured as starRP.
RP Address	Displays the RP address
Type	Displays the type of entry: (*,*, rp)/(*,G) or (S,G)
Spt Bit	Specifies whether to forward on (*,*, rp)/(*,G) or on (S,G) state. It is updated when the (S,G) data comes on the RPF interface toward the source.
Incoming Intf	Displays the interface on which the traffic comes in. It can be the RPF interface to the RP (if starg) or the source (if sg).
Num Oifs	Displays the number of interfaces in the inherited outgoing interface list. An inherited list inherits the state from other types.
Flags	Displays the different lists that this interface belongs to
Keepalive Timer Exp	The keepalive timer is applicable only for (S,G) entries. The (S,G) keepalive timer is updated by data being forwarded using this (S,G) Forwarding state. It is used to keep (S,G) state alive in the absence of explicit (S,G) joins.

Label	Description
MRIB Next Hop	Displays the next hop address toward the RP
MRIB Src Flags	Displays the MRIB information about the source. If the entry is of type starg or starstarrp, it will contain information about the RP for the group.
Up Time	Displays the time since this source group entry was created
Resolved By	Displays the route table used for RPF check
Up JP State	Displays the upstream join prune state for this entry on the interface. PIM join prune messages are sent by the downstream routers toward the RPF neighbor.
Up JP Expiry	Displays the minimum amount of time remaining before this entry will be aged out
Up JP Rpt	Displays the join prune Rpt state for this entry on the interface. PIM join/prune messages are sent by the downstream routers toward the RPF neighbor. (S,G, rpt) state is a result of receiving (S,G, rpt) JP message from the downstream router on the RP tree.
Up JP Rpt Override	Displays the value used to delay triggered Join (S,G, rpt) messages to prevent implosions of triggered messages. If this has a non-zero value, it means that the router was in 'not Pruned' state and it saw a prune (S,G, rpt) message being sent to RPF (S,G, rpt). If the router sees a join (S,G, rpt) override message being sent by some other router on the LAN while the timer is still non-zero, it cancels the override timer. If it does not see a join (S,G, rpt) message, then on expiry of the override timer, it sends it's own join (S,G, rpt) message to RPF (S,G, rpt). A similar scenario exists when RPF (S,G, rpt) changes to become equal to RPF (*,G).
Register State	Specifies the register state. The register state is kept at the source DR. When the host starts sending multicast packets and if there are no entries programmed for that group, the source DR sends a register packet to the RP (g). Register state transition happen based on the register stop timer and the response received from the RP.
Register Stop Exp	Displays the time remaining before the register state might transition to a different state
Register from Anycast RP	Displays if the register packet for that group has been received from one of the RP from the anycast-RP set
RPF Neighbor	Displays the address of the RPF neighbor
Outgoing Intf List	Displays a list of interfaces on which data is forwarded

interface

Syntax

interface [*ip-int-name* | *int-ip-address*] [**group** [*group-ip-address*] **source** *ip-address*] [**type** {**starstarrp** | **starg** | **sg**}] [**detail**] [**family**]

Context

show>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays PIM interface information and the (S,G)/(*,G)/(*, *, rp) state of the interface.

Parameters

ip-int-name

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

ip-address

Displays the interface information associated with the specified IP address.

group *group-ip-address*

Specifies the IP multicast group address for which this entry contains information.

source *ip-address*

Specifies the source address for which this entry contains information.

If the type is starg, the value of this object will be zero.

If the type is starstarrp, the value of this object will be address of the RP.

type

Specifies the type of this entry.

Values **starstarrp, starg, sg**

detail

Displays detailed interface information.

family

Displays IPv4 information for the interface.

Output

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

Sample output

```
*7210 SAS>show>router>pim# interface
```

```
=====
PIM Interfaces ipv4
=====
Interface          Adm  Opr  DR Prty      Hello Intvl  Mcast Send
DR
-----
system             Up   Up   1           30           auto
 10.5.5.5
loopback1          Up   Up   1           30           auto
 10.1.1.5
toG_1              Up   Down 1           30           auto
toIxia_Ntw_1       Up   Up   1           30           auto
 10.2.1.5
toIxia_Ntw_2       Up   Up   1           30           auto
 10.2.2.5
toR_1              Up   Down 1           30           auto
  N/A
toIxia_1           Up   Down 1           30           auto
  N/A
toLAN_1            Up   Up   1           30           auto
 10.1.1.5
-----
Interfaces : 124
=====
*7210 SAS>show>router>pim#
```

Table 18: Output fields: PIM interface

Label	Description
Admin State	Displays the administrative state for PIM protocol on this interface
Oper State	Displays the current operational state of PIM protocol on this interface
DR	Displays the designated router on this PIM interface
DR Priority	Displays the priority value sent in PIM Hello messages and that is used by routers to elect the designated router (DR)
Hello Intvl	Indicates the frequency at which PIM Hello messages are transmitted on this interface

mc-ecmp-balance

Syntax

mc-ecmp-balance

Context

show>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays multicast balance information.

neighbor

Syntax

neighbor [*ip-address* | *ip-int-name* [**address** *neighbor-ip-address*]] [**detail**] [**family**]

Context

show>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays PIM neighbor information.

This can be important if an interface has more than one adjacency. For example, a LAN-interface configuration with three routers connected and all are running PIM on their LAN interfaces. These routers then have two adjacencies on their LAN interface, each with different neighbors. If the **address** *address* parameter is not defined in this example, then the **show** command output would display two adjacencies.

Parameters

neighbor *ip-int-name*

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

neighbor *ip-address*

Displays the interface information associated with the specified IP address.

address *ip-address*

Specifies the IP address of the neighbor, on the other side of the interface.

detail

Displays detailed neighbor information.

family

Displays IPv4 information for the specified neighbor.

Output

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

Sample output

```
ALA-1>show>router>pim# neighbor
```

```

=====
PIM Neighbor ipv4
=====
Interface      Nbr DR Prty    Up Time      Expiry Time   Hold Time
Nbr Address
-----
toB_1          1              0d 00:31:36   0d 00:01:40   105
 10.1.1.2
toE_1          1              0d 00:32:04   0d 00:01:42   105
 10.1.1.5
toE_10         1              0d 00:32:04   0d 00:01:42   105
 10.1.10.5
toE_11         1              0d 00:32:04   0d 00:01:42   105
 10.1.11.5
toE_12         1              0d 00:32:04   0d 00:01:42   105
 10.1.12.5
toE_13         1              0d 00:32:04   0d 00:01:42   105
 10.1.13.5
toE_14         1              0d 00:32:04   0d 00:01:42   105
 10.1.14.5
toE_15         1              0d 00:32:05   0d 00:01:41   105
 10.1.15.5
ALA-1#

```

Table 19: Output fields: PIM neighbor

Label	Description
Interface	Displays the neighbor interface name
Nbr DR Priority	Displays the value of the neighbor DR priority which is received in the hello message
Nbr Address	Displays the neighbor address
Expiry Time	Displays the minimum time remaining before this PIM neighbor will be aged out 0 — Means that this neighbor will never be aged out. This happens when the PIM neighbor sends a Hello message with holdtime set to '0xffff'.
Hold Time	Displays the value of the hold time present in the hello message
DR Priority	Displays the value of the neighbor DR priority which is received in the hello message
Tracking Support	Displays whether the T bit in the LAN prune delay option was present in the hello message. This indicates the neighbor capability to disable join message suppression
LAN Delay	Displays the value of the LAN delay field present in the hello message received from the neighbor
Gen Id	Displays a randomly generated 32-bit value that is regenerated each time PIM forwarding is started or restarted on the interface, including when the router restarts. When a hello message with a new GenID is received from a neighbor, any old hello

Label	Description
	information about that neighbor is discarded and superseded by the information from the new hello message.
Override Intvl (ms)	Displays the value of the override interval present in the Hello message

rp

Syntax

rp [**family** | *ip-address*]

Context

show>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the rendezvous point (RP) set information built by the router.

Parameters

family

Displays IPv4 information.

ip-address

Specifies the IP address of the RP.

Output

The following output is an example of PIM RP information, and [Table 20: Output fields: PIM RP](#) describes the output fields.

Sample output

```
A:ALA-1# show router pim rp
=====
PIM RP Set ipv4
=====
Group Address      RP Address      Type      Priority  Holdtime  Expirytime
-----
224.0.0.0/4        239.200.200.4   Dynamic   192       150
                   10.1.7.1        Static    1         N/A
-----
Group Prefixes : 1
=====
A:ALA-1#
```

Table 20: Output fields: PIM RP

Label	Description
Group Address	Displays the multicast group address of the entry
RP Address	Displays the address of the Rendezvous Point (RP)
Type	Specifies whether the entry was learned through the Bootstrap mechanism or if it was statically configured
Priority	Displays the priority for the specified group address. The higher the value, the higher the priority.
Holdtime	Displays the value of the hold time present in the BSM message

rp-hash

Syntax

rp-hash *ip-address*

Context

show>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command hashes the RP for the specified group from the RP set.

Parameters

ip-address

Displays specific multicast group addresses.

Output

The following output is an example of RP hash information, and [Table 21: Output fields: PIM RP hash](#) describes the output fields.

Sample output

```
A:ALA-1# show router pim rp-hash 239.101.0.0
=====
PIM Group-To-RP mapping
=====
Group Address      RP Address      Type
-----
239.101.0.0        239.200.200.4   Bootstrap
=====
```

```
A:ALA-1#  
  
A:ALA-1# show router pim rp-hash 239.101.0.6  
=====
```

PIM Group-To-RP mapping		
=====		
Group Address	RP Address	Type

239.101.0.6	239.200.200.4	Bootstrap
=====		

```
A:ALA-1#
```

Table 21: Output fields: PIM RP hash

Label	Description
Group Address	Displays the multicast group address of the entry
RP Address	Displays the address of the Rendezvous Point (RP)
Type	Specifies whether the entry was learned through the Bootstrap mechanism or if it was statically configured

statistics

Syntax

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

Context

show>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays statistics for a particular PIM instance.

Parameters

ip-int-name

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

int-ip-address

Displays the interface information associated with the specified IP address.

mpls-ip-name

Identifies the system created IP-MPLS tunnel interfaces, when using NG-MVPN with BGP based signaling and using P2MP LSPs setup using RSVP or mLDP.

family

Displays IPv4 information.

Output

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

Sample output

```
A:dut-g>show>router>pim# statistics

=====
PIM Statistics ipv4
=====
Message Type      Received      Transmitted   Rx Errors
-----
Hello             9690          9735          0
Join Prune        2441          6855          0
Asserts           589           0             0
Register          0             0             0
Null Register     0             0             0
Register Stop     0             0             0
BSM               0             0             0
Total Packets     12720         16590
-----
General Statistics
-----
Rx Invalid Register      : 0
Rx Neighbor Unknown      : 0
Rx Bad Checksum Discard  : 0
Rx Bad Encoding          : 0
Rx Bad Version Discard   : 0
Rx BSM Router Alert Drops : 0
Rx BSM Wrong If Drops    : 0
Rx Invalid Join Prune    : 0
Rx Unknown PDU Type      : 0
Join Policy Drops        : 0
Register Policy Drops    : 0
Bootstrap Import Policy Drops : 0
Bootstrap Export Policy Drops : 0
PDU Drops on Non-PIM/Down Intf : 0
-----
Source Group Statistics
-----
(S,G)                  : 435
(*,G)                  : 251
(*,*,RP)               : 0
=====
A:dut-g>show>router>pim#
```

Table 22: Output fields: PIM statistics

Label	Description
PIM Statistics	Displays the PIM statistics for a particular interface
Message Type	Displays the type of message

Label	Description
	<p>Hello — Displays the number of PIM hello messages received or transmitted on this interface</p> <p>Asserts — Displays the number of PIM assert messages received or transmitted on this interface</p> <p>Register — Displays the number of register messages received or transmitted on this interface</p> <p>Null Register — Displays the number of PIM null register messages received or transmitted on this interface</p> <p>Register Stop — Displays the number of PIM register stop messages received or transmitted on this interface</p> <p>BSM — Displays the number of PIM Bootstrap messages (BSM) received or transmitted on this interface</p> <p>Candidate RP Adv — Displays the number of candidate RP advertisements</p> <p>Total Packets — Displays the total number of packets transmitted and received on this interface</p>
Received	Displays the number of messages received on this interface
Transmitted	Displays the number of multicast data packets transmitted on this interface
Rx Errors	Displays the total number of receive errors
General Interface Statistics	Displays the general PIM interface statistics
Register TTL Drop	Displays the number of multicast data packets that could not be encapsulated in Register messages because the time to live (TTL) was zero
Tx Register MTU Drop	Displays the number of bootstrap messages received on this interface but were dropped
Rx Invalid Register	Displays the number of invalid PIM register messages received on this interface
Rx Neighbor Unknown	Displays the number of PIM messages (other than hello messages) that were received on this interface and were rejected because the adjacency with the neighbor router was not already established
Rx Bad Checksum Discard	Displays the number of PIM messages received on this interface which were discarded because of bad checksum
Rx Bad Encoding	Displays the number of PIM messages with bad encodings received on this interface

Label	Description
Rx Bad Version Discard	Displays the number of PIM messages with bad versions received on this interface
Rx CRP No Router Alert	Displays the number of candidate-rp advertisements (C-RP-Adv) received on this interface which had no router alert option set
Rx Invalid Join Prune	Displays the number of invalid PIM join prune messages received on this interface
Rx Unknown PDU Type	Displays the number of packets received with an unsupported PIM type
Join Policy Drops	Displays the number of times the join policy match resulted in dropping PIM join-prune message or one of the source groups contained in the message
Register Policy Drops	Displays the number of times the register policy match resulted in dropping PIM Register messages
Bootstrap Import Policy Drops	Displays the number of Bootstrap messages received on this interface that were dropped because of the bootstrap import policy
Bootstrap Export Policy Drops	Displays the number of Bootstrap messages that were not transmitted on this interface because of the bootstrap export policy
Source Group Statistics	Displays source group statistics
(S,G)	Displays the number of entries in which the type is (S,G)
(* ,G)	Displays the number of entries in which the type is (* ,G)
(* ,*,RP)	Displays the number of entries in which the type is (* , *, rp)

status

Syntax

status [**detail**] [*family*]

Context

show>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the PIM status. The oper status reflects the combined operational status of the IPv4 PIM protocol status. If both are down, the oper status will be reflected as down. If IPv4 reflects up, the oper status will reflect up.

If PIM is not enabled, the following message appears:

```
A:NYC# show router pim status
MINOR: CLI PIM is not configured.
A:NYC#
```

Parameters

- detail**
Displays detailed status information.
- family**
Displays IPv4 information.

Output

The following output is an example of PIM status information, and [Table 23: Output fields: PIM status](#) describes the output fields.

Sample output

```
A:dut-g>show>router>pim# status

=====
PIM Status ipv4
=====
Admin State                : Up
Oper State                  : Up

IPv4 Admin State           : Up
IPv4 Oper State            : Up

BSR State                   : Accept Any

Elected BSR
  Address                   : None
  Expiry Time               : N/A
  Priority                   : N/A
  Hash Mask Length          : 30
  Up Time                   : N/A
  RPF Intf toward E-BSR    : N/A

Candidate BSR
  Admin State               : Down
  Oper State                 : Down
  Address                   : None
  Priority                   : 0
  Hash Mask Length          : 30

SSM-Default-Range          : Enabled
SSM-Assert-Comp-Mode       : Disabled
SSM-Group-Range            : None

MC-ECMP-Hashing            : Disabled
```

```
Policy : None
RPF Table : rtable-u
Non-DR-Attract-Traffic : Disabled
=====
A:dut-g>show>router>pim#
```

Table 23: Output fields: PIM status

Label	Description
Admin State	Displays the administrative status of PIM
Oper State	Displays the current operating state of this PIM protocol instance
BSR State	Displays the state of the router with respect to the bootstrap mechanism
Address	Displays the address of the elected bootstrap router
Expiry Time	Displays the time remaining before the router sends the next Bootstrap message
Priority	Displays the priority of the elected bootstrap router. The higher the value, the higher the priority.
Hash Mask Length	Displays the hash mask length of the bootstrap router
Up Time	Displays the time since the current E-BSR became the bootstrap router
RPF Intf toward	Displays the RPF interface toward the elected BSR. The value is zero if there is no elected BSR in the network.
Address	Displays the address of the candidate BSR router
Expiry Time	Displays the time remaining before the router sends the next Bootstrap message
Priority	Displays the priority of the Bootstrap router. The higher the value, the higher the priority.
Hash Mask Length	Displays the hash mask length of the candidate bootstrap router
Up Time	Displays the time since becoming the bootstrap router
Admin State	Displays the administrative status of CRP
Oper State	Displays the current operating state of the CRP mechanism
Address	Displays the local RP address
Priority	Displays the CRP's priority for becoming a rendezvous point (RP). A 0 value is the highest priority.

Label	Description
Holdtime	Displays the hold time of the candidate RP. It is used by the bootstrap router to timeout the RP entries if it does not listen to another CRP advertisement within the hold time period.
Policy	Displays the PIM policies for a particular PIM instance
Default Group	Displays the default core group address
RPF Table	Displays the route table used for RPF check
MC-ECMP-Hashing	Displays if hash-based multicast balancing of traffic over ECMP links is enabled or disabled

2.4.2.4 Clear commands

database

Syntax

database [**group** *grp-ip-address* [**source** *src-ip-address*]]

database interface {*ip-int-name*|*ip-address*} [**group** *grp-ip-address* [**source** *src-ip-address*]]

database host *ip-address* [**group** *grp-ip-address* [**source** *src-ip-address*]]

database host all [**group** *grp-ip-address* [**source** *src-ip-address*]]

database group-interface all

Context

clear>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears IGMP or PIM database statistics on a specified interface or IP address.

Parameters

interface *ip-int-name*

Clears the IGMP or PIM database on the specified interface.

interface *ip-address*

Clears the IGMP or PIM database on the specified IP address.

group *grp-ip-address*

Clears the multicast group address or zero in the specified address group.

source *src-ip-address*

Clears the IGMP or PIM database from the specified source IP address.

group-interface all

Clears the IGMP database on all group interfaces.

database

Syntax

database [**interface** *ip-int-name*|*mt-int-name*|*int-ip-address*] [**group** *grp-ip-address* [**source** *ip-address*]]
[*family*]

Context

clear>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears IGMP or PIM database statistics on a specified interface or IP address.

Parameters

interface *ip-int-name*

Clears the IGMP or PIM database on the specified interface.

interface *ip-address*

Clears the IGMP or PIM database on the specified IP address.

group *group-ip-address*

Clears the multicast group address(ipv4) or zero in the specified address group.

source *ip-address*

Clears the IGMP or PIM database from the specified source IP address.

family

Clears IPv4 information.

statistics

Syntax

statistics group-interface [**fwd-service** *service-id*] *ip-int-name*

statistics group-interface all

statistics host *ip-address*

statistics host all

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

Context

clear>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears IGMP statistics on a specified interface or IP address.

An interface and a group or source cannot be specified at the same time.

Parameters

group-interface interface-name

Clears the IGMP statistics on the specifies group interface.

group-interface all

Clears the IGMP statistics on all group interfaces.

fwd-service service-id

Clears the IGMP statistics on the specified service ID.

Values 1 to 2147483647 | *svc-name: 64 char max*

host ip-address

Clears the IGMP statistics on the specified host.

host all

Clears the IGMP statistics on all hosts.

interface ip-int-name

Clears IGMP statistics on the specified interface.

interface ip-address

Clears IGMP statistics on the specified IP address.

statistics

Syntax

statistics [[[**interface** *ip-int-name* | *ip-address* | *mt-int-name*]]] {[**group** *grp-ip-address* [**source** *ip-address*]]} [*family*]

Context

clear>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears PIM statistics on a specified interface or IP address.

An interface and group or source cannot be specified at the same time.

Parameters

interface *ip-int-name*

Clears PIM statistics on the specified interface.

interface *ip-address*

Clears PIM statistics on the specified IP address.

group *grp-ip-address*

When only the group address is specified and no source is specified, (*,G) statistics are cleared. When the group address is specified along with the source address, then the (S,G) statistics are reset to zero.

source *ip-address*

When the source address is specified along with the group address, then the (S,G) statistics are reset to zero.

family

Clears IPv4 information.

version

Syntax

version group-interface [fwd-service *service-id*] *ip-int-name*

version group-interface all

version host *ip-address*

version host all

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

Context

clear>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears IGMP statistics on a specified interface or IP address.

Parameters

group-interface *interface-name*

Clears the IGMP version on the specifies group interface.

group-interface all

Clears the IGMP version on all group interfaces.

fwd-service service-id

Clears the IGMP version on the specified service ID.

Values 1 to 2147483647 | *svc-name: 64 char max*

host ip-address

Clears the IGMP version on the specified host.

host all

Clears the IGMP version on all hosts.

interface ip-int-name

Clears IGMP version on the specified interface.

interface ip-address

Clears IGMP version on the specified IP address.

neighbor

Syntax

neighbor [*ip-int-name* | *ip-address*] [*family*]

Context

clear>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears PIM neighbor data on a specified interface or IP address.

Parameters

ip-int-name

Clears PIM neighbor on the specified interface.

ip-address

Clears PIM neighbor on the specified IP address.

family

Clears IPv4 information.

igmp-snooping

Syntax

igmp-snooping

Context

clear>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context clear IGMP snooping-related data.

port-db

Syntax

port-db {**sap** *sap-id* | **sdp** *sdp-id:vc-id*} [**group** *grp-address* [**source** *ip-address*]]

Context

clear>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears the information about the IGMP snooping port database.

Parameters

sap *sap-id*

Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value. The *sap-id* can be in one of the following formats:

Encapsulation type	Syntax	Example
null	port-id	1/1/3
dot1q	port-id :qtag1	1/1/3:100
qinq	port-id :qtag1.qtag2	1/1/3:100.200

qtag1, qtag2

The encapsulation value on the specified port ID.

Values 0 to 4094

sdp sdp-id

Clears only IGMP snooping entries associated with the specified mesh SDP or spoke-SDP. For a spoke-SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id

The virtual circuit ID on the SDP ID for which to clear information.

Values 1 to 4294967295

Default for mesh SDPs only, all VC IDs

group grp-address

Clears IGMP snooping statistics matching the specified group address.

source ip-address

Clears IGMP snooping statistics matching one particular source within the multicast group.

querier

Syntax

querier

Context

clear>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears information about the IGMP snooping queriers for the VPLS service.

statistics

Syntax

statistics [sap sap-id | sdp sdp-id:vc-id]

Context

clear>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears IGMP snooping statistics for the VPLS service.

Parameters

sap *sap-id*

Displays IGMP snooping statistics for a specific SAP. The *sap-id* can be in one of the following formats:

Encapsulation type	Syntax	Example
null	port-id	1/1/3
dot1q	port-id :qtag1	1/1/3:100
qinq	port-id :qtag1.qtag2	1/1/3:100.200

qtag1, qtag2

The encapsulation value on the specified port ID.

Values 0 to 4094

sdp *sdp-id*

Displays the IGMP snooping statistics for a specific spoke or mesh SDP.

Values 1 to 17407

vc-id

The virtual circuit ID on the SDP ID for which to display information.

Values 1 to 4294967295

Default for mesh SDPs only, all VC IDs

2.4.2.5 Debug commands

2.4.2.5.1 Debug IGMP commands

group-interface

Syntax

[no] group-interface [fwd-service *service-id*] [*ip-int-name*]

Context

debug>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for IGMP group interfaces.

Parameters

service-id

Displays information associated with the specified service ID.

Values 1 to 2147483647 | *svc-name: 64 char max*

ip-int-name

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

interface

Syntax

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

Context

debug>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for IGMP interfaces.

The **no** form of this command disables the IGMP interface debugging for the specifies interface name or IP address.

Parameters

ip-int-name

Displays the information associated with the specified IP interface name.

ip-address

Displays the information associated with the specified IP address.

misc

Syntax

[no] **misc**

Context

debug>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for IGMP miscellaneous.

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

Output

The following output is an example of IGMP miscellaneous information.

Sample output

```
A:ALA-CA# debug router igmp misc
*A:ALA-CA# show debug
debug
  router
    igmp
      misc
    exit
  exit
exit
*A:ALA-CA#
```

packet

Syntax

packet [query | v1-report | v2-report | v3-report | v2-leave] host *ip-address*

packet [query | v1-report | v2-report | v3-report | v2-leave] [*ip-int-name* | *ip-address*]

no packet [query | v1-report | v2-report | v3-report | v2-leave] [*ip-int-name* | *ip-address*]

no packet [query | v1-report | v2-report | v3-report | v2-leave] host *ip-address*

Context

debug>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables/disables debugging for IGMP packets.

Parameters

query

Specifies to log the IGMP group- and source-specific queries transmitted and received on this interface.

v1-report

Specifies to log IGMP V1 reports transmitted and received on this interface.

v2-report

Specifies to log IGMP V2 reports transmitted and received on this interface.

v3-report

Specifies to log IGMP V3 reports transmitted and received on this interface.

v2-leave

Specifies to log the IGMP Leaves transmitted and received on this interface.

ip-int-name

Displays the information associated with the specified IP interface name.

ip-address

Displays the information associated with the specified IP address.

2.4.2.5.2 Debug PIM commands

adjacency

Syntax

[no] adjacency

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for PIM adjacencies.

all

Syntax

all [group *grp-ip-address*] [source *ip-address*] [detail]

no all

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for all the PIM modules.

Parameters

group *grp-ip-address*

Debugs information associated with all PIM modules.

Values IPv4 address

source *ip-address*

Debugs information associated with all PIM modules.

Values IPv4 address

detail

Debugs detailed information about all PIM modules.

assert

Syntax

assert [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]

no assert

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for PIM assert mechanism.

Parameters

group *grp-ip-address*

Debugs information associated with the PIM assert mechanism.

Values multicast group address (ipv4)

source *ip-address*

Debugs information associated with the PIM assert mechanism.

Values source address (ipv4)

detail

Debugs detailed information about the PIM assert mechanism.

bgp

Syntax

bgp [**source** *ip-address*] [**group** *group-ip-address*] [**peer** *peer-ip-address*]

no bgp

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for PIM/BGP specific interoperation.

Parameters

ip-address

Debugs BGP information associated with the specified source.

Values source address (ipv4)

group-ip-address

Debugs BGP information associated with the specified group.

Values group address (ipv4)

peer-ip-address

Debugs BGP information associated with the specified peer.

Values peer address (ipv4)

bsr

Syntax

bsr [**detail**]

no bsr

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for the PIM bootstrap mechanism.

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

Parameters

detail

Debugs detailed information about the PIM assert mechanism.

data

Syntax

data [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]

no data

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for PIM data exception.

Parameters

group *grp-ip-address*

Debugs information associated with the specified data exception.

Values multicast group address (ipv4)

source *ip-address*

Debugs information associated with the specified data exception.

Values source address (ipv4)

detail

Debugs detailed IP data exception information.

db

Syntax

db [*group grp-ip-address*] [*source ip-address*] [*detail*]

no db

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for the PIM database.

Parameters

group *grp-ip-address*

Debugs information associated with the specified database.

Values multicast group address (ipv4) or zero

source *ip-address*

Debugs information associated with the specified database.

Values source address (ipv4)

detail

Debugs detailed IP database information.

interface

Syntax

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

no interface

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for the PIM interface.

Parameters

ip-int-name

Debugs the information associated with the specified IP interface name.

Values IPv4 interface address

ip-address

Debugs the information associated with the specified IP address.

detail

Debugs detailed IP interface information.

jp

Syntax

jp [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]

no jp

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for the PIM join-prune mechanism.

Parameters

group *grp-ip-address*

Debugs information associated with the specified join-prune mechanism.

Values multicast group address (ipv4) or zero

source *ip-address*

Debugs information associated with the specified join-prune mechanism.

Values source address (ipv4)

detail

Debugs detailed join-prune mechanism information.

mrrib

Syntax

mrrib [**group** *grp-ip-address*] [**source** *ip-address*] [**detail**]

no mrib

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for PIM MRIB.

Parameters

group *grp-ip-address*

Debugs information associated with the specified PIM MRIB.

Values multicast group address (ipv4)

source *ip-address*

Debugs information associated with the specified PIM MRIB.

Values source address (ipv4)

detail

Debugs detailed MRIB information.

msg

Syntax

msg [**detail**]

no msg

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for PIM messaging.

Parameters

detail

Debugs detailed messaging information.

packet

Syntax

packet [**hello** | **register** | **register-stop** | **jp** | **bsr** | **assert**] [*ip-int-name* | *ip-address*]
no packet

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for PIM packets.

Parameters

hello | ***register*** | ***register-stop*** | ***jp*** | ***bsr*** | ***assert*** | ***crp***
PIM packet types.

ip-int-name

Debugs the information associated with the specified IP interface name.

Values IPv4 interface address

ip-address

Debugs the information associated with the specified IP address of a particular packet type.

red

Syntax

red [**detail**]
no red

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for PIM redundancy messages to the standby CPM.

Parameters

detail

Displays detailed redundancy information.

register

Syntax

register [group *grp-ip-address*] [**source** *ip-address*] [**detail**]

no register

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for the PIM register mechanism.

Parameters

group *grp-ip-address*

Debugs information associated with the specified PIM register.

Values multicast group address (ipv4)

source *ip-address*

Debugs information associated with the specified PIM register.

Values source address (ipv4)

detail

Debugs detailed register information.

rtm

Syntax

rtm [**detail**]

no rtm

Context

debug>router>pim

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables debugging for PIM RTM.

Parameters

detail

Debugs detailed RTM information.

3 OSPF

This chapter provides information about configuring the Open Shortest Path First (OSPF) protocol.



Note:

- OSPFv3 is not supported for use as a PE-CE routing protocol on any of the platforms as described in this document.
- The platforms as described in this document allow for the configuration of a single instance at any time. The instance ID can be any number other than 0. This enables these platforms to be used in a network where multi-instance OSPF is deployed, and the node needs to use an instance ID other than the default instance ID of 0.
- On the 7210 SAS-K 2F6C4T, scaling is designed so that the platforms can fit into an OSPF stub area or NSSA area.

3.1 Configuring OSPF

OSPF (Open Shortest Path First) is a hierarchical link state protocol. OSPF is an interior gateway protocol (IGP) used within large autonomous systems (ASs). OSPF routers exchange state, cost, and other relevant interface information with neighbors. The information exchange enables all participating routers to establish a network topology map. Each router applies the Dijkstra algorithm to calculate the shortest path to each destination in the network. The resulting OSPF forwarding table is submitted to the routing table manager to calculate the routing table.

When a router is started with OSPF configured, OSPF, along with the routing-protocol data structures, is initialized and waits for indications from lower-layer protocols that its interfaces are functional. Nokia's implementation of OSPF conforms to OSPF Version 2 specifications presented in RFC 2328, *OSPF Version 2*. Routers running OSPF can be enabled with minimal configuration. All default and command parameters can be modified.

Key OSPF features are:

- backbone areas
- stub areas
- Not-So-Stubby areas (NSSAs)
- virtual links
- authentication
- route redistribution
- routing interface parameters
- OSPF-TE extensions (Nokia's implementation allows MPLS fast reroute)
- addressing semantics have been removed from OSPF packets and the basic link-state advertisements (LSAs); new LSAs have been created to carry IPv6 addresses and prefixes
- OSPF3 runs on a per-link basis, instead of on a per-IP-subnet basis

- unlike OSPFv2, OSPFv3 authentication relies on IPV6's authentication header and encapsulating security payload
- most packets in OSPF for IPv6 are almost as compact as those in OSPF for IPv4, even with the larger IPv6 addresses

The 7210 SAS-K 2F6C4T and the 7210 SAS-K 3SFP+ 8C support IGP-LDP synchronization on OSPF routes. See the "IGP-LDP and Static Route-LDP Synchronization on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C" in the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide* for more information.

3.1.1 OSPF areas

The hierarchical design of OSPF allows a collection of networks to be grouped into a logical area. An area's topology is concealed from the rest of the AS which significantly reduces OSPF protocol traffic. With the proper network design and area route aggregation, the size of the route-table can be drastically reduced which results in decreased OSPF route calculation time and topological database size.


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 is used.

Routers that belong to more than one area are called area border routers (ABRs). An ABR maintains a separate topological database for each area it is connected to. Every router that belongs to the same area has an identical topological database for that area.

3.1.1.1 Backbone area

The OSPF backbone area, area 0.0.0.0, must be contiguous and all other areas must be connected to the backbone area. The backbone distributes routing information between areas. If it is not practical to connect an area to the backbone (see area 0.0.0.5 in the following figure), the ABRs (such as routers Y and Z) must be connected via a virtual link. The two ABRs form a point-to-point-like adjacency across the transit area (see area 0.0.0.4).

Figure 2: Backbone area



3.1.1.2 Stub area

A stub area is a designated area that does not allow external route advertisements. Routers in a stub area do not maintain external routes. A single default route to an ABR replaces all external routes. This OSPF implementation supports the optional summary route (type-3) advertisement suppression from other areas into a stub area. This feature further reduces topological database sizes and OSPF protocol traffic, memory usage, and CPU route calculation time.

In [Figure 2: Backbone area](#), areas 0.0.0.1, 0.0.0.2 and 0.0.0.5 could be configured as stub areas. A stub area cannot be designated as the transit area of a virtual link and a stub area cannot contain an AS boundary router. An AS boundary router exchanges routing information with routers in other ASs.

3.1.1.3 Not-So-Stubby Area

Another OSPF area type is called a Not-So-Stubby area (NSSA). NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. External routes learned by OSPF routers in the NSSA area are advertised as type-7 LSAs within the NSSA area and are translated by ABRs into type-5 external route advertisements for distribution into other areas of the OSPF domain. An NSSA area cannot be designated as the transit area of a virtual link.

In [Figure 2: Backbone area](#), area 0.0.0.3 could be configured as a NSSA area.

3.1.1.3.1 OSPF super backbone

The 7210 SAS PE routers have implemented a version of the BGP/OSPF interaction procedures as defined in RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*. The features included in this RFC are:

- loop prevention
- handling LSAs received from the CE
- sham links
- managing VPN-IPv4 routes received by BGP

VPN routes can be distributed among the PE routers by BGP. If the PE uses OSPF to distribute routes to the CE router, the standard procedures governing BGP/OSPF interactions causes routes from one site to be delivered to another in type 5 LSAs, as AS-external routes.

The MPLS VPN super backbone behaves like an additional layer of hierarchy in OSPF. The PE-routers that connect the respective OSPF areas to the super backbone function as OSPF Area Border Routers (ABR) in the OSPF areas to which they are attached. To achieve full compatibility, they can also behave as AS Boundary Routers (ASBR) in non-stub areas.

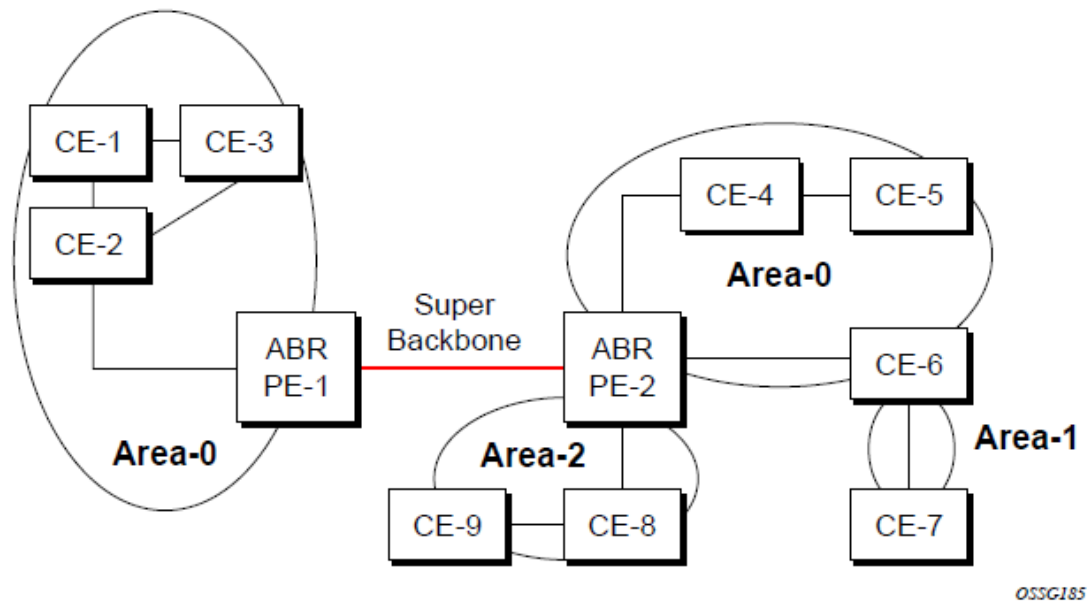
The PE-routers insert inter-area routes from other areas into the area in which the CE-router is present. The CE-routers are not involved at any level nor are they aware of the super backbone or of other OSPF areas present beyond the MPLS VPN super backbone.

The CE always assumes the PE is an ABR:

- If the CE is in the backbone then the CE router assumes that the PE is an ABR linking one or more areas to the backbone.
- If the CE is not in the backbone, then the CE believes that the backbone is on the other side of the PE.
- As such, the super backbone looks like another area to the CE.

In the following figure, the PEs are connected to the MPLS-VPN super backbone. To be able to distinguish if two OSPF instances are in fact the same and require Type 3 LSAs to be generated or are two separate routing instances where type 5 external LSAs need to be generated the concept of a domain-id is introduced.

Figure 3: PEs connected to an MPLS-VPN super backbone



The domain ID is carried with the MP-BGP update and indicates the source OSPF Domain. When the routes are being redistributed into the same OSPF Domain, the concepts of previously described super backbone apply and Type 3 LSAs should be generated. If the OSPF domain does not match, then the route type will be external.

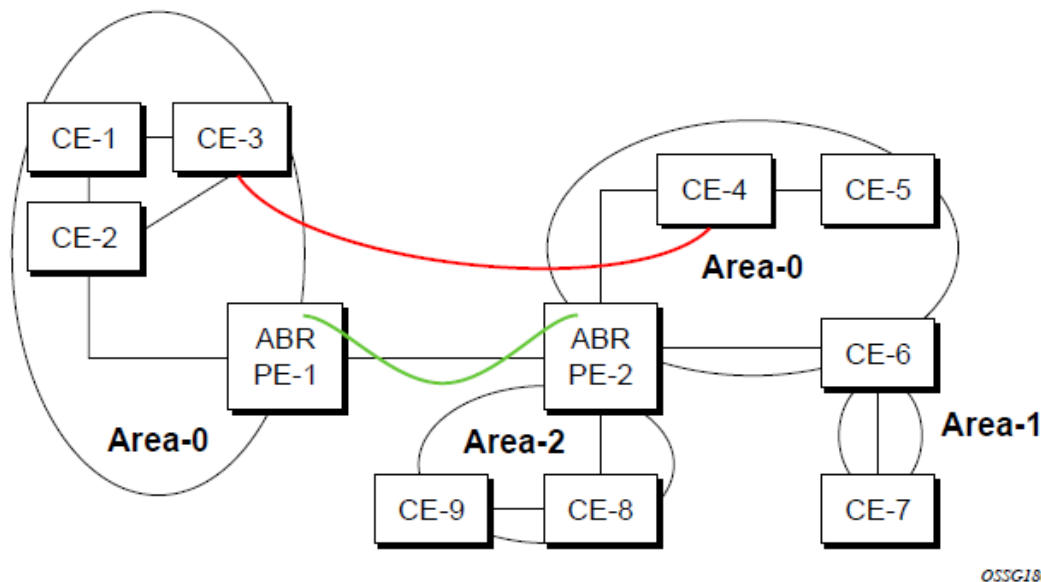
Configuring the super backbone (not the sham links) makes all destinations learned by PEs with matching domain IDs inter-area routes.

When configuring sham links, these links become intra-area routes if they are present in the same area.

3.1.1.3.2 Sham links

The following figure shows the red link between CE-3 and CE-4 could be a low speed OC-3/STM-1 link but because it establishes a intra-area route connection between the CE-3 and CE-4 the potentially high-speed PE-1 to PE-2 connection will not be utilized. Even with a super backbone configuration it is regarded as a inter-area connection.

Figure 4: Sham links



The establishment of the (green) sham-link is also constructed as an intra-area link between PE routers, a normal OSPF adjacency is formed and the link-state database is exchanged across the MPLS-VPN. As a result, the desired intra-area connectivity is created, at this time the cost of the green and red links can be managed such that the red link becomes a standby link only in case the VPN fails.

A sham link is only required if a back door link (shown as the red link in the preceding figure) is present; otherwise, configuring an OSPF super backbone will probably suffice.

3.1.1.3.3 Implementing the OSPF super backbone

With the OSPF super backbone architecture, the continuity of OSPF routing is preserved:

- The OSPF intra-area LSAs (type-1 and type-2) advertised by the CE are inserted into the MPLS-VPN super backbone by redistributing the OSPF route into MP-BGP by the PE adjacent to the CE.
- The MP-BGP route is propagated to other PE-routers and inserted as an OSPF route into other OSPF areas. Considering the PEs across the super backbone always act as ABRs they will generate inter area route OSPF summary LSAs, Type 3.
- The inter-area route can now be propagated into other OSPF areas by other customer owned ABRs within the customer site.
- Customer Area 0 (backbone) routes when carried across the MPLS-VPN using MPBGP will appear as Type 3 LSAs even if the customer area remains area 0 (backbone).

A BGP extended community (OSPF domain ID) provides the source domain of the route. This domain ID is not carried by OSPF but carried by MP-BGP as an extended community attribute.

If the configured extended community value matches the receiving OSPF domain, then the OSPF super backbone is implemented.

From a BGP perspective, the cost is copied into the MED attribute.

3.1.1.3.4 Loop avoidance

If a route sent from a PE router to a CE router could then be received by another PE router from one of its own CE routers then it is possible for routing loops to occur. RFC 4577 specifies several methods of loop avoidance.

3.1.1.3.5 DN-BIT

When a Type 3 LSA is sent from a PE router to a CE router, the DN bit in the LSA options field is set. This is used to ensure that if any CE router sends this Type 3 LSA to a PE router, the PE router will not redistribute it further.

When a PE router needs to distribute to a CE router a route that comes from a site outside the latter's OSPF domain, the PE router presents itself as an ASBR (Autonomous System Border Router), and distributes the route in a type 5 LSA. The DN bit MUST be set in these LSAs to ensure that they will be ignored by any other PE routers that receive them.

DN-BIT loop avoidance is also supported.

3.1.1.3.6 Route tag

If a particular VRF in a PE is associated with an instance of OSPF, then by default it is configured with a special OSPF route tag value called the VPN route tag. This route tag is included in the Type 5 LSAs that the PE originates and sends to any of the attached CEs. The configuration and inclusion of the VPN Route Tag is required for backward compatibility with deployed implementations that do not set the DN bit in Type 5 LSAs.

3.1.2 OSPFv3 authentication

OSPFv3 authentication requires IPv6 IPsec and supports the following:

- IPsec transport mode
- AH and ESP
- Manual keyed IPsec Security Association (SA)
- Authentication Algorithms MD5 and SHA1

To pass OSPFv3 authentication, OSPFv3 peers must have matching inbound and outbound SAs configured using the same SA parameters such as SPI, keys and related parameters. The implementation must allow the use of one SA for both inbound and outbound directions.

The re-keying procedure defined in RFC 4552, *Authentication/Confidentiality for OSPFv3*, supports the following:

- For every router on the link, create an additional inbound SA for the interface being re-keyed using a new SPI and the new key.
- For every router on the link, replace the original outbound SA with one using the new SPI and key values. The SA replacement operation must be atomic with respect to sending OSPFv3 packet on the link, so that no OSPFv3 packets are sent without authentication or encryption.
- For every router on the link, remove the original inbound SA.

The key rollover procedure automatically starts when the operator changes the configuration of the inbound static-SA or bidirectional static-SA under an interface or virtual link. Within the KeyRolloverInterval time period, OSPF3 accepts packets with both the previous inbound static-SA and the new inbound static-SA, and the previous outbound static-SA should continue to be used. When the timer expires, OSPF3 only accepts packets with the new inbound static-SA and for outgoing OSPF3 packets, the new outbound static-SA is used instead.

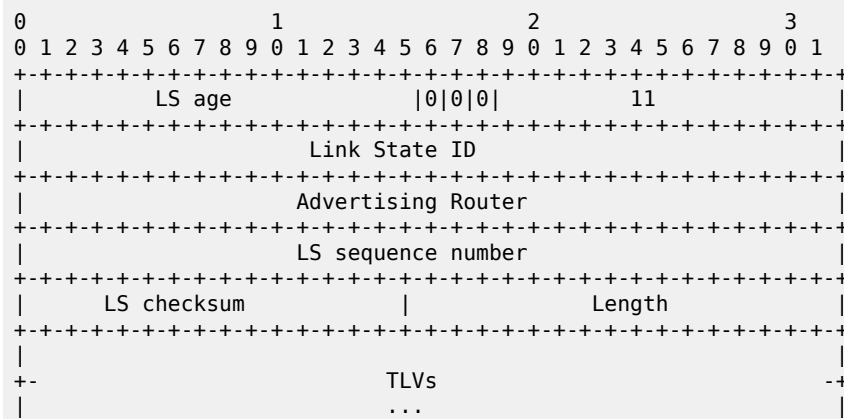
3.1.3 OSPFv3 graceful restart helper

This feature extends the Graceful Restart helper function supported on OSPFv2 protocols to OSPFv3:


The primary difference between graceful restart helper for OSPFv2 and OSPFv3 is in OSPFv3 a different grace-LSA format is used.

The graceful restart helper mode allows SR OS-based systems to provide a grace period to other routers which have requested it, during which the SR OS systems will continue to use routes authored by or transiting the router requesting the grace period. This is typically used when another router is rebooting the control plane but the forwarding plane is expected to continue to forward traffic based on the previously available FIB.

The grace-LSA format for OSPF restart (GRACE) LSA format is:



The Link State ID of a grace-LSA in OSPFv3 is the Interface ID of the interface originating the LSA. The format of each TLV is:



3.1.4 Virtual links

The backbone area in an OSPF AS must be contiguous and all other areas must be connected to the backbone area. Sometimes, this is not possible. You can use virtual links to connect to the backbone through a non-backbone area.

Figure 2: Backbone area shows routers Y and Z as the start and end points of the virtual link while area 0.0.0.4 is the transit area. To configure virtual links, the router must be an ABR. Virtual links are identified by the router ID of the other endpoint, another ABR. These two endpoint routers must be attached to a common area, called the transit area. The area through which you configure the virtual link must have full routing information.

Transit areas pass traffic from an area adjacent to the backbone or to another area. The traffic does not originate in, nor is it destined for, the transit area. The transit area cannot be a stub area or a NSSA area.

Virtual links are part of the backbone, and behave as if they were unnumbered point-to-point networks between the two routers. A virtual link uses the intra-area routing of its transit area to forward packets. Virtual links are brought up and down through the building of the shortest-path trees for the transit area.

3.1.5 Neighbors and adjacencies

A router uses the OSPF Hello protocol to discover neighbors. A neighbor is a router configured with an interface to a common network. The router sends hello packets to a multicast address and receives hello packets in return.

In broadcast networks, a designated router and a backup designated router are elected. The designated router is responsible for sending link-state advertisements (LSAs) describing the network, which reduces the amount of network traffic.

The routers attempt to form adjacencies. An adjacency is a relationship formed between a router and the designated or backup designated router. For point-to-point networks, no designated or backup designated router is elected. An adjacency must be formed with the neighbor.

To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.

When the link-state databases of two neighbors are synchronized, the routers are considered to be fully adjacent. When adjacencies are established, pairs of adjacent routers synchronize their topological databases. Not every neighboring router forms an adjacency. Routing protocol updates are only sent to and received from adjacencies. Routers that do not become fully adjacent remain in the two-way neighbor state.

3.1.6 Link-state advertisements

Link-state advertisements (LSAs) describe the state of a router or network, including router interfaces and adjacency states. Each LSA is flooded throughout an area. The collection of LSAs from all routers and networks form the protocol's topological database.

The distribution of topology database updates take place along adjacencies. A router sends LSAs to advertise its state according to the configured interval and when the router's state changes. These packets include information about the router's adjacencies, which allows detection of non-operational routers.

When a router discovers a routing table change or detects a change in the network, link state information is advertised to other routers to maintain identical routing tables. Router adjacencies are reflected in the contents of its link state advertisements. The relationship between adjacencies and the link states allow the protocol to detect non-operating routers. Link state advertisements flood the area. The flooding mechanism ensures that all routers in an area have the same topological database. The database consists of the collection of LSAs received from each router belonging to the area.

OSPF sends only the part that has changed and only when a change has taken place. From the topological database, each router constructs a tree of shortest paths with itself as root. OSPF distributes routing information between routers belonging to a single AS.

3.1.7 Metrics

In OSPF, all interfaces have a cost value or routing metric used in the OSPF link-state calculation. OSPF uses cost values to determine the best path to a particular destination: the lower the cost value, the more likely the interface will be used to forward data traffic.

3.1.8 Authentication

All OSPF protocol exchanges can be authenticated. This means that only trusted routers can participate in autonomous system routing. Nokia's implementation of OSPF supports plain text and Message Digest 5 (MD5) authentication (also called simple password).

MD5 allows an authentication key to be configured per network. Routers in the same routing domain must be configured with the same key. When the MD5 hashing algorithm is used for authentication, MD5 is used to verify data integrity by creating a 128-bit message digest from the data input. It is unique to that data. Nokia's implementation of MD5 allows the migration of an MD5 key by using a key ID for each unique key.

By default, authentication is not enabled on an interface.

3.1.9 Multiple OSPF instances



Note:

Nokia recommends using only a single instance of OSPFv2. This allows for the use of different instance IDs, if required by the customer.

3.1.9.1 Route export policies for OSPF

Route policies allow specification of the source OSPF process ID in the **from** and **to** parameters in the **config>router>policy-options>policy-statement>entry>from** context, for example **from protocol ospf instance-id**.

If an *instance-id* is specified, only routes installed by that instance are picked up for announcement. If no *instance-id* is specified, then only routes installed by the base instance will be announced. The **all** keyword announces routes installed by all instances of OSPF.

When announcing internal (intra/inter-area) OSPF routes from another process, the default type should be type-1, and metric set to the route metric in RTM. For AS-external routes, by default the route type (type-1/2) should be preserved in the originated LSA, and metric set to the route metric in RTM. By default,

the tag value should be preserved when an external OSPF route is announced by another process. All these can be changed with explicit action statements.

Export policy should allow a match criteria based on the OSPF route hierarchy (for example, only intra-area, only inter-area, only external, only internal (intra/inter-area)). There must also be a possibility to filter based on existing tag values.

3.1.9.2 Preventing route redistribution loops

The legacy method for this was to assign a tag value to each OSPF process and mark each external route originated within that domain with that value. However, because the tag value must be preserved throughout different OSPF domains, this only catches loops that go back to the originating domain and not where looping occurs in a remote set of domains. To prevent this type of loop, the route propagation information in the LSA must be accumulative. The following method has been implemented:

- The OSPF tag field in the AS-external LSAs is treated as a bit mask, instead of a scalar value. That is, each bit in the tag value can be independently checked, set or reset as part of the routing policy.
- When a set of OSPF domains are provisioned in a network, each domain is assigned a specific bit value in the 32-bit tag mask. When an external route is originated by an ASBR using an internal OSPF route in a specific domain, a corresponding bit is set in the AS-external LSA. As the route gets redistributed from one domain to another, more bits are set in the tag mask, each corresponding to the OSPF domain the route visited. Route redistribution looping is prevented by checking the corresponding bit as part of the export policy--if the bit corresponding to the announcing OSPF process is already set, the route is not exported there.

From the CLI perspective, this involves adding a set of **from tag** and **action tag** commands that allow for bit operations.

3.1.10 IP subnets

OSPF enables the flexible configuration of IP subnets. Each distributed OSPF route has a destination and mask. A network mask is a 32-bit number that indicates the range of IP addresses residing on a single IP network/subnet. This specification displays network masks as hexadecimal numbers; for example, the network mask for a class C IP network is displayed as 0xfffff00. Such a mask is often displayed as 255.255.255.0.

Two different subnets with same IP network number have different masks, called variable length subnets. A packet is routed to the longest or most specific match. Host routes are considered to be subnets whose masks are all ones (0xffffffff).

3.1.11 Preconfiguration recommendations

Before configuring OSPF, the router ID must be available. The router ID is a 32-bit number assigned to each router running OSPF. This number uniquely identifies the router within an AS. OSPF routers use the router IDs of the neighbor routers to establish adjacencies. Neighbor IDs are learned when Hello packets are received from the neighbor.

Before configuring OSPF parameters, ensure that the router ID is derived by one of the following methods:

- Define the value in the **config>router router-id** context.

- Define the system interface in the **config>router>interface** *ip-int-name* context (used if the router ID is not specified in the **config>router** *router-id* context).

A system interface must have an IP address with a 32-bit subnet mask. The system interface is used as the router identifier by higher-level protocols such as OSPF and IS-IS. The system interface is assigned during the primary router configuration process when the interface is created in the logical IP interface context.

- If you do not specify a router ID, then the last four bytes of the MAC address are used.

3.2 IP Fast-Reroute (IP FRR) for OSPF and IS-IS prefixes



Note:

Only LDP FRR is supported. LDP FRR uses the LFA computed for IP prefixes to determine the backup path to use for LDP FEC that are installed in the MPLS tables. This section is here only for completeness of description for this feature.

This feature provides for the use of the Loop-Free Alternate (LFA) backup next-hop for forwarding in-transit and CPM generated IP packets when the primary next-hop is not available. This means that a node resumes forwarding IP packets to a destination prefix without waiting for the routing convergence.

When any of the following events occurs, IGP instructs in the fast path the IOM, the forwarding engine to enable the LFA backup next-hop:

- OSPF/IS-IS interface goes operationally down: physical or local admin shutdown.
- Timeout of a BFD session to a next-hop when BFD is enabled on the OSPF/IS-IS interface.

IP FRR is supported on IPv4 and IPv6 OSPF/IS-IS prefixes forwarded in the base router instance to a network IP interface or to an IES SAP interface or spoke interface. It is also supported for VPRN VPN-IPv4 OSPF prefixes and VPN-IPv6 OSPF prefixes forwarded to a VPRN SAP interface or spoke interface.

The LFA next-hop precomputation by IGP is described in RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*.

3.2.1 IP FRR/LFA configuration



Note:

IP FRR is not supported on 7210 SAS nodes. LFA is supported on 7210 SAS nodes that support LDP FRR.

The user first enables Loop-Free Alternate (LFA) computation by SPF under the IS-IS routing protocol level or under the OSPF routing protocol instance level:

config>router>isis>loopfree-alternate config>router>ospf>loopfree-alternate

The preceding commands instruct the IGP SPF to attempt to precompute both a primary next-hop and an LFA next-hop for every learned prefix. When found, the LFA next-hop is populated into the RTM along with the primary next-hop for the prefix.

3.2.1.1 Reducing the scope of the LFA calculation by SPF

The user can instruct IGP to not include all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.

```
config>router>isis>level>loopfree-alternate-exclude
```

```
config>router>ospf>area>loopfree-alternate-exclude
```

The user can also exclude a specific IP interface from being included in the LFA SPF computation by IS-IS or OSPF:

```
config>router>isis>interface>loopfree-alternate-exclude
```

```
config>router>ospf>area>interface>loopfree-alternate-exclude
```

Note that when an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When the user excludes an interface from the LFA SPF in OSPF, it is excluded in all areas. However, the preceding OSPF command can only be executed under the area in which the specified interface is primary and after enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command fails.

3.2.2 ECMP considerations

Whenever the SPF computation determined there is more than one primary next-hop for a prefix, it will not program any LFA next-hop in RTM. Therefore, IP prefixes will resolve to the multiple primary next-hops in this case which provides the required protection.

3.2.3 IP FRR and RSVP shortcut (IGP shortcut)

When both IGP shortcut and LFA are enabled in IS-IS or OSPF, and IP FRR is also enabled, then the following additional IP FRR capabilities are supported:

- A prefix which is resolved to a direct primary next-hop can be backed up by a tunneled LFA next-hop.
- A prefix which is resolved to a tunneled primary next-hop will not have an LFA next-hop. It will rely on RSVP FRR for protection.

The LFA SPF is extended to use IGP shortcuts as LFA next-hops as described in [OSPF and IS-IS support for Loop-Free Alternate calculation](#).

3.2.4 IP FRR and BGP next-hop resolution


An LFA backup next-hop will be able to protect the primary next-hop to reach a prefix advertised by a BGP neighbor. The BGP next-hop will therefore remain up when the FIB switches from the primary IGP next-hop to the LFA IGP next-hop.

3.2.5 OSPF and IS-IS support for Loop-Free Alternate calculation

SPF computation in IS-IS and OSPF is enhanced to compute LFA alternate routes for each learned prefix and populate it in RTM.

The following figure shows a simple network topology with point-to-point (P2P) interfaces and highlights three routes to reach router R5 from router R1.

Figure 5: Topology example with primary and LFA routes




OSSG712

The primary route is via R3. The LFA route via R2 has two equal cost paths to reach R5. The path by way of R3 protects against failure of link R1-R3. This route is computed by R1 by checking that the cost for R2 to reach R5 by way of R3 is lower than the cost by way of routes R1 and R3. This condition is referred to as the "loop-free criterion".

The path by way of R2 and R4 can be used to protect against the failure of router R3. However, with the link R2-R3 metric set to 5, R2 sees the same cost to forward a packet to R5 by way of R3 and R4. Therefore R1 cannot guarantee that enabling the LFA next-hop R2 will protect against R3 node failure. This means that the LFA next-hop R2 provides link-protection only for prefix R5. If the metric of link R2-R3 is changed to 8, then the LFA next-hop R2 provides node protection since a packet to R5 will always go over R4. That is, it is required that R2 becomes loop-free with respect to both the source node R1 and the protected node R3.

Consider now the case where the primary next-hop uses a broadcast interface as shown in the following figure.

Figure 6: Topology example with broadcast interfaces



OSSG713

In order for next-hop R2 to be a link-protect LFA for route R5 from R1, it must be loop-free with respect to the R1-R3 link Pseudo-Node (PN). However, because R2 has also a link to that PN, its cost to reach R5 by way of the PN or router R4 are the same. Therefore R1 cannot guarantee that enabling the LFA next-hop R2 will protect against a failure impacting link R1-PN because this may cause the entire subnet represented by the PN to go down. If the metric of link R2-PN is changed to 8, then R2 next-hop will be an LFA providing link protection.

The following are the detailed equations for this criterion as provided in RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*:

- **Rule 1**

Link-protect LFA backup next-hop (primary next-hop R1-R3 is a P2P interface):

$$\text{Distance_opt}(\text{R2}, \text{R5}) < \text{Distance_opt}(\text{R2}, \text{R1}) + \text{Distance_opt}(\text{R1}, \text{R5})$$

and,

$$\text{Distance_opt}(\text{R2}, \text{R5}) \geq \text{Distance_opt}(\text{R2}, \text{R3}) + \text{Distance_opt}(\text{R3}, \text{R5})$$

- **Rule 2**

Node-protect LFA backup next-hop (primary next-hop R1-R3 is a P2P interface):

$$\text{Distance_opt}(\text{R2}, \text{R5}) < \text{Distance_opt}(\text{R2}, \text{R1}) + \text{Distance_opt}(\text{R1}, \text{R5})$$

and,

$$\text{Distance_opt}(\text{R2}, \text{R5}) < \text{Distance_opt}(\text{R2}, \text{R3}) + \text{Distance_opt}(\text{R3}, \text{R5})$$

- **Rule 3**

Link-protect LFA backup next-hop (primary next-hop R1-R3 is a broadcast interface):

$$\text{Distance_opt}(R2, R5) < \text{Distance_opt}(R2, R1) + \text{Distance_opt}(R1, R5)$$

and,

$$\text{Distance_opt}(R2, R5) < \text{Distance_opt}(R2, \text{PN}) + \text{Distance_opt}(\text{PN}, R5)$$

where; PN stands for the R1-R3 link Pseudo-Node.

For the case of P2P interface, if SPF finds multiple LFA next-hops for a specific primary next-hop, it follows the following selection algorithm:

1. It will pick the node-protect type in favor of the link-protect type.
2. If there is more than one LFA next-hop within the selected type, then it will pick one based on the least cost.
3. If more than one LFA next-hop with the same cost results from step 2, then SPF will select the first one. This is not a deterministic selection and will vary following each SPF calculation.

For the case of a broadcast interface, a node-protect LFA is not necessarily a link protect LFA if the path to the LFA next-hop goes over the same PN as the primary next-hop. Similarly, a link protect LFA may not guarantee link protection if it goes over the same PN as the primary next-hop. The selection algorithm when SPF finds multiple LFA next-hops for a specific primary next-hop is modified as follows:

1. The algorithm splits the LFA next-hops into two sets:
 - The first set consists of LFA next-hops which do not go over the PN used by primary next-hop.
 - The second set consists of LFA next-hops which do go over the PN used by the primary next-hop.
2. If there is more than one LFA next-hop in the first set, it will pick the node-protect type in favor of the link-protect type.
3. If there is more than one LFA next-hop within the selected type, then it will pick one based on the least cost.
4. If more than one LFA next-hop with equal cost results from step 3, SPF will select the first one from the remaining set. This is not a deterministic selection and will vary following each SPF calculation.
5. If no LFA next-hop results from step 4, SPF will rerun steps 2-4 using the second set.

Note this algorithm is more flexible than strictly applying Rule 3; that is, the link protect rule in the presence of a PN and specified in RFC 5286. A node-protect LFA which does not avoid the PN, that is, does not guarantee link protection, can still be selected as a last resort. The same thing, a link-protect LFA which does not avoid the PN may still be selected as a last resort.

Both the computed primary next-hop and LFA next-hop for a specific prefix are programmed into RTM.

3.2.5.1 Loop-Free Alternate calculation in the presence of IGP shortcuts

To expand the coverage of the LFA backup protection in a network, RSVP LSP based IGP shortcuts can be placed selectively in parts of the network and be used as an LFA backup next-hop.

When IGP shortcut is enabled in IS-IS or OSPF on a specific node, all RSVP LSP originating on this node and with a destination address matching the router-id of any other node in the network are included in the main SPF by default.

To limit the time it takes to compute the LFA SPF, the user must explicitly enable the use of an IGP shortcut as LFA backup next-hop using one of a couple of new optional argument for the existing LSP level IGP shortcut command:

config router mpls lsp igp-shortcut [lfa-only]

The **lfa-only** option allows an LSP to be included in the LFA SPF only such that the introduction of IGP shortcuts does not impact the main SPF decision. For a specific prefix, the main SPF always selects a direct primary next-hop. The LFA SPF will select a an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.

Therefore the selection algorithm in Section 1.3 when SPF finds multiple LFA next-hops for a specific primary next-hop is modified as follows:

1. The algorithm splits the LFA next-hops into two sets:
 - the first set consists of direct LFA next-hops
 - the second set consists of tunneled LFA next-hops. after excluding the LSPs which use the same outgoing interface as the primary next-hop.
2. The algorithms continues with first set if not empty, otherwise it continues with second set.
3. If the second set is used, the algorithm selects the tunneled LFA next-hop which endpoint corresponds to the node advertising the prefix:
 - If more than one tunneled next-hop exists, it selects the one with the lowest LSP metric.
 - If still more than one tunneled next-hop exists, it selects the one with the lowest tunnel-id.
 - If none is available, it continues with rest of the tunneled LFAs in second set.
4. Within the selected set, the algorithm splits the LFA next-hops into two sets:
 - The first set consists of LFA next-hops which do not go over the PN used by primary next-hop.
 - The second set consists of LFA next-hops which go over the PN used by the primary next-hop.
5. If there is more than one LFA next-hop in the selected set, it will pick the node-protect type in favor of the link-protect type.
6. If there is more than one LFA next-hop within the selected type, then it will pick one based on the least total cost for the prefix. For a tunneled next-hop, it means the LSP metric plus the cost of the LSP endpoint to the destination of the prefix.
7. If there is more than one LFA next-hop within the selected type (ecmp-case) in the first set, it will select the first direct next-hop from the remaining set. This is not a deterministic selection and will vary following each SPF calculation.
8. If there is more than one LFA next-hop within the selected type (ecmp-case) in the second set, it will pick the tunneled next-hop with the lowest cost from the endpoint of the LSP to the destination prefix. If there remains more than one, it will pick the tunneled next-hop with the lowest tunnel-id.

3.2.5.2 Loop-Free Alternate calculation for inter-area/inter-level prefixes

When SPF resolves OSPF inter-area prefixes or IS-IS inter-level prefixes, it will compute an LFA backup next-hop to the same exit area/border router as used by the primary next-hop.

3.3 Loop-Free Alternate Shortest Path First (LFA SPF) policies

An LFA SPF policy allows the user to apply specific criteria, such as admin group and SRLG constraints, to the selection of a LFA backup next-hop for a subset of prefixes that resolve to a specific primary next-hop. The feature introduces the concept of route next-hop template to influence LFA backup next-hop selection.

3.3.1 Configuration of route next-hop policy template

The LFA SPF policy consists of applying a route next-hop policy template to a set of prefixes.

The user first creates a route next-hop policy template under the global router context:

configure>router>route-next-hop-policy>template *template-name*

A policy template can be used in both IS-IS and OSPF to apply the specific criteria described in the next subsections to prefixes protected by LFA. Each instance of IS-IS or OSPF can apply the same policy template to one or more prefix lists and to one or more interfaces.

The commands within the route next-hop policy use the **begin-commit-abort** model introduced with BFD templates. The following are the steps to create and modify the template:

- To create a template, the user enters the name of the new template directly under **route-next-hop-policy** context.
- 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.
- The user enters the editing mode by executing the **begin** command under **route-next-hop-policy** context. The user can then edit and change any number of route next-hop policy templates. However, the parameter value will still be stored temporarily in the template module until the **commit** is executed under the **route-next-hop-policy** context. Any temporary parameter changes will be lost if the user enters the **abort** command before the **commit** command.
- 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, the **abort** command if entered will have no effect on the prior deletion or creation of a template.

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

3.3.1.1 Configuring affinity or admin group constraint in route next-hop policy

Administrative groups (admin groups), also known as affinity, are used to tag IP 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, or all links which have bandwidth higher than 10G, or all links which are dedicated to a specific service.

The user first configures locally on each router the name and identifier of each admin group:

config>router>if-attribute>admin-group *group-name value group-value*

A maximum of 32 admin groups can be configured per system.

Next the user configures the admin group membership of the IP interfaces used in LFA. The user can apply admin groups to a network IP interface.

config>router> interface>if-attribute>admin-group *group-name [group-name...(up to 5 max)]*

The user can add as many admin groups as configured to a specific IP interface. The preceding command can be applied multiple times.

Note that 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.

The **no** form of the **admin-group** command under the interface deletes one or more of the admin-group memberships of the interface. It deletes all memberships if no group name is specified.

Finally, the user adds the admin group constraint into the route next-hop policy template:

configure router route-next-hop-template template *template-name*

include-group *group-name* [**pref** 1]

include-group *group-name* [**pref** 2]

exclude-group *group-name*

Each group is entered individually. The **include-group** statement instructs the LFA SPF selection algorithm to pick up a subset of LFA next-hops among the links which belong to one or more of the specified admin groups. A link which does not belong to at least one of the admin-groups is excluded. However, a link can still be selected if it belongs to one of the groups in a **include-group** statement but also belongs to other groups which are not part of any **include-group** statement in the route next-hop policy.

The **pref** option is used to provide a relative preference for the admin group to select. A lower preference value means that LFA SPF will first attempt to select a LFA backup next-hop which is a member of the corresponding admin group. If none is found, then the admin group with the next higher preference value is evaluated. If no preference is configured for a specific admin group name, then it is supposed to be the least preferred, that is, numerically the highest preference value.

When evaluating multiple **include-group** statements within the same preference, any link which belongs to one or more of the included admin groups can be selected as an LFA next-hop. There is no relative preference based on how many of those included admin groups the link is a member of.

The **exclude-group** statement prunes all links belonging to the specified admin group before making the LFA backup next-hop selection for a prefix.

If the same group name is part of both **include** and **exclude** statements, the **exclude** statement will win. In other words, the **exclude** statement can be viewed as having an implicit preference value of 0.

Note the admin-group criterion is applied before running the LFA next-hop selection algorithm. The modified LFA next-hop selection algorithm is shown in Section 7.5.

3.3.1.2 Configuring SRLG group constraint in route next-hop policy

Shared Risk Loss Group (SRLG) is used to tag IP interfaces which share a specific fate with the same identifier. For example, an SRLG group identifier could represent all links which use separate fibers but are carried in the same fiber conduit. If the conduit is accidentally cut, all the fiber links are cut which means all IP interfaces using these fiber links will fail. Therefore 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.

The user first configures locally on each router the name and identifier of each SRLG group:

configure>router>if-attribute>srlg-group *group-name* **value** *group-value*

A maximum of 1024 SRLGs can be configured per system.

Next the user configures the admin group membership of the IP interfaces used in LFA. The user can apply SRLG groups to a network IP interface.

config>router>interface>if-attribute>srlg-group *group-name* [*group-name...*(up to 5 max)]

The user can add a maximum of 64 SRLG groups to a specific IP interface. The same preceding command can be applied multiple times.

Note that 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.

The **no** form of the **srlg-group** command under the interface deletes one or more of the SRLG memberships of the interface. It deletes all SRLG memberships if no group name is specified.

Finally, the user adds the SRLG constraint into the route next-hop policy template:

configure router route-next-hop-template template template-name

srlg-enable

When this command is applied to a prefix, the LFA SPF will select a LFA next-hop, among the computed ones, which uses an outgoing interface that does not participate in any of the SRLGs of the outgoing interface used by the primary next-hop.

Note the SRLG and admin-group criteria are applied before running the LFA next-hop selection algorithm. The modified LFA next-hop selection algorithm is shown in Section 7.5.

3.3.1.3 Interaction of IP and MPLS admin group and SRLG

The LFA SPF policy feature generalizes the use of admin-group and SRLG to other types of interfaces. To that end, it is important that the new IP admin groups and SRLGs be compatible with the ones already supported in MPLS. The following rules are implemented:

- The definition of admin groups and SRLGs are moved under the new **config>router>if-attribute** context. When upgrading customers to the release which supports the feature, all user configured admin groups and SRLGs under **config>router>mpls** context will automatically be moved into the new context. The configuration of admin groups and SRLGs under the **config>router>mpls** context in CLI is deprecated.
- The binding of an MPLS interface to a group, that is, configuring membership of an MPLS interface in a group, continues to be performed under **config>router>mpls>interface** context.
- The binding of a local or remote MPLS interface to an SRLG in the SRLG database continues to be performed under the **config>router>mpls>srlg-database** context.
- The binding of an ISIS/OSPF interface to a group is performed in the **config>router>interface>if-attribute** context. This is used by ISIS or OSPF in route next-hop policies.
- Only the admin groups and SRLGs bound to an MPLS interface context or the SRLG database context are advertised in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF.

3.3.1.4 Configuring protection type and next-hop type preference in route next-hop policy template

The user can select if link protection or node protection is preferred in the selection of a LFA next-hop for all IP prefixes and LDP FEC prefixes to which a route next-hop policy template is applied. The default in SR OS implementation is node protection. The implementation will fall back to the other type if no LFA next-hop of the preferred type is found.

The user can also select if IP backup next-hop. The default in SR OS implementation is to prefer IP next-hop as only IP backup next-hop is supported on the 7210 SAS.

The following options are therefore added into the route next-hop policy template:

configure router route-nh-template template *template-name*

protection-type {link | node}

nh-type {ip | tunnel}

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

3.3.2 Application of route next-hop policy template to an interface

After the route next-hop policy template is configured with the desired policies, the user can apply it to all prefixes which primary next-hop uses a specific interface name. The following command achieves that:

config>router>isis>interface>lfa-policy-map route-nh-template *template-name*

config>router>ospf>area>interface>lfa-policy-map route-nh-template *template-name*

When a route next-hop policy template is applied to an interface in IS-IS, it is applied in both level 1 and level 2. When a route next-hop policy template is applied to an interface in OSPF, it is applied in all areas. However, the preceding CLI command in an OSPF interface context can only be executed under the area in which the specified interface is primary and then applied in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

If the user excluded the interface from LFA using the command **loopfree-alternate-exclude**, the LFA policy if applied to the interface has no effect.

Finally, if the user applied a route next-hop policy template to a loopback interface or to the system interface, the command will not be rejected but it will result in no action taken.

3.3.3 Excluding prefixes from LFA SPF

In the current 7210 SAS implementation, the user can exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

This feature adds the ability to exclude prefixes from a prefix policy which matches on prefixes or on IS-IS tags:

config>router>isis>loopfree-alternate-exclude prefix-policy *prefix-policy* [*prefix-policy..* up to 5]

config>router>ospf>loopfree-alternate-exclude prefix-policy *prefix-policy* [*prefix-policy..* up to 5]

Example

The prefix policy is configured as in the existing SR OS implementation:

```
config
  router
    policy-options
      [no] prefix-list prefix-list1
        prefix 10.225.16.0/24 prefix-length-
range 32-32
      [no] policy-statements prefix-policy1
        entry 10
        from
          prefix-list "prefix-
list1"
    exit
```

```
                                exit      action accept
                                exit
                                exit
                                default-action reject
exit
```

The default action of the preceding **loopfree-alternate-exclude** command when not explicitly specified by the user in the prefix policy is a "reject". Therefore, regardless of whether the user explicitly added the statement "default-action reject" to the prefix policy, a prefix that did not match an entry in the policy is accepted into LFA SPF.

3.3.4 Modification to LFA next-hop selection algorithm

This feature modifies the LFA next-hop selection algorithm. The SRLG and admin-group criteria are applied before running the LFA next-hop selection algorithm. That is, links which do not include one or more of the admin-groups in the include-group statements and links which belong to admin-groups which have been explicitly excluded using the exclude-group statement, and the links which belong to the SRLGs used by the primary next-hop of a prefix are first pruned.

This pruning applies only to IP next-hops. Tunnel next-hops can have the admin group or SRLG constraint applied to them under MPLS. For example, if a tunnel next-hop is using an outgoing interface which belongs to a specific SRLG ID, the user can enable the **srlg-frr** option under the **config>router>mpls** context to be sure the RSVP LSP FRR backup LSP will not use an outgoing interface with the same SRLG ID. A prefix which is resolved to a tunnel next-hop is protected by the RSVP FRR mechanism and not by the IP FRR mechanism. Similarly, the user can include or exclude admin-groups for the RSVP LSP and its FRR bypass backup LSP in MPLS context. The admin-group constraints will, however, be applied to the selection of the outgoing interface of both the LSP primary path and its FRR bypass backup path.

The following is the modified LFA selection algorithm which is applied to prefixes resolving to a primary next-hop which uses a specific route next-hop policy template.

- Split the LFA next-hops into two sets:
 - IP or direct next-hops.
 - Tunnel next-hops after excluding the LSPs which use the same outgoing interface as the primary next-hop.
- Prune the IP LFA next-hops which use the following links:
 - links which do not include one or more of the admin-groups in the include group statements in the route next-hop policy template.
 - links which belong to admin-groups which have been explicitly excluded using the **exclude-group** statement in the route next-hop policy template.
 - links which belong to the SRLGs used by the primary next-hop of a prefix.
- Continue with the set indicated in the **nh-type** value in the route next-hop policy template if not empty, otherwise continue with the other set.
- Within IP next-hop set:
 - prefer LFA next-hops which do not go over the Pseudo-Node (PN) used by the primary next-hop
 - Within selected subset prefer the node-protect type or the link-protect type according to the value of the **protection-type** option in the route next-hop policy template.

- Within the selected subset, select the best admin-groups according to the preference specified in the value of the **include-group** option in the route next-hop policy template.
- Within selected subset, select lowest **total cost** of a prefix.
- If same **total cost**, select lowest **router-id**.
- If same **router-id**, select lowest **interface-index**.
- Within tunnel next-hop set, select tunnel next-hops which endpoint corresponds to the node owning or advertising the prefix.
- Within selected subset, select the one with the lowest cost (lowest LSP metric).
- If same lowest cost, select tunnel with lowest tunnel-index.
 - If none is available, continue with rest of the tunnel LFA next-hop set.
 - Prefer LFA next-hops which do not go over the Pseudo-Node (PN) used by the primary next-hop.
 - Within selected subset prefer the node-protect type or the link-protect type according to the value of the **protection-type** in the route next-hop policy template.
 - Within selected subset, select lowest **total cost** of a prefix. For a tunnel next-hop, it means the LSP metric plus the cost of the LSP endpoint to the destination of the prefix.
 - If same **total cost**, select lowest **endpoint to destination cost**.
 - If same **endpoint to destination cost**, select lowest **router-id**.

3.4 Segment routing in shortest path forwarding

OSPF can be configured with segment routing in shortest path forwarding using the same procedures as those used to configure IS-IS. See [Segment routing in shortest path forwarding](#) in the IS-IS section for more information.

3.4.1 LFA protection using segment routing backup node SID



Note:

Backup node SID configuration is not supported on the 7210 SAS-K 2F6C4T or 7210 SAS-K 3SFP+ 8C. The 7210 SAS operates as the AGN node, and the 7750 SR router must be configured as the ABR with the backup node SID configured on it.

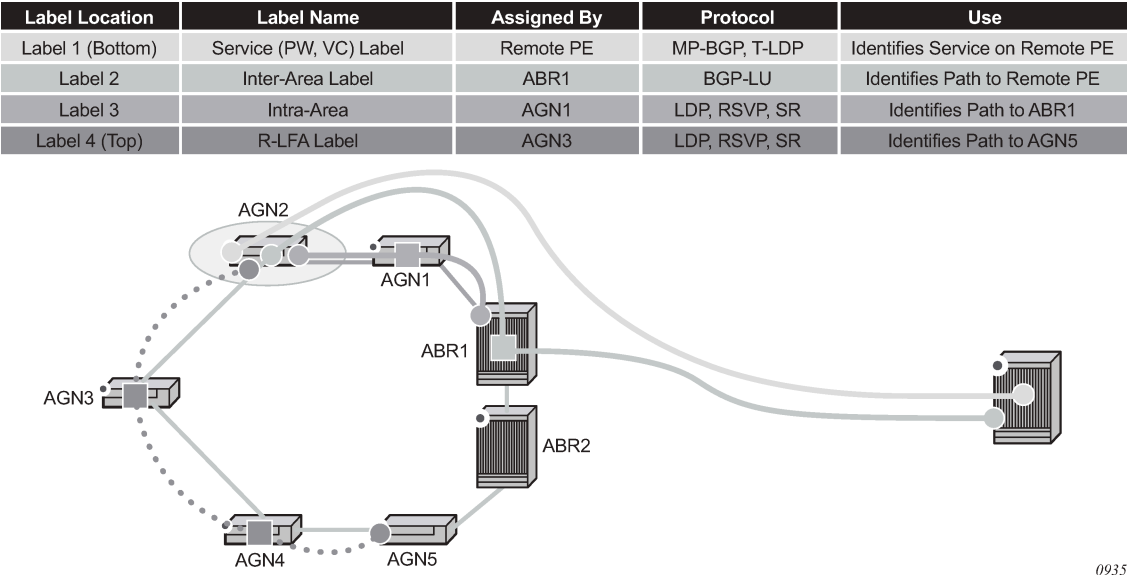
In MPLS deployments across multiple IGP areas or domains, such as in seamless MPLS design, it is challenging to provision FRR local protection in access and metro domains that use a ring, square, or partial mesh topology. To implement IP, LDP, or SR FRR in these topologies, the remote LFA feature must be implemented. Remote LFA provides a segment routing (SR) tunneled LFA next hop for an IP prefix, an LDP tunnel, or an SR tunnel. For prefixes outside of the area or domain, the access or aggregation router must push four labels: service label, BGP label for the destination PE, LDP/RSVP/SR label to reach the exit ABR/ASBR, and one label for the remote LFA next hop. Small routers deployed in these parts of the network have limited MPLS label stack size support.

[Figure 7: Label stack for remote LFA in ring topology](#) shows the label stack required for the primary next hop and the remote LFA next hop computed by aggregation node AGN2 for the inter-area prefix of a remote PE. For an inter-area BGP label unicast route prefix for which ABR1 is the primary exit ABR, AGN2 resolves the prefix to the transport tunnel of ABR1 and therefore, uses the remote LFA next hop of ABR1

for protection. The primary next hop uses two transport labels plus a service label. The remote LFA next hop for ABR1 uses PQ node AGN5 and pushes three transport labels plus a service label.

Seamless MPLS with Fast Restoration requires up to four labels to be pushed by AGN2, as shown in the following figure.

Figure 7: Label stack for remote LFA in ring topology



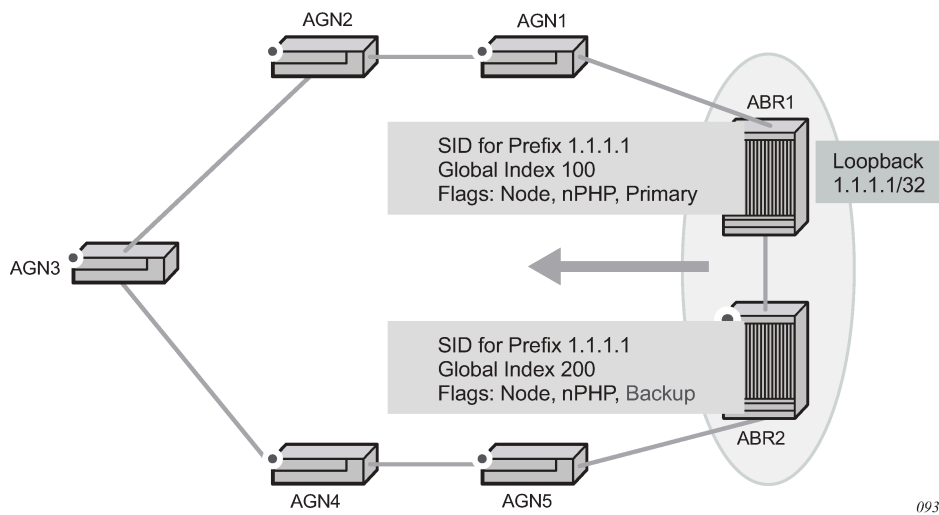
The objective of the LFA protection with a backup node SID feature is to reduce the label stack pushed by AGN2 for BGP label unicast inter-area prefixes. If link AGN2-AGN1 fails, packets are directed away from the failure and forwarded toward ABR2, which acts as the backup for ABR1 (and the other way around when ABR2 is the primary exit ABR for the BGP label unicast inter-area prefix). This requires ABR2 to advertise a special label for the loopback of ABR1 that will attract packets normally destined for ABR1. These packets are forwarded by ABR2 to ABR1 via the inter-ABR link.

As a result, AGN2 will push the label advertised by ABR2 to back up ABR1 in addition to the BGP label for the remote PE and the service label. This ensures that the label stack size for the LFA next hop is the same as that of the primary next hop. It is also the same size as the remote LFA next hop for the local prefix within the ring.

3.4.1.1 Detailed operation of LFA protection using backup node SID

- As shown in the following figure, LFA for seamless MPLS supports environments where the boundary routers are either:
- ABR nodes that connect with iBGP multiple domains, each using a different area of the same IGP instance
 - ASBR nodes that connect domains running different IGP instances and use iBGP within a domain and eBGP to the other domains

Figure 8: Backup ABR node SID



The following steps describe the configuration and behavior of LFA Protection using Backup Node SID, as shown in the preceding figure:

1. The user configures node SID 100 in ABR1 for its loopback prefix 1.1.1.1/32. This is the regular node SID. ABR1 advertises the prefix SID sub-TLV for this node SID in the IGP and installs the ILM using a unique label.
2. Each router receiving the prefix sub-TLV for node SID 100 resolves it as described in [Segment routing in shortest path forwarding](#). Changes to the programming of the backup NHLFE of node SID 100 based on receiving the backup node SID for prefix 1.1.1.1/32 are defined in [Duplicate SID handling](#).
3. The user configures a backup node SID 200 in ABR2 for the loopback 1.1.1.1/32 of ABR1. The SID value must be different from that assigned by ABR1 for the same prefix. ABR2 installs the ILM, which performs a swap operation from the label of SID 200 to that of SID 100. The ILM must point to a direct link and next hop to reach 1.1.1.1/32 of ABR1 as its primary next hop. The IGP examines all adjacencies established in the same area as that of prefix 1.1.1.1/32 and determines which ones have ABR1 as a direct neighbor and with the best cost. If more than one adjacency has the best cost, the IGP selects the one with the lowest interface index. If there is no adjacency to reach ABR2, the prefix SID for the backup node is flushed and is not resolved. This prevents the use of any non-direct path to reach ABR1. As a result, any received traffic on the ILM of SID 200 traffic will be blackholed.
4. If resolved, ABR2 advertises the prefix SID sub-TLV for this backup node SID 200 and indicates in the SR Algorithm field that a modified SPF algorithm, referred to as "Backup-constrained-SPF", is required to resolve this node SID.
5. Each router receiving the prefix sub-TLV for the backup node SID 200 performs the following resolution steps. These steps do not require a CLI command to be enabled:
 - a. The router determines which router is being backed up. This is achieved by checking the router ID owner of the prefix sub-TLV that was advertised with the same prefix but without the backup flag and which is used as the best route for the prefix. In this case, it should be ABR1. Then the router runs a modified SPF by removing node ABR1 from the topology to resolve the backup node SID 200. The primary next hop should point to the path to ABR2 in the counter clockwise direction of the ring. The router will not compute an LFA or a remote LFA for node SID 200 because the main SPF used a modified topology.

- b. The router installs the ILM and primary NHLFE for the backup node SID.

Only a swap label operation is configured by all routers for the backup node SID. There is no push operation, and no tunnel for the backup node SID is added into the TTM.

- c. The router programs the backup node SID as the LFA backup for the SR tunnel to node SID of 1.1.1.1/32 of ABR1. In other words, each router overrides the remote LFA backup for prefix 1.1.1.1/32, which is normally PQ node AGN5.

- d. If the router is adjacent to ABR1, for example AGN1, it also programs the backup node SID as the LFA backup for the protection of any adjacency SID to ABR1.

6. When node AGN2 resolves a BGP label route for an inter-area prefix for which the primary ABR exit router is ABR1, it will use the backup node SID of ABR1 as the remote LFA backup instead of the SID to the PQ node (AGN5 in this example) to save on the pushed label stack.

AGN2 continues to resolve the prefix SID for any remote PE prefix that is summarized into the local area of AGN2 as usual. AGN2 programs a primary next hop and a remote LFA next hop. Remote LFA will use AGN5 as the PQ node and will push two labels, as it would for an intra-area prefix SID. There is no need to use the backup node SID for this prefix SID and force its backup path to go to ABR1. The backup path may exit from ABR2 if the cost from ABR2 to the destination prefix is shorter.

7. If the user excludes a link from LFA in the IGP instance (**config>router>ospf>area>interface>loopfree-alternate-exclude**), a backup node SID that resolves to that interface will not be used as a remote LFA backup in the same way as regular LFA or PQ remote LFA next hop behavior.
8. If the OSPF neighbor of a router is put into overload or if the metric of an OSPF interface to that neighbor is set to LSInfinity (0xFFFF), a backup node SID that resolves to that neighbor will not be used as a remote LFA backup in the same way as regular LFA or PQ remote LFA next hop behavior.
9. LFA policy is supported for IP next hops only. It is not supported with tunnel next hops such as IGP shortcuts or remote LFA tunnels. A backup node SID is also a tunnel next hop and, therefore, a user-configured LFA policy is not applied to check constraints such as admin-groups and SRLG against the outgoing interface of the selected backup node SID.

3.4.1.2 Duplicate SID handling

If the IGP issues or receives an LSA/LSP containing a prefix SID sub-TLV for a node SID or a backup node SID with a SID value that is a duplicate of an existing SID or backup node SID, the resolution in the following table is followed.

Table 24: Handling of duplicate SIDs

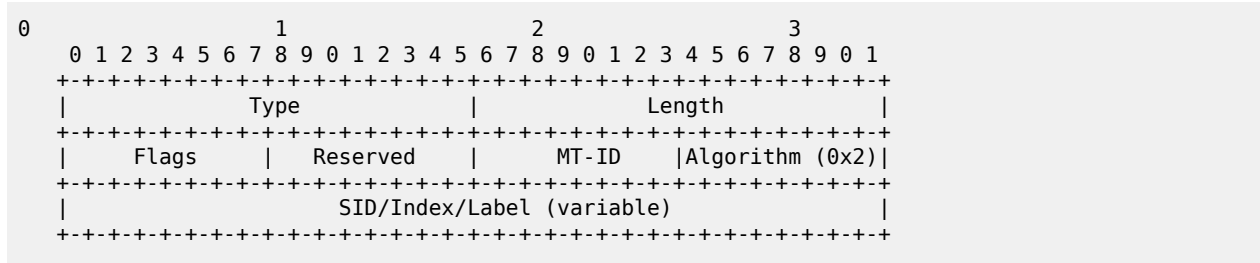
Old LSA/LSP	New LSA/LSP			
	Backup node SID	Local backup node SID	Node SID	Local node SID
Backup Node SID	Old	New	New	New
Local Backup Node SID	Old	Equal	New	New
Node SID	Old	Old	Equal/Old ⁹	Equal/New ¹⁰

Old LSA/LSP	New LSA/LSP			
	Backup node SID	Local backup node SID	Node SID	Local node SID
Local Node SID	Old	Old	Equal/Old ⁹	Equal/Old ⁹

3.4.1.3 OSPF control plane extensions

All routers supporting OSPF control plane extensions must advertise support of the new algorithm “Backup-constrained-SPF” of value 2 in the SR-Algorithm TLV, which is advertised in the Router Information Opaque LSA. This is in addition to the default supported algorithm “IGP-metric-based-SPF” of value 0. The following shows the encoding of the prefix SID sub-TLV to indicate a node SID of type backup and to indicate the modified SPF algorithm in the SR Algorithm field. The values used in the Flags field and in the Algorithm field are SR OS proprietary.

The new Algorithm (0x2) field and values are used by this feature.




The following table lists OSPF control plane extension flag values.

Table 25: OSPF control plane extension fields

Field	Value
Type	2
Length	variable
Flags	1 octet field

The following flags are defined; the “B” flag is new:



⁹ Equal/Old means the following:

- If the prefix is duplicate, it is equal and no change is needed. Keep the old LSA/LSP.
- If the prefix is not duplicate, still keep the old LSA/LSP.

¹⁰ Equal/New means the following:

- If the prefix is duplicate, it is equal and no change is needed. Keep the old LSA/LSP.
- If the prefix is not duplicate, pick a new prefix and use the new LSA/LSP.

```
| |NP|M |E |V |L |B| |
+---+---+---+---+---+---+
```

The following table describes OSPF control plane extension flags.

Table 26: OSPF control plane extension flags

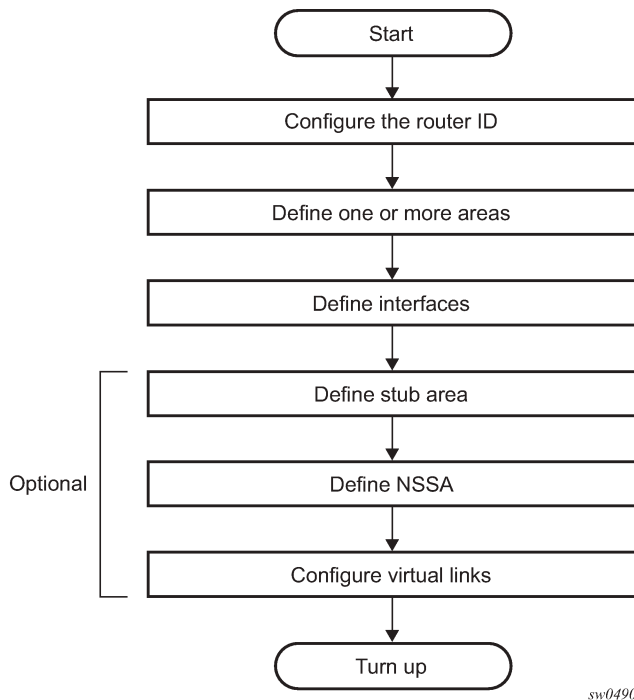
Flag	Description
NP-Flag	No-PHP flag If set, the penultimate hop must not pop the prefix SID before delivering the packet to the node that advertised the prefix SID.
M-Flag	Mapping Server Flag If set, the SID is advertised from the Segment Routing Mapping Server functionality as described in <i>I-D.filsfils-spring-segment-routing-ldp-interop</i> .
E-Flag	Explicit-Null Flag If set, any upstream neighbor of the prefix SID originator must replace the prefix SID with a prefix SID having an Explicit-NULL value (0 for IPv4) before forwarding the packet.
V-Flag	Value/Index Flag If set, the prefix SID carries an absolute value. If not set, the prefix SID carries an index.
L-Flag	Local/Global Flag If set, the value/index carried by the prefix SID has local significance. If not set, then the value/index carried by this sub-TLV has global significance.
B-Flag	This flag is used by the Protection using backup node SID feature. If set, the SID is a backup SID for the prefix. This value is SR OS proprietary.
Other bits	Reserved These must be zero when sent and are ignored when received.
MT-ID	Multi-Topology ID, as defined in RFC 4915
Algorithm	One octet identifying the algorithm the prefix SID is associated with. A value of (0x2) indicates the modified SPF algorithm, which removes from the topology the node that is backed up by the backup node SID. This value is SR OS proprietary.
SID/Index/Label	Based on the V and L flags, it contains either: <ul style="list-style-type: none"> a 32-bit index defining the offset in the SID/Label space advertised by this router

Flag	Description
	<ul style="list-style-type: none">a 24-bit label where the 20 rightmost bits are used for encoding the label value

3.5 OSPF configuration process overview

The following figure shows the process to provision basic OSPF parameters.

Figure 9: OSPF configuration and implementation flow



3.6 Configuration notes

This section describes OSPF configuration caveats.

3.6.1 General

- Before OSPF can be configured, the router ID must be configured.
- The basic OSPF configuration includes at least one area and an associated interface.
- All default and command parameters can be modified.

3.6.1.1 OSPF defaults

The following list summarizes the OSPF configuration defaults:

- By default, a router has no configured areas.
- An OSPF instance is created in the administratively enabled state.

3.7 Configuring OSPF with CLI

This section provides information to configure Open Shortest Path First (OSPF) using the command line interface.

3.8 OSPF configuration guidelines

Configuration planning is essential to organize routers, backbone, non-backbone, stub, NSSA areas, and transit links. OSPF provides essential defaults for basic protocol operability. You can configure or modify commands and parameters. OSPF is not enabled by default.

The minimal OSPF parameters which should be configured to deploy OSPF are:

- **Router ID**

Each router running OSPF must be configured with a unique router ID. The router ID is used by OSPF routing protocols in the routing table manager.

When configuring a new router ID, protocols will not automatically be restarted with the new router ID. Shut down and restart the protocol to initialize the new router ID.

- **OSPF instance**

OSPF instances must be defined when configuring multiple instances and/or the instance being configured is not the base instance.

- **An area**

At least one OSPF area must be created. An interface must be assigned to each OSPF area.

- **Interfaces**

An interface is the connection between a router and one of its attached networks. An interface has state information associated with it, which is obtained from the underlying lower level protocols and the routing protocol. An interface to a network has associated with it a single IP address and mask (unless the network is an unnumbered point-to-point network). An interface is sometimes also referred to as a link.

3.9 Basic OSPF configuration

This section provides information to configure the basic parameter for OSPF and OSPFv3 as well as configuration examples of common configuration tasks.

The minimal OSPF parameters that need to be configured are:

- **router ID**

If a *router-id* is not configured in the **config>router** context, the router's system interface IP address is used.

- one or more areas
- interfaces (interface "system")

Example: Basic OSPF configuration output

```
SAS-A>config>router>ospf# info
-----
      area 0.0.0.0
        interface "system"
        exit
      exit
      area 0.0.0.20
        nssa
        exit
        interface "to-104"
        priority 10
        exit
      exit
      area 0.0.1.1
      exit
-----
SAS-A>config>router>ospf#
```

Example: Basic OSPFv3 configuration output

```
SAS-A>config>router>ospf3# info
-----
      asbr
      overload
        lsa-arrival 50000
        exit
      exit
      export "OSPF-Export"
      area 0.0.0.0
        interface "system"
        exit
      exit
      area 0.0.1.20
        nssa
        exit
        interface "SR1-2"
        exit
      exit
-----
SAS-A>config>router>ospf3#
```

3.9.1 Configuring the router ID

The router ID uniquely identifies the router within an AS. In OSPF, routing information is exchanged between autonomous systems, groups of networks that share routing information. It can be set to be the same as the loopback (system interface) address. Subscriber services also use this address as far-end

router identifiers when service distribution paths (SDPs) are created. The router ID is used by both OSPF and BGP routing protocols. A router ID can be derived by:

- Defining the value in the **config>router** *router-id* context.
- Defining the system interface in the **config>router>interface** *ip-int-name* context (used if the router ID is not specified in the **config>router** *router-id* context).
- Inheriting the last four bytes of the MAC address.
- 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.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is (re) 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.

Example: Router ID configuration output

```
A:ALA-B>config>router# info
#-----
# IP Configuration
#-----
        interface "system"
            address 10.10.10.104/32
        exit
        interface "to-103"
            address 10.0.0.104/24
            port 1/1/1
        exit
router-id 10.10.10.104
...
#-----
A:ALA-B>config>router#
```

3.10 Configuring OSPF components

The following section describes the syntax used to configure the OSPF components.

3.10.1 Configuring OSPF parameters

Example: Basic OSPF configuration output

```
A:SAS-12>config>router>ospf# info
-----
        asbr
        overload
        overload-on-boot timeout 60
        traffic-engineering
        export "OSPF-Export"
        graceful-restart
            helper-disable
        exit
-----
A:SAS-12>config>router>ospf# ex
```

3.10.2 Configuring OSPFv3 parameters

Use the following syntax to configure OSPF3 parameters.

```
config>router# ospf3
  asbr
  export policy-name [policy-name...(up to 5 max)]
  external-db-overflow limit seconds
  external-preference preference
  overload [timeout seconds]
  overload-include-stub
  overload-on-boot [timeout seconds]
  preference preference
  reference-bandwidth bandwidth-in-kbps
  router-id ip-address
  no shutdown
  timers
    lsa-arrival lsa-arrival-time
    lsa-generate max-lsa-wait [lsa-initial-wait lsa-initial-wait [lsa-second-wait lsa-second-wait]]
    spf-wait max-spf-wait [spf-initial-wait spf-initial-wait [spf-second-wait spf-second-wait]]
```

Example: OSPF3 configuration output

```
A:SAS-12>config>router>ospf3# info
-----
      asbr
      overload
      timers
        lsa-arrival 50000
      exit
      export "OSPF-Export"
-----
A:SAS-12>config>router>ospf3#
```

OSPF also supports the concept of multi-instance OSPFv2 and OSPFv3 which allows separate instances of the OSPF protocols to run independently within SR OSs.

Separate instances are created by adding a different instance ID as the optional parameter to the **config>router>ospf** and **config>router>ospf3** commands. When this is done a separate OSPF instance is created which maintains separate link state databases for each instance.

3.10.3 Configuring an OSPF and OSPFv3 area

An OSPF area consists of routers configured with the same area ID. To include a router in a specific area, the common area ID must be assigned and an interface identified.

If a network consists of multiple areas you must also configure a backbone area (0.0.0.0) on at least one router. The backbone is comprised of the area border routers and other routers not included in other areas. The backbone distributes routing information between areas. The backbone is considered to be a participating area within the autonomous system. To maintain backbone connectivity, there must be at least one interface in the backbone area or have a virtual link configured to another router in the backbone area.

The minimal configuration must include an area ID and an interface. Modifying other command parameters are optional.

Use the following syntax to configure an OSPF area.

```
ospf ospf-instance
  area area-id
    area-range ip-prefix/mask [advertise|not-advertise]
    blackhole-aggregate
```

Use the following syntax to configure an OSPFv3 area.

```
ospf ospf-instance
  ospf3
    area area-id
      area-range ip-prefix/mask [advertise|not-advertise]
      blackhole-aggregate
```

Example: OSPF area configuration output

```
A:ALA-A>config>router>ospf# info
-----
      area 0.0.0.0
      exit
      area 0.0.0.20
      exit
-----
ALA-A>config>router>ospf#A:
```

3.10.4 Configuring an OSPF and OSPFv3 stub area

Configure stub areas to control external advertisements flooding and to minimize the size of the topological databases on an area's routers. A stub area cannot also be configured as an NSSA.

By default, summary route advertisements are sent into stub areas. The **no** form of the summary command disables sending summary route advertisements and only the default route is advertised by the ABR. This example retains the default so the command is not entered.

If this area is configured as a transit area for a virtual link, then existing virtual links of a non-stub or NSSA area are removed when its designation is changed to NSSA or stub.

Stub areas for OSPFv3 are configured the same as for OSPF.

Use the following syntax to configure an OSPF stub area.

```
ospf
  area area-id
  stub
  default-metric metric
  summaries
```

The following is a sample stub configuration output.

```
SAS-12>config>router>ospf>area># info
-----
...
      area 0.0.0.0
      exit
```

```

        area 0.0.0.20
    stub
        exit
    exit
...
-----
SAS-12>config>router>ospf#

```

3.10.5 Configuring a Not-So-Stubby Area

You must explicitly configure an area to be a Not-So-Stubby Area (NSSA) area. NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is an NSSA has the capability to flood external routes it learns throughout its area and by an area border router to the entire OSPF domain. An area cannot be both a stub area and an NSSA.

If this area is configured as a transit area for a virtual link, then existing virtual links of a non-stub or NSSA area are removed when its designation is changed to NSSA or stub.

Use the following syntax to configure an NSSA for OSPF.

```

ospf ospf-instance
    area area-id
    nssa
        area-range ip-prefix/mask [advertise|not-advertise]
        originate-default-route [type-7]
        redistribute-external
        summaries

```

Use the following syntax to configure stub areas for the OSPFv3.

```

ospf ospf-instance
    ospf3
        area area-id
        nssa
            area-range ip-prefix/mask [advertise|not-advertise]
            originate-default-route [type-7]
            redistribute-external
            summaries

```

Example: NSSA configuration output

```

A:SAS-12>config>router>ospf# info
-----
    asbr
    overload
    overload-on-boot timeout 60
    traffic-engineering
    export "OSPF-Export"
    exit
    area 0.0.0.0
    exit
    area 0.0.0.20
        stub
        exit
    exit
    area 0.0.0.25

```

```

        nssa
        exit
    exit
-----
A:SAS-12>config>router>ospf#

```

Example: OSPFv3 NSSA configuration output

```

A:SAS-12>config>router>ospf3# info
-----
    asbr
    overload
    timers
        lsa-arrival 50000
    exit
    export "OSPF-Export"
    area 0.0.0.0
    exit
    area 0.0.0.20
        stub
        exit
    exit
    area 0.0.0.25
        nssa
        exit
    exit
    area 4.3.2.1
    exit
-----
A:SAS-12>config>router>ospf3#

```

3.10.6 Configuring a virtual link

The backbone area (area 0.0.0.0) must be contiguous and all other areas must be connected to the backbone area. If it is not practical to connect an area to the backbone then the area border routers must be connected via a virtual link. The two area border routers will form a point-to-point-like adjacency across the transit area. A virtual link can only be configured while in the area 0.0.0.0 context.

The **router-id** parameter specified in the **virtual-link** command must be associated with the virtual neighbor, that is, enter the virtual neighbor's router ID, not the local router ID. The transit area cannot be a stub area or an NSSA.

Use the following syntax to configure a virtual link.

```

ospf ospf-instance
    area area-id
    virtual-link router-id transit-area area-id
        authentication-key [authentication-key|hash-key] [hash]
        authentication-type [password|message-digest]
        dead-interval seconds
        hello-interval seconds
        message-digest-key key-id md5 [key|hash-key] [hash|hash2]
        retransmit-interval seconds
        transit-delay
        no shutdown

```

Example: Virtual link configuration output

```
A:SAS-12>config>router>ospf# info
-----
asbr
overload
overload-on-boot timeout 60
traffic-engineering
export "OSPF-Export"
graceful-restart
    helper-disable
exit
area 0.0.0.0
    virtual-link 1.2.3.4 transit-area 1.2.3.4
        hello-interval 9
        dead-interval 40
    exit
exit
area 0.0.0.20
    stub
    exit
exit
area 0.0.0.25
    nssa
    exit
exit
area 1.2.3.4
    exit
-----
A:SAS-12>config>router>ospf#
```

3.10.7 Configuring an interface

In OSPF, an interface can be configured to act as a connection between a router and one of its attached networks. An interface includes state information that was obtained from underlying lower level protocols and from the routing protocol. An interface to a network is associated with a single IP address and mask. If the address is merely changed, then the OSPF configuration is preserved.

The **passive** command enables the passive property to and from the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol. By default, only interface addresses that are configured for OSPF are advertised as OSPF interfaces. The passive parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol. When enabled, the interface will ignore ingress OSPF protocol packets and not transmit any OSPF protocol packets.

An interface can be part of more than one area, as specified in RFC 5185. To do this, add the keyword **secondary** when creating the interface.

Use the following syntax to configure an OSPF interface.

```
ospf ospf-instance
    area area-id
interface ip-int-name
    advertise-subnet
    authentication-key [authentication-key|hash-key] [hash|hash2]
    authentication-type [password|message-digest]
    dead-interval seconds
    hello-interval seconds
    interface-type {broadcast|point-to-point}
    message-digest-key key-id md5 [key|hash-key][hash|hash2]
```

```
metric metric
mtu bytes
passive
priority number
retransmit-interval seconds
no shutdown
transit-delay seconds
```

The following is a sample interface configuration output.

```
A:ALA-49>config>router>ospf# info
-----
    asbr
    overload
    overload-on-boot timeout 60
    traffic-engineering
    export "OSPF-Export"
    graceful-restart
        helper-disable
    exit
    area 0.0.0.0
        virtual-link 1.2.3.4 transit-area 1.2.3.4
            hello-interval 9
            dead-interval 40
        exit
        interface "system"
        exit
    exit
    area 0.0.0.20
        stub
        exit
        interface "to-103"
        exit
    exit
    area 0.0.0.25
        nssa
        exit
    exit
    area 1.2.3.4
    exit
    area 4.3.2.1
        interface "SR1-3"
        exit
    exit
    area 4.3.2.1
        interface "SR1-3" secondary
        exit
    exit
-----
A:ALA-49>config>router>ospf# area 0.0.0.20
```

3.10.8 Configuring authentication

Authentication must be explicitly configured. The following authentication commands can be configured on the interface level or the virtual link level:

- **authentication-key**

Configures the password used by the OSPF interface or virtual-link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.

- **authentication-type**

Enables authentication and specifies the type of authentication to be used on the OSPF interface, either password or message digest.

- **message-digest-key**

Use this command when the **message-digest** keyword is selected in the **authentication-type** command. The Message Digest 5 (MD5) hashing algorithm is used for authentication. MD5 is used to verify data integrity by creating a 128-bit message digest from the data input. It is unique to that specific data.

An special checksum is included in transmitted packets and are used by the far-end router to verify the packet by using an authentication key (a password). Routers on both ends must use the same MD5 key.

MD5 can be configured on each interface and each virtual link. If MD5 is enabled on an interface, then that interface accepts routing updates only if the MD5 authentication is accepted. Updates that are not authenticated are rejected. A router accepts only OSPF packets sent with the same **key-id** value defined for the interface.

When the hash parameter is not used, non-encrypted characters can be entered. When configured using the **message-digest-key** command, then all keys specified in the command are stored in encrypted format in the configuration file using the **hash** keyword. When using the **hash** keyword the password must be entered in encrypted form. Hashing cannot be reversed. Issue the **no message-digest-key key-id** command and then re-enter the command without the **hash** parameter to configure an unhashed key.

The following CLI commands are displayed to illustrate the key authentication features. These command parameters can be defined at the same time interfaces and virtual-links are being configured. See [Configuring an interface](#) and [Configuring a virtual link](#).

Use the following syntax to configure authentication.

```
ospf ospf-instance
  area area-id
  interface ip-int-name
    authentication-key [authentication-key|hash-key] [hash]
    authentication-type [password|message-digest]
    message-digest-key key-id md5 key[hash]
  virtual-link router-id transit-area area-id
    authentication-key [authentication-key|hash-key] [hash]
    authentication-type [password|message-digest]
    message-digest-key key-id md5 key[hash]
```

Example: Authentication configuration output

```
A:ALA-49>config>router>ospf# info
-----
asbr
overload
overload-on-boot timeout 60
traffic-engineering
export "OSPF-Export"
graceful-restart
  helper-disable
exit
area 0.0.0.0
  virtual-link 1.2.3.4 transit-area 1.2.3.4
    hello-interval 9
    dead-interval 40
  exit
interface "system"
```



```
        exit
    exit
    area 0.0.0.20
        stub
    exit
    interface "to-103"
    exit
exit
area 0.0.0.25
    nssa
    exit
exit
area 0.0.0.40
    interface "test1"
        authentication-type password
        authentication-key "3WErEDozxyQ" hash
    exit
exit
area 1.2.3.4
exit
-----
A:ALA-49>config>router>ospf#
```

```
A:ALA-49>config>router>ospf# info
-----
asbr
overload
overload-on-boot timeout 60
traffic-engineering
export "OSPF-Export"
graceful-restart
    helper-disable
exit
area 0.0.0.0
    virtual-link 10.0.0.1 transit-area 0.0.0.1
        authentication-type message-digest
        message-digest-key 2 md5 "Mi6BQAFi3MI" hash
    exit
    virtual-link 1.2.3.4 transit-area 1.2.3.4
        hello-interval 9
        dead-interval 40
    exit
    interface "system"
    exit
exit
area 0.0.0.1
exit
area 0.0.0.20
    stub
    exit
    interface "to-103"
    exit
exit
area 0.0.0.25
    nssa
    exit
exit
area 0.0.0.40
    interface "test1"
        authentication-type password
        authentication-key "3WErEDozxyQ" hash
    exit
exit
```

```
        area 1.2.3.4
        exit
-----
A:ALA-49>config>router>ospf#
```

3.10.9 Assigning a designated router

A designated router is elected according to the priority number advertised by the routers. When a router starts up, it checks for a current designated router. If a designated router is present, then the router accepts that designated router, regardless of its own priority designation. When a router fails, then new designated and backup routers are elected according to their priority numbers.

The **priority** command is only used if the interface is a broadcast type. The designated router is responsible for flooding network link advertisements on a broadcast network to describe the routers attached to the network. A router uses hello packets to advertise its priority. The router with the highest priority interface becomes the designated router. A router with priority 0 is not eligible to be a designated router or a backup designated router. At least one router on each logical IP network or subnet must be eligible to be the designated router. By default, routers have a priority value of 1.

Use the following syntax to configure the designated router.

```
ospf ospf-instance
    area area-id
    interface ip-int-name
        priority number
```

Example: Priority designation output

```
A:ALA-49>config>router>ospf# info
-----
    asbr
    overload
    overload-on-boot timeout 60
    traffic-engineering
    export "OSPF-Export"
    graceful-restart
        helper-disable
    exit
    area 0.0.0.0
        virtual-link 10.0.0.1 transit-area 0.0.0.1
            authentication-type message-digest
            message-digest-key 2 md5 "Mi6BQAFi3MI" hash
        exit
        virtual-link 1.2.3.4 transit-area 1.2.3.4
            hello-interval 9
            dead-interval 40
        exit
        interface "system"
        exit
    exit
    area 0.0.0.1
    exit
    area 0.0.0.20
        stub
        exit
        interface "to-103"
        exit
    exit
    area 0.0.0.25
```

```

        nssa
        exit
        interface "if2"
            priority 100
        exit
    exit
    area 0.0.0.40
        interface "test1"
            authentication-type password
            authentication-key "3WErEDozxyQ" hash
        exit
    exit
    area 1.2.3.4
    exit
-----
A:ALA-49>config>router>ospf#

```

3.10.10 Configuring route summaries

Area border routers send summary (type 3) advertisements into a stub area or NSSA to describe the routes to other areas. This command is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub or NSSA.

By default, summary route advertisements are sent into the stub area or NSSA. The **no** form of the **summaries** command disables sending summary route advertisements and, in stub areas, the default route is advertised by the area border router.

The following CLI commands are displayed to illustrate route summary features. These command parameters can be defined at the same time stub areas and NSSAs are being configured. See [Configuring an OSPF and OSPFv3 stub area](#) and [Configuring a Not-So-Stubby Area](#).

Route summaries for OSPF3 are configured the same as for OSPF.

Use the following syntax to configure a route summary.

```

ospf ospf-instance
    area area-id
    stub
        summaries
    nssa
        summaries

```

Example: Stub route summary configuration output

```

A:SAS-12>config>router>ospf# info
-----
    asbr
    overload
    overload-on-boot timeout 60
    traffic-engineering
    export "OSPF-Export"
    graceful-restart
        helper-disable
    exit
    area 0.0.0.0
        virtual-link 10.0.0.1 transit-area 0.0.0.1
        authentication-type message-digest
        message-digest-key 2 md5 "Mi6BQAFi3MI" hash

```

```
        exit
        virtual-link 1.2.3.4 transit-area 1.2.3.4
            hello-interval 9
            dead-interval 40
        exit
        interface "system"
        exit
    exit
    area 0.0.0.1
    exit
    area 0.0.0.20
        stub
        exit
        interface "to-103"
        exit
    exit
    area 0.0.0.25
        nssa
        exit
        interface "if2"
            priority 100
        exit
    exit
    area 0.0.0.40
        interface "test1"
            authentication-type password
            authentication-key "3WErEDozxyQ" hash
        exit
    exit
    area 1.2.3.4
    exit
-----
A:SAS-12>config>router>ospf#
```

Example: Stub route summary OSPv3 configuration output

```
A:SAS-12>config>router>ospf3# info
-----
    asbr
    overload
    timers
        lsa-arrival 50000
    exit
    export "OSPF-Export"
    area 0.0.0.0
        virtual-link 4.3.2.1 transit-area 4.3.2.1
        exit
        interface "system"
        exit
    exit
    area 0.0.0.20
        stub
        exit
        interface "SR1-2"
        exit
    exit
    area 0.0.0.25
        nssa
        exit
    exit
    area 4.3.2.1
    exit
```

```
-----  
A:SAS-12>config>router>ospf3#
```

3.10.11 Configuring route preferences

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs the preference value is used to decide which route is installed in the forwarding table if several protocols calculate routes to the same destination. The route with the lowest preference value is selected.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as described in the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

Table 27: Route preference defaults by route type

Route type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF/ OSPFv3 internal	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF/ OSPFv3 external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes



Note:

Preference for OSPF and OSPFv3 internal routes is configured with the **preference** command.

The following CLI commands are displayed to illustrate route preference features. The command parameters can be defined at the same time you are configuring OSPF. See [Configuring OSPF components](#).

Route parameters for OSPFv3 are configured the same as for OSPF.

Use the following syntax to configure a route preference.

```
ospf ospf-instance  
    preference preference  
    external-preference preference
```

Example: Route preference configuration output

```
A:ALA-49>config>router>ospf# info
```

```
-----
asbr
overload
overload-on-boot timeout 60
traffic-engineering
preference 9
external-preference 140
export "OSPF-Export"
graceful-restart
    helper-disable
exit
area 0.0.0.0
    virtual-link 10.0.0.1 transit-area 0.0.0.1
        authentication-type message-digest
        message-digest-key 2 md5 "Mi6BQAFi3MI" hash
    exit
    virtual-link 1.2.3.4 transit-area 1.2.3.4
        hello-interval 9
        dead-interval 40
    exit
    interface "system"
    exit
exit
area 0.0.0.1
exit
area 0.0.0.20
    stub
    exit
    interface "to-103"
    exit
exit
area 0.0.0.25
    nssa
    exit
    interface "if2"
        priority 100
    exit
exit
area 0.0.0.40
    interface "test1"
        authentication-type password
        authentication-key "3WErEDozxyQ" hash
    exit
exit
area 1.2.3.4
exit
-----
```

3.11 OSPF configuration management tasks

This section describes the OSPF configuration management tasks.

3.11.1 Modifying a router ID

Because the router ID is defined in the **config>router** context, not in the OSPF configuration context, the protocol instance is not aware of the change. Re-examine the plan detailing the router ID. Changing the router ID on a device could cause configuration inconsistencies if associated values are not also modified.

After you have changed a router ID, manually shut down and restart the protocol using the **shutdown** and **no shutdown** commands in order for the changes to be incorporated.

Use the following syntax to change a router ID number.

```
config>router# router-id router-id
```

Example: NSSA router ID modification output

```
A:ALA-49>config>router# info
-----
IP Configuration
-----
      interface "system"
        address 10.10.10.104/32
      exit
      interface "to-103"
        address 10.0.0.103/24
        port 1/1/1
      exit
router-id 10.10.10.104
-----
A:ALA-49>config>router#
```

3.11.2 Deleting a router ID

You can modify a router ID, but you cannot delete the parameter. When the **no router router-id** command is issued, the router ID reverts to the default value, the system interface address (which is also the loopback address). If a system interface address is not configured, then the last 32 bits of the chassis MAC address is used as the router ID.

3.11.3 Modifying OSPF parameters

You can change or remove existing OSPF parameters in the CLI or NMS. The changes are applied immediately.

Example

The following is a sample OSPF modification in which an interface is removed and another interface added.

```
config>router# ospf 1
config>router>ospf# area 0.0.0.20
config>router>ospf>area# no interface "to-103"
config>router>ospf>area# interface "to-HQ"
config>router>ospf>area>if$ priority 50
config>router>ospf>area>if# exit
config>router>ospf>area# exit
```

Example

The following is a sample OSPF configuration output with the modifications entered in the previous sample.

```
A:ALA-49>config>router>ospf# info
```

```
-----
asbr
overload
overload-on-boot timeout 60
traffic-engineering
preference 9
external-preference 140
export "OSPF-Export"
graceful-restart
    helper-disable
exit
area 0.0.0.0
    virtual-link 10.0.0.1 transit-area 0.0.0.1
        authentication-type message-digest
        message-digest-key 2 md5 "Mi6BQAFi3MI" hash
    exit
    virtual-link 1.2.3.4 transit-area 1.2.3.4
        hello-interval 9
        dead-interval 40
    exit
    interface "system"
    exit
exit
area 0.0.0.1
exit
area 0.0.0.20
    stub
    exit
    interface "to-HQ"
        priority 50
    exit
exit
area 0.0.0.25
    nssa
    exit
    interface "if2"
        priority 100
    exit
exit
area 0.0.0.40
    interface "test1"
        authentication-type password
        authentication-key "3WErEDozxyQ" hash
    exit
exit
area 1.2.3.4
exit
-----
A:ALA-49>config>router>ospf#
```

3.12 OSPF command reference

3.12.1 Command hierarchies

- [Configuration commands for OSPF](#)
- [Show commands](#)
- [Clear commands](#)

- [Debug commands](#)

3.12.1.1 Configuration commands for OSPF

```

config
- router
- [no] ospf [router-id]
- advertise-router-capability {link | area | as}
- no advertise-router-capability
- [no] area area-id
- area-range ipv4-prefix/mask | ipv6-prefix/prefixlength [advertise | not-
advertise]
- no area-range ipv4-prefix/mask | ipv6-prefix/prefixlength
- [no] blackhole-aggregate
- [no] interface ip-int-name [secondary]
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- authentication-type {password | message-digest}
- no authentication-type
- bfd-enable [remain-down-on-failure]
- no bfd-enable
- dead-interval seconds
- no dead-interval
- hello-interval seconds
- no hello-interval
- interface-type {broadcast | point-to-point}
- no interface-type
- [no] loopfree-alternate-exclude
- message-digest-key key-id md5 [key | hash-key] [hash | hash2]
- no message-digest-key key-id
- metric metric
- no metric
- mtu bytes
- no mtu
- node-sid index value
- node-sid label value
- no node-sid
- [no] passive
- priority number
- no priority
- retransmit-interval seconds
- no retransmit-interval
- [no] shutdown
- transit-delay seconds
- no transit-delay
- [no] loopfree-alternate-exclude
- [no] nssa
- area-range ipv4-prefix/mask | ipv6-prefix/prefixlength [advertise | not-
advertise]
- no area-range ipv4-prefix/mask | ipv6-prefix/prefixlength
- originate-default-route [type-7]
- no originate-default-route
- [no] redistribute-external
- [no] summaries
- [no] stub
- default-metric metric
- no default-metric
- [no] summaries
- [no] virtual-link router-id transit-area area-id
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key

```

```

- authentication-type {password | message-digest}
- no authentication-type
- dead-interval seconds
- no dead-interval
- hello-interval seconds
- no hello-interval
- message-digest-key key-id md5 [key | hash-key] [hash | hash2]
- no message-digest-key key-id
- retransmit-interval seconds
- no retransmit-interval
- [no] shutdown
- transit-delay seconds
- no transit-delay
- [no] asbr [trace-path domain-id]
- [no] compatible-rfc1583
- [no] disable-ldp-sync
- export policy-name [ policy-name...(up to 5 max)]
- no export
- export-limit number [log percentage]
- no export-limit
- external-db-overflow limit seconds
- no external-db-overflow
- external-preference preference
- no external-preference
- [no] graceful-restart
-   [no] helper-disable
- loopfree-alternate [remote-lfa]
- loopfree-alternate remote-lfa [max-pq-cost value]
- no loopfree-alternate
- loopfree-alternate-exclude prefix-policy prefix-policy [prefix-policy ... (up to
5)]
- no loopfree-alternate-exclude
- overload [timeout seconds]
- no overload
- [no] overload-include-stub
- overload-on-boot [timeout seconds]
- no overload-on-boot
- preference preference
- no preference
- reference-bandwidth bandwidth-in-kbps
- no reference-bandwidth
- router-id ip-address
- no router-id
- segment-routing
- no segment-routing
-   prefix-sid-range {global | start-label label-value max-index index-value}
-   no prefix-sid-range
-   tunnel-mtu bytes
-   no tunnel-mtu
-   tunnel-table-pref preference
-   no tunnel-table-pref
-   [no] shutdown
- [no] shutdown
- timers
-   [no] lsa-arrival lsa-arrival-time
-   [no] lsa-generate max-lsa-wait [lsa-initial-wait [lsa-second-wait]]
-   [no] spf-wait max-spf-wait [spf-initial-wait [spf-second-wait]]
- [no] traffic-engineering
- [no] ospf3 router-id
-   [no] area area-id
-     area-range ipv4-prefix/mask | ipv6-prefix/prefixlength [advertise | not-
advertise]
-     no area-range ipv4-prefix/mask | ipv6-prefix/prefixlength
-   [no] blackhole-aggregate

```

```

- [no] interface ip-int-name [secondary]
  - authentication bidirectional sa-name
  - authentication inbound sa-name outbound sa-name
  - no authentication
  - dead-interval seconds
  - no dead-interval
  - hello-interval seconds
  - no hello-interval
  - interface-type {broadcast | point-to-point}
  - no interface-type
  - [no] loopfree-alternate-exclude
  - metric metric
  - no metric
  - mtu bytes
  - no mtu
  - [no] passive
  - priority number
  - no priority
  - retransmit-interval seconds
  - no retransmit-interval
  - [no] shutdown
  - transit-delay seconds
  - no transit-delay
- key-rollover-interval
- no key-rollover-interval
- [no] nssa
  - area-range ipv4-prefix/mask | ipv6-prefix/prefixlength [advertise | not-
advertise]
    - no area-range ipv4-prefix/mask | ipv6-prefix/prefixlength
    - originate-default-route [type-nssa]
    - no originate-default-route
    - [no] redistribute-external
    - [no] summaries
- [no] stub
  - default-metric metric
  - no default-metric
  - [no] summaries
- [no] virtual-link router-id transit-area area-id
  - authentication bidirectional sa-name
  - authentication inbound sa-name outbound sa-name
  - no authentication
  - dead-interval seconds
  - no dead-interval
  - hello-interval seconds
  - no hello-interval
  - retransmit-interval seconds
  - no retransmit-interval
  - [no] shutdown
  - transit-delay seconds
  - no transit-delay
- [no] asbr [trace-path domain-id]
- export policy-name [policy-name...(up to 5 max)]
- no export
- export-limit number [log percentage]
- no export-limit
- external-db-overflow limit seconds
- no external-db-overflow
- external-preference preference
- no external-preference
- [no] graceful-restart
  - [no] helper-disable
- import policy-name [policy-name...(up to 15 max)]
- no import policy-name
- overload [timeout seconds]

```

```

- no overload
- [no] overload-include-stub
- overload-on-boot [timeout seconds]
- no overload-on-boot
- preference preference
- no preference
- reference-bandwidth bandwidth-in-kbps
- reference-bandwidth [tbps Tera-bps] [gbps Giga-bps] [mbps Mega-bps] [kbps Kilo-
bps]

- no reference-bandwidth
- router-id ip-address
- no router-id
- [no] shutdown
- timers
  - [no] lsa-arrival lsa-arrival-time
  - [no] lsa-generate max-lsa-wait [lsa-initial-wait [lsa-second-wait]]
  - [no] spf-wait max-spf-wait [spf-initial-wait [spf-second-wait]]

```

3.12.1.2 Show commands

```

show
- router
  - ospf
    - area [area-id] [detail]
    - database [type {router | network | summary | asbr-summary | external | nssa |
all} [area area-id] [adv-router router-id] [link-state-id] [detail]
    - interface [ip-int-name | ip-address] [detail]]
    - interface [area area-id] [detail]
    - neighbor [ip-int-name] [router-id] [detail]
    - prefix-sids [ip-prefix [/prefix-length]] [sid sid] [adv-router router-id]
    - range [area-id]
    - routes [ip-prefix[/prefix-length]] [type] [detail] [alternative] [summary]
  [exclude-shortcut]
    - spf [interface-name] [detail]
    - spf interface-name remote ip-address [detail]
    - spf
    - statistics
    - status
    - virtual-link [detail]
    - virtual-neighbor [remote ip-address] [detail]
  - ospf3
    - area [area-id] [detail]
    - database [type {router | network | summary | asbr-summary | external | nssa |
all} [area area-id] [adv-router router-id] [link-state-id] [detail]
    - interface [ip-int-name | ip-address] [detail]]
    - interface [area area-id] [detail]
    - neighbor [ip-int-name] [router-id] [detail]
    - range [area-id]
    - routes [ip-prefix[/prefix-length]] [type] [detail] [alternative] [summary]
  [exclude-shortcut]
    - spf
    - statistics
    - status
    - virtual-link [detail]
    - virtual-neighbor [remote ip-address] [detail]

```

3.12.1.3 Clear commands

```
clear
- router
- ospf ospf-instance
- database [purge]
- export
- neighbor [ip-int-name | ip-address]
- statistics
- ospf3
- database [purge]
- export
- neighbor [ip-int-name | ip-address]
- statistics
```

3.12.1.4 Debug commands

```
debug
- router
- ospf ospf-instance
- area area-id
- no area
- area-range ip-address
- no area-range
- cspf ip-addr
- no cspf
- [no] graceful-restart
- interface [ip-int-name | ip-address]
- no interface
- leak ip-address
- no leak
- lsdb [type] [ls-id] [adv-rtr-id] [area area-id]
- no lsdb
- [no] misc
- neighbor [ip-int-name | router-id]
- no neighbor
- nssa-range ip-address
- no nssa-range
- packet [packet-type] [ip-address]
- no packet
- rtm ip-addr
- no rtm
- spf [type] [dest-addr]
- no spf
- virtual-neighbor ip-address
- no virtual-neighbor
- ospf3
- area area-id
- no area
- area-range ip-address
- no area-range
- [no] graceful-restart
- interface [ip-int-name | ip-address]
- no interface
- leak ip-address
- no leak
- lsdb [type] [ls-id] [adv-rtr-id] [area area-id]
- no lsdb
- [no] misc
```

```
- neighbor [ip-int-name | router-id]
- no neighbor
- nssa-range ip-address
- no nssa-range
- packet [packet-type] [ip-address]
- no packet
- rtm ip-addr
- no rtm
- spf [type] [dest-addr]
- no spf
- virtual-neighbor ip-address
- no virtual-neighbor
```

3.12.2 Command descriptions

3.12.2.1 Configuration commands

3.12.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

```
config>router>ospf
config>router>ospf>area>interface
config>router>ospf>area>virtual-link
config>router>ospf>segment-routing
config>router>ospf3
config>router>ospf3>area>interface
config>router>ospf3>area>virtual-link
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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

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.

Special Cases

OSPF Protocol

The Open Shortest Path First (OSPF) protocol is created in the **no shutdown** state.

OSPF Interface

When an IP interface is configured as an OSPF interface, OSPF on the interface is in the **no shutdown** state by default.

3.12.2.1.2 OSPF global commands

```
ospf
```

Syntax

```
[no] ospf [ospf-instance]
```

Context

```
config>router
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the router ID for OSPF.

The router ID configured in the base instance of OSPF overrides the router ID configured in the **config>router** context.

The default value for the base instance is inherited from the configuration in the **config>router** context. When that is not configured the following applies:

- The system uses the system interface address (which is also the loop-back address).
- If a system interface address is not configured, use the last 32 bits of the chassis MAC address.

This is a required command when configuring multiple instances and the instance being configured is not the base instance. When configuring multiple instances of OSPF, there is a risk of loops because networks are advertised by multiple domains configured with multiple interconnections to one another. To avoid this from happening all routers in a domain should be configured with the same domain-id. Each domain (OSPF instance) should be assigned a specific bit value in the 32-bit tag mask.

The default value for non-base instances is 0.0.0.0 and is invalid, in this case the instance of OSPF will not start. When configuring a new router ID, the instance is not automatically restarted with the new router ID.

The next time the instance is initialized, the new router ID is used.

Issue the **shutdown** and **no shutdown** commands for the instance for the new router ID to be used, or reboot the entire router

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



Note:

- The platforms as described in this document allow for the configuration of a single OSPF instance at any time. The instance ID can be any number other than 0. This enables these platforms to be used in a network where multi-instance OSPF is deployed, and the node must use an instance ID other than the default instance ID of 0.
- The number of OSPF instances supported on the 7210 SAS differs depending on the platform. Contact a Nokia representative for information about the supported scaling limits.

Default

no ospf

Parameters

ospf-instance

Specifies a unique integer that identifies a specific instance of a version of the OSPF protocol running in the router instance specified by the router ID.

Values 0 to 31

ospf3

Syntax

[no] ospf3

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure OSPF to support IPv6.

Before OSPFv3 can be activated on the router, the router ID must be configured. The router ID is derived by one of the following methods:

- defining the value using the **config>router>router-id** *ip-address* command
- defining the system interface using the **config>router>interface** *ip-int-name* command (used if the router ID is not specified with the **config>router>router-id** *ip-address* command)
- inheriting the last 4 bytes of the MAC address

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. To force the new router ID, issue the **shutdown** and **no shutdown** commands for OSPFv3 or restart the entire router.

The **no** form of this command disables OSPF support for IPv6.

advertise-router-capability

Syntax

advertise-router-capability {link | area | as}

no advertise-router-capability

Context

config>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the advertisement of router capabilities to its neighbors for informational and troubleshooting purposes. A router information (RI) LSA as defined in RFC 4970 advertises the following capabilities:

- OSPF graceful restart capable: no
- OSPF graceful restart helper: yes, when enabled
- OSPF stub router support: yes
- OSPF traffic engineering support: yes, when enabled
- OSPF point-to-point over LAN: yes
- OSPF experimental TE: no

The **link**, **area**, and **as** keywords control the scope of the capability advertisements.

The **no** form of this command disables this advertisement capability.

Default

no advertise-router-capability

Parameters

link

Keyword specifying to advertise only over local links and not flood beyond.

area

Keyword specifying to advertise only within the area of origin.

as

Keyword specifying to advertise throughout the entire autonomous system.

asbr

Syntax

[no] asbr [trace-path domain-id]

Context

config>router>ospf

config>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description


This command configures the router as an Autonomous System Boundary Router (ASBR) if the router is to be used to export routes from the Routing Table Manager (RTM) into this instance of OSPF. When a router is configured as an ASBR, the export policies into this OSPF domain take effect. If no policies are configured, no external routes are redistributed into the OSPF domain.

The **no** form of this command removes the ASBR status and withdraws the routes redistributed from the RTM into this OSPF instance from the link state database.

When configuring multiple instances of OSPF, there is a risk of loops because networks are advertised by multiple domains configured with multiple interconnections to one another. To avoid loops, all routers in a domain should be configured with the same domain ID. Each domain (OSPF instance) should be assigned a specific bit value in the 32-bit tag mask.

When an external route is originated by an ASBR using an internal OSPF route in a specific domain, the corresponding bit is set in the AS-external LSA. As the route gets redistributed from one domain to another, more bits are set in the tag mask, each corresponding to the OSPF domain the route visited. Route redistribution looping is prevented by checking the corresponding bit as part of the export policy; if the bit corresponding to the announcing OSPF process is already set, the route is not exported there. The following figure shows the checking of the corresponding bit.

Figure 10: Checking corresponding bit



Domain IDs are incompatible with any other use of normal tags. The domain ID should be configured with a value between 1 and 31 by each router in a specific OSPF domain (OSPF instance).

When an external route is originated by an ASBR using an internal OSPF route in a specific domain, the corresponding bit (1 to 31) is set in the AS-external LSA.

The **no** form of this command removes the ASBR status and withdraws the routes redistributed from the routing table into OSPF from the link-state database.

Default

no asbr

Parameters

domain-id

Specifies the domain ID.

Values 1 to 31

Default 0

compatible-rfc1583

Syntax

[no] compatible-rfc1583

Context

config>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables OSPF summary and external route calculations in compliance with RFC 1583 and earlier RFCs.

RFC 1583 and earlier RFCs use a different method to calculate summary and external route costs. To avoid routing loops, all routers in an OSPF domain should perform the same calculation method.

Although it would be favorable to require all routers to run a more current compliancy level, this command allows the router to use obsolete methods of calculation.

The **no** form of this command enables the post-RFC 1583 method of summary and external route calculation.

Default

compatible-rfc1583

disable-ldp-sync

Syntax

[no] disable-ldp-sync

Context

config>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command disables the IGP-LDP synchronization feature on all interfaces participating in the OSPF routing protocol. When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces that have the IGP-LDP synchronization enabled if the currently advertised cost is different. It will disable IGP-LDP synchronization for all interfaces. This command does not delete the interface configuration.

For information about LDP synchronization, see "IGP-LDP and static route-LDP Synchronization on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C" and the **ldp-sync** and **ldp-sync-timer** commands in the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*.

The **no** form of this command restores the default settings and re-enables IGP-LDP synchronization on all interfaces participating in the OSPF routing protocol and for which the **ldp-sync-timer** is configured.

Default

no disable-ldp-sync

export

Syntax

export *policy-name* [*policy-name*...(5 maximum)]

no export

Context

config>router>ospf

config>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command associates export route policies to determine which routes are exported from the route table to OSPF or OSPFv3. Export polices are only in effect if OSPF or OSPFv3 is configured as an ASBR.

If no export policy is specified, non-OSPF or OSPFv3 routes are not exported from the routing table manager to OSPF or OSPFv3.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

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

Default

no export

Parameters

policy-name

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

export-limit

Syntax

export-limit *number* [**log** *percentage*]

no export-limit

Context

config>router>ospf

config>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum number of routes (prefixes) that can be exported into OSPF or OSPFv3 from the route table.

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

Default

no export-limit

Parameters

number

Specifies the maximum number of routes (prefixes) that can be exported into OSPF or OSPFv3 from the route table.

Values 1 to 4294967295

log *percentage*

Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.

Values 1 to 100

external-db-overflow

Syntax

external-db-overflow *limit seconds*
no external-db-overflow

Context

config>router>ospf
config>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures limits on the number of non-default AS-external LSA entries that can be stored in the link-state database (LSDB) and specifies a wait timer before processing these entries after the limit is exceeded.

The *limit* value specifies the maximum number of non-default AS-external LSA entries that can be stored in the LSDB. Placing a limit on the non-default AS-external LSAs in the LSDB protects the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reaches or exceeds the configured *limit*, the table is in an overflow state. When in an overflow state, the router will not originate any new AS-external LSAs and will withdraws all the self-originated non-default external LSAs.

The *seconds* value specifies the amount of time to wait after an overflow state before regenerating and processing non-default AS-external LSAs. The waiting period acts like a dampening period preventing the router from continuously running Shortest Path First (SPF) calculations caused by the excessive number of non-default AS-external LSAs.

The **external-db-overflow** command must be set identically on all routers attached to any regular OSPF or OSPFv3 area. OSPF or OSPFv3 stub areas and not-so-stubby areas (NSSAs) are excluded.

The **no** form of this command disables limiting the number of non-default AS-external LSA entries.

Default

no external-db-overflow

Parameters

limit

Specifies the maximum number of non-default AS-external LSA entries that can be stored in the LSDB before going into an overflow state expressed as a decimal integer.

Values 1 to 2147483674

interval

Specifies the number of seconds after entering an overflow state before attempting to process non-default AS-external LSAs, expressed as a decimal integer.

Values 0 to 2147483674

external-preference

Syntax

external-preference *preference*

no external-preference

Context

config>router>ospf

config>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the preference for OSPF or OSPFv3 external routes.

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs the preference is used to decide which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is determined by the default preference as defined in the [Table 28: Route preference defaults by route type](#). If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

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

Default

150

Parameters

preference

Specifies the preference for external routes expressed as a decimal integer. Defaults for different route types are listed in the following table.

Table 28: Route preference defaults by route type

Route type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF/ OSPFv3 internal	10	Yes
IS-IS level 1 internal	15	Yes

Route type	Preference	Configurable
IS-IS level 2 internal	18	Yes
OSPF/OSPFv3 external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes



Note:
Preference for OSPF or OSPFv3 internal routes is configured with the **preference** command.

Values 1 to 255

graceful-restart

Syntax
[no] graceful-restart

Context
config>router>ospf
config>router>ospf3

Platforms
Supported on all 7210 SAS platforms as described in this document

Description
This command enables graceful-restart for OSPF or OSPFv3. When the control plane of a GR-capable router fails, the neighboring routers (GR helpers) temporarily preserve adjacency information, so packets continue to be forwarded through the failed GR router using the last known routes. If the control plane of the GR router comes back up within the GR timer, the routing protocols reconverges to minimize service interruption.

The **no** form of this command disables graceful restart and removes all graceful restart configurations in the OSPF or OSPFv3 instance.

Default
no graceful-restart

helper-disable

Syntax

[no] helper-disable

Context

config>router>ospf>graceful-restart

config>router>ospf3>graceful-restart

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command disables the helper support for graceful restart.

When **graceful-restart** is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart) or be a restarting router or both. The router supports only helper mode. This facilitates the graceful restart of neighbors but will not act as a restarting router.

The **no** form of this command enables helper support and is the default when **graceful-restart** is enabled.

Default

helper-disable

import

Syntax

import [*policy-name*...(up to 15 max)]

no import *policy-name*

Context

config>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the import route policy for an OSPF3 instance.

The **no** form of this command removes the policy association with the OSPF3 instance.

Default

no import

Parameters

policy-name

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

loopfree-alternate

Syntax

loopfree-alternate [**remote-lfa**]

loopfree-alternate remote-lfa [**max-pq-cost** *value*]

no loopfree-alternate

Context

config>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables Loop-Free Alternate (LFA) computation by SPF for the OSPF routing protocol instance.

This command instructs the IGP SPF to precompute both a primary next hop and an LFA next hop for every learned prefix. When found, the LFA next hop is populated into the routing table along with the primary next-hop for the prefix.

The remote LFA next hop calculation by the IGP LFA SPF is enabled by using the **remote-lfa** option. When this option is enabled in an IGP instance, SPF performs the remote LFA additional computation following the regular LFA next-hop calculation when the latter results in no protection for one or more prefixes that are resolved to a specific interface.

Remote LFA extends the protection coverage of LFA-FRR to any topology by automatically computing and establishing or tearing down shortcut tunnels, also referred to as repair tunnels, to a remote LFA node. Doing this puts the packets back into the shortest path without looping them back to the node that forwarded them over the repair tunnel. The remote LFA node is referred to as a PQ node. A repair tunnel can, in theory, be an RSVP LSP, an LDP-in-LDP tunnel, or a segment routing tunnel. Using segment routing repair tunnels is restricted to the remote LFA node.

The remote LFA algorithm is a per-link LFA SPF calculation and is not per-prefix, like the regular LFA calculation. It provides protection to all destination prefixes that share the protected link by using the neighbor on the other side of the protected link as a proxy for those prefixes.

Default

no loopfree-alternate

Parameters

remote-lfa

Keyword to enable the remote LFA next-hop calculation by the IGP LFA SPF.

value

Specifies the maximum IGP cost from the router that is performing the remote LFA calculation to the candidate P or Q node.

Values 0 to 4294967295

loopfree-alternate-exclude

Syntax

loopfree-alternate-exclude prefix-policy *prefix-policy* [*prefix-policy* ... (up to 5)]

no loopfree-alternate-exclude

Context

config>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.

The implementation already allows the user to exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority. The prefix tag will, however, be used in the main SPF. Note that prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action of the **loopfree-alternate-exclude** command, when not explicitly specified by the user in the prefix policy, is a "reject". Therefore, regardless if the user did or did not explicitly add the statement "default-action reject" to the prefix policy, a prefix that did not match any entry in the policy will be accepted into LFA SPF.

The **no** form of this command deletes the exclude prefix policy.

Parameters

prefix-policy

Specifies the name of the prefix policy. The specified name must have been previously defined. 32 characters maximum.

overload

Syntax

overload [timeout seconds]
no overload

Context

config>router>ospf
config>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command changes the overload state of the local router so that it appears to be overloaded. When overload is enabled, the router can participate in OSPF or OSPFv3 routing, but is not used for transit traffic. Traffic destined to directly attached interfaces continues to reach the router.

To put the IGP in an overload state, enter a timeout value. The IGP will enter the overload state until the timeout timer expires or a **no overload** command is executed.

If the **overload** command is encountered during the execution of an [overload-on-boot](#) command, this command takes precedence. This could occur as a result of a saved configuration file where both parameters are saved. When the file is saved by the system the **overload-on-boot** command is saved after the **overload** command. However, if **overload-on-boot** is configured under OSPF or OSPFv3 with no timeout value configured, the router will remain in overload state indefinitely after a reboot.

The **no** form of this command reverts to the default. When the **no overload** command is executed, the overload state is terminated regardless of the reason the protocol entered overload state.

Default

no overload

Parameters

timeout seconds

Specifies the number of seconds to reset overloading.

Values 60 to 1800

Default 60

overload-include-stub

Syntax

[no] **overload-include-stub**

Context

```
config>router>ospf  
config>router>ospf3
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command determines whether the OSPF or OSPFv3 stub networks should be advertised with a maximum metric value when the system goes into overload state for any reason. When enabled, the system uses the maximum metric value. When this command is enabled and the router is in overload, all stub interfaces, including loopback and system interfaces, will be advertised at the maximum metric.

Default

no overload-include-stub

overload-on-boot

Syntax

```
overload-on-boot [timeout seconds]  
no overload
```

Context

```
config>router>ospf  
config>router>ospf3
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the IGP upon bootup in the overload state until one of the following events occurs:

- the timeout timer expires.
- a manual override of the current overload state is entered with the **no overload** command.

When the router is in an overload state, the router is used only if there is no other router to reach the destination. The **no overload** command does not affect the **overload-on-boot** function.

The **no** form of this command removes the **overload-on-boot** function from the configuration.

Parameters

timeout *seconds*

Specifies the number of seconds to reset overloading.

Values 60 to 1800

preference

Syntax

preference *preference*

no preference

Context

config>router>ospf

config>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the preference for OSPF or OSPFv3 internal routes.

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide which route will be used.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is based on the default preference as defined in [Table 28: Route preference defaults by route type](#). If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

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

Default

preference 10

Parameters

preference

Specifies the preference for internal routes, expressed as a decimal integer.

Values 1 to 255

reference-bandwidth

Syntax

reference-bandwidth *reference-bandwidth*

reference-bandwidth [**tbps** *Tera-bps*] [**gbps** *Giga-bps*] [**mbps** *Mega-bps*] [**kbps** *Kilo-bps*]

no reference-bandwidth

Context

```
config>router>ospf  
config>router>ospf3
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the reference bandwidth used to calculate the default costs of interfaces based on their underlying link speed.

The default interface cost is calculated as follows:

cost = reference-bandwidth # bandwidth

The default *reference-bandwidth* is 100,000,000 Kb/s or 100 Gb/s, so the default auto-cost metrics for various link speeds are as follows:

- 10 Mb/s link default cost of 10000
- 100 Mb/s link default cost of 1000
- 1 Gb/s link default cost of 100
- 10 Gb/s link default cost of 10

The **reference-bandwidth** command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the **metric** *metric* command in the **config>router>ospf>area>interface** context.

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

Default

reference-bandwidth 100000000

Parameters

bandwidth-in-kbps

Specifies the reference bandwidth in kilobits per second, expressed as a decimal integer.

Values 1 to 4000000000

Tera-bps

Specifies the reference bandwidth in terabits per second, expressed as a decimal integer.

Values 1 to 4

Giga-bps

Specifies the reference bandwidth in gigabits per second, expressed as a decimal integer.

Values 1 to 4

Mega-bps

Specifies the reference bandwidth in megabits per second, expressed as a decimal integer.

Values 1 to 999

Kilo-bps

Specifies the reference bandwidth in kilobits per second, expressed as a decimal integer.

Values 1 to 999

router-id

Syntax

router-id *ip-address*

no router-id

Context

config>router>ospf

config>router>osp3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the router ID for the OSPF or OSPFv3 instance.

When configuring the router ID in the base instance of OSPF or OSPFv3, it overrides the router ID configured in the **config>router** context.

The default value for the base instance is inherited from the configuration in the **config>router** context. If the router ID in the **config>router** context is not configured, the following applies.

- The system uses the system interface address (which is also the loopback address).
- If a system interface address is not configured, use the last 32 bits of the chassis MAC address.

When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for the instance, or reboot the entire router.

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

Default

router-id 0.0.0.0

Parameters

ip-address

Specifies a 32-bit unsigned integer uniquely identifying the router in the Autonomous System.

timers

Syntax

timers

Context

config>router>ospf

config>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure OSPF or OSPFv3 timers. Timers control the delay between receipt of an LSA requiring a Dijkstra (Shortest Path First (SPF)) calculation and the minimum time between successive SPF calculations.

Changing the timers affects CPU utilization and network reconvergence times. Lower values reduce the convergence time but increase CPU utilization. Higher values reduce CPU utilization but increase the reconvergence time.

lsa-arrival

Syntax

lsa-arrival *lsa-arrival-time*

no lsa-arrival

Context

config>router>ospf>timers

config>router>ospf3>timers

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the minimum delay that must pass between receipt of the same LSAs arriving from neighbors.

Nokia recommends that the neighbor configured **lsa-generate** *lsa-second-wait* interval is equal or greater than the **lsa-arrival** timer configured here.

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

Default

no lsa-arrival

Parameters

lsa-arrival-time

Specifies the timer, in milliseconds. Values entered that do not match this requirement are rejected.

Values 0 to 600000

lsa-generate

Syntax

lsa-generate *max-lsa-wait* [*lsa-initial-wait* [*lsa-second-wait*]]

no lsa-generate-interval

Context

config>router>ospf>timers

config>router>ospf3>timers

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command customizes the throttling of OSPF or OSPFv3 LSA generation. Timers that determine when to generate the first, second, and subsequent LSAs can be controlled with this command. Subsequent LSAs are generated at increasing intervals of the *lsa-second-wait* timer until a maximum value is reached.

Configuring the **lsa-arrival** interval to equal or less than the *lsa-second-wait* interval configured in the **lsa-generate** command is recommended.

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

Default

no lsa-generate

Parameters

max-lsa-wait

Specifies the maximum interval, in milliseconds, between two consecutive occurrences of an LSA being generated

Values 10 to 600000

Default 5000

lsa-initial-wait

Specifies the first waiting period between LSA originates, in milliseconds. When the LSA exceeds the *lsa-initial-wait* timer value and the topology changes, there is no wait period and the LSA is immediately generated.

When an LSA is generated, the initial wait period commences. If within the specified *lsa-initial-wait* period and another topology change occurs, the *lsa-initial-wait* timer applies.

Values 10 to 600000

Default 5000

lsa-second-wait

Specifies the hold time in milliseconds between the first and second LSA generation. The next topology change is subject to this second wait period. With each subsequent topology change, the wait time doubles (this is 2x the previous wait time.). This assumes that each failure occurs within the relevant wait period.

Values 10 to 600000

Default 5000

spf-wait

Syntax

spf-wait *max-spf-wait* [*spf-initial-wait* [*spf-second-wait*]]

no spf-wait

Context

config>router>ospf>timers

config>router>ospf3>timers

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the maximum interval between two consecutive SPF calculations in milliseconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, the next SPF will run after 2000 milliseconds, and then next SPF will run after 4000 milliseconds, and so on, until it reaches the **spf-wait** value. The SPF interval will stay at the **spf-wait** value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to the *spf-initial-wait* value.

The timer must be entered in increments of 100 milliseconds. Values entered that do not match this requirement will be rejected.

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

Default

no spf-wait

Parameters

max-spf-wait

Specifies the maximum interval, in milliseconds, between two consecutive SPF calculations.

Values 10 to 120000

Default 1000

spf-initial-wait

Specifies the initial SPF calculation delay, in milliseconds, after a topology change.

Values 10 to 100000

Default 1000

spf-second-wait

Specifies the hold time, in milliseconds, between the first and second SPF calculation.

Values 10 to 100000

Default 1000

segment-routing

Syntax

segment-routing

no segment-routing

Context

config>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure segment routing parameters within an IGP instance.

Segment routing adds to OSPF routing protocols the ability to perform shortest path routing and source routing using the concept of the abstract segment. A segment can represent a local prefix of a node, a specific adjacency of the node (interface or next-hop), a service context, or a specific explicit path over the network. For each segment, the IGP advertises an identifier referred to as a segment ID (SID).

When segment routing is used together with the MPLS data plane, the SID is a standard MPLS label. A router forwarding a packet using segment routing pushes one or more MPLS labels.

Segment routing using MPLS labels is used in both shortest path routing applications and traffic engineering applications. This command configures the shortest path forwarding application.

After segment routing is configured in the OSPF instance, the router performs the following operations.

1. Advertises the segment routing capability sub-TLV to routers in all areas and levels of this IGP instance. However, only neighbors with which it established an adjacency will interpret the SID and label range information and use it for calculating the label to swap to or push for a given resolved prefix SID.
2. Advertises the assigned index for each configured node SID in the new prefix SID sub-TLV with the N-flag (node-SID flag) set. The segment routing module then programs the incoming label map (ILM) with a pop operation for each local node SID in the data path.
3. Automatically assigns and advertises an adjacency SID label for each formed adjacency over a network IP interface in the new adjacency SID sub-TLV. The segment routing module programs the ILM with a pop operation (in effect with a swap to an implicit null label operation), for each advertised adjacency SID.
4. Resolves received prefixes, and if a prefix SID sub-TLV exists, the segment routing module programs the ILM with a swap operation and an LTN with a push operation both pointing to the primary/LFA NHLFE. An SR tunnel is also added to the TTM.

When the user enables segment routing in an IGP instance, the main SPF and LFA SPF are computed, and the primary next-hop and LFA backup next-hop for a received prefix are added to the RTM without the label information advertised in the prefix SID sub-TLV.

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

prefix-sid-range

Syntax

prefix-sid-range {**global** | **start-label** *label-value* **max-index** *index-value*}

no prefix-sid-range

Context

config>router>ospf>segment-routing

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the prefix SID index range and offset label value for an IGP instance.

The user must configure the prefix SID index range and the offset label value that this IGP instance uses. Because each prefix SID represents a network global IP address, the SID index for a prefix must be unique in the network. Therefore, all routers in the network are expected to configure and advertise the same prefix SID index range for an IGP instance. However, the label value used by each router to represent this prefix, that is, the label programmed in the ILM, can be local to that router by the use of an offset label, referred to as a start label, as in the following:

Local Label (Prefix SID) = start-label + {SID index}

The label operation in the network becomes similar to LDP when operating in the independent label distribution mode (RFC 5036), with the difference that the label value used to forward a packet to each downstream router is computed by the upstream router based on the advertised prefix SID index using the preceding formula.

There are two mutually exclusive modes of operation for the prefix SID range on the router. In the global mode of operation, the user configures the global value and this IGP instance assumes the start label value is the lowest label value in the SRGB and the prefix SID index range size equal to the range size of the SRGB. When one IGP instance selects the global option for the prefix SID range, all IGP instances on the system are restricted to do the same. The user must shut down the segment routing context and delete the **prefix-sid-range** command in all IGP instances to change the SRGB. After the SRGB is changed, the user must re-enter the **prefix-sid-range** command. The SRGB range change fails if an already allocated SID index or label goes out of range.

In the per-instance mode of operation, the user partitions the SRGB into non-overlapping sub-ranges among the IGP instances. The user therefore configures a subset of the SRGB by specifying the start label value and the prefix SID index range size. All resulting net label values (start-label + index) must be within the SRGB or the configuration fails.

The code checks for overlaps of the resulting net label value range across IGP instances and strictly enforces that these ranges do not overlap. The user must shut down the segment routing context of an IGP instance to change the SID index or label range of that IGP instance using the **prefix-sid-range** command.

In addition, any range change will fail if an already allocated SID index or label goes out of range. The user can, however, change the SRGB on the fly as long as it does not reduce the current per-IGP instance SID index or label range defined in the **prefix-sid-range** command. Otherwise, the user must shut down the segment routing context of the IGP instance and delete and reconfigure the **prefix-sid-range** command.

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

Default

no prefix-sid-range

Parameters

start-label label-value

Specifies the label offset for the SR label range of this IGP instance.

Values 0 to 524287

max-index index-value

Specifies the maximum value of the prefix SID index range for this IGP instance.

Values 1 to 524287

global

Keyword to enable global operation mode.

tunnel-mtu

Syntax

tunnel-mtu *bytes*

no tunnel-mtu

Context

config>router>ospf>segment-routing

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum transmission unit (MTU) of all SR tunnels within each IGP instance.

The MTU of an SR tunnel populated into the TTM is determined like in the case of an IGP tunnel; for example, LDP LSP, based on the outgoing interface MTU minus the label stack size. Remote LFA can add at least two more labels to the tunnel for a total of three labels. There is no default value. If the user does not configure an SR tunnel MTU, the MTU is determined by IGP.

The MTU of the SR tunnel in bytes is determined as follows:

$$\text{SR_Tunnel_MTU} = \text{MIN} \{ \text{Cfg_SR_MTU}, \text{IGP_Tunnel_MTU} - (1 + \text{frr-overhead}) * 4 \}$$

Where:

- Cfg_SR_MTU is the MTU configured by the user for all SR tunnels within a specific IGP instance using this command. If no value was configured by the user, the SR tunnel MTU will be determined by the following IGP interface calculation.
- IGP_Tunnel_MTU is the minimum of the IS-IS or OSPF interface MTU among all the ECMP paths or among the primary and LFA backup paths of this SR tunnel.
- frr-overhead is set to 1 if **segment-routing** and **remote-lfa** options are enabled in the IGP instance. Otherwise, it is set to 0.

The SR tunnel MTU is dynamically updated whenever any of the preceding parameters used in its calculation changes. This includes when the set of tunnel next hops changes, or the user changes the configured SR MTU or interface MTU value.

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

Default

no tunnel-mtu

Parameters

bytes

Specifies the size of the MTU in bytes.

Values 512 to 9198

tunnel-table-pref

Syntax

tunnel-table-pref *preference*

no tunnel-table-pref

Context

config>router>ospf>segment-routing

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the TTM preference of shortest path SR tunnels created by the IGP instance. The TMM preference is used in the case of VPRN auto-bind or BGP transport tunnels when the new tunnel binding commands are configured to the **any** value, which parses the TTM for tunnels in the protocol preference order. The user can either use the global TTM preference or list the tunnel types they want to use. When they list the tunnel types explicitly, the TTM preference is used to select one type over the other. In both cases, a fallback to the next preferred tunnel type is performed if the selected one fails. Also, a reversion to a more preferred tunnel type is performed as soon as one is available.

The segment routing module adds an SR tunnel entry to the TTM for each resolved remote node SID prefix and programs the datapath with the corresponding LTN with the push operation pointing to the primary and LFA backup NHLFEs.

The default preference for shortest path SR tunnels in the TTM is set lower than LDP tunnels but higher than BGP tunnels to allow controlled migration of customers without disrupting their current deployment when they enable segment routing. The following is the setting of the default preference for various tunnel types. This includes the preference of SR tunnels based on the shortest path (referred to as SR-OSPF).

The global default TTM preference for the tunnel types is as follows:

- ROUTE_PREF_RSVP 7
- ROUTE_PREF_SR_TE 8
- ROUTE_PREF_LDP 9
- ROUTE_PREF_OSPF_TTM 10
- ROUTE_PREF_ISIS_TTM 11
- ROUTE_PREF_BGP_TTM 12
- ROUTE_PREF_GRE 255

The default value for SR-OSPF is the same regardless of whether one or more OSPF instances programmed a tunnel for the same prefix. The selection of an SR tunnel in this case is based on the lowest IGP instance ID.

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

Default

no tunnel-table-pref

Parameters

preference

Specifies the integer value to represent the preference of OSPF SR tunnels in the TTM.

Values 1 to 255

Default 10

traffic-engineering

Syntax

[no] traffic-engineering

Context

config>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables traffic engineering route calculations constrained by nodes or links.

Traffic engineering enables the router to perform route calculations constrained by nodes or links. The traffic engineering capabilities of this router are limited to calculations based on link and nodal constraints.

The **no** form of this command disables traffic-engineered route calculations.

Default

no traffic-engineering

3.12.2.1.3 OSPF area commands

area

Syntax

[no] area *area-id*

Context

config>router>ospf

config>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an OSPF or OSPF3 area. An area is a collection of network segments within an AS that have been administratively grouped together. The area ID can be specified in dotted decimal notation or as a 32-bit decimal integer.

The **no** form of this command deletes the specified area from the configuration. Deleting the area also removes the OSPF or OSPF3 configuration of all the interfaces, virtual-links, and address-ranges that are currently assigned to this area.

Default

no area

Parameters

area-id

Specifies the OSPF or OSPF3 area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

Values	ip-address: a.b.c.d (dotted decimal)
	area: 0 to 4294967295 (decimal integer)

area-range

Syntax

area-range *ipv4-prefix/mask* | *ipv6-prefix/prefix-length* [**advertise** | **not-advertise**]

no area-range *ipv4-prefix/mask* | *ipv6-prefix/prefix-length*

Context

config>router>ospf>area

config>router>ospf>area>nssa

config>router>ospf3>area

config>router>ospf3>area>nssa

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, the range is configured to be advertised or not advertised into other areas. Multiple range commands may be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.

ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.

The **no** form of this command deletes the range advertisement or non-advertisement.

Default

no area-range

Special Cases

NSSA Context

Specifies that the range applies to external routes (via type-7 LSAs) learned within the NSSA when the routes are advertised to other areas as type-5 LSAs.

Area Context

If this command is not entered under the NSSA context, the range applies to summary LSAs even if the area is an NSSA

Parameters

ipv4-prefix/mask

Specifies the IPv4 prefix for the range in dotted-decimal notation and the subnet mask for the range, expressed as a decimal integer.

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

ipv6-prefix/prefix-length

Specifies the IPv6 prefix for the range in hexadecimal notation, and the prefix length for the range, expressed as a decimal integer.

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

advertise | not-advertise

Specifies whether to advertise the summarized range of addresses to other areas.

Default advertise

blackhole-aggregate

Syntax

[no] blackhole-aggregate

Context

```
config>router>ospf>area  
config>router>ospf3>area
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command installs a low priority blackhole route for the entire aggregate. Existing routes that make up the aggregate will have a higher priority and only the components of the range for which no route exists are blackholed.

It is possible that when performing area aggregation, addresses may be included in the range for which no actual route exists. This can cause routing loops. To avoid this problem, configure the **blackhole-aggregate** option.

The **no** form of this command removes this option.

Default

blackhole-aggregate

default-metric

Syntax

```
default-metric metric  
no default-metric
```

Context

```
config>router>ospf>area>stub  
config>router>ospf3>area>stub
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the metric used by the area border router (ABR) for the default route into a stub area.

The default metric should only be configured on an ABR of a stub area.

An ABR generates a default route if the area is a stub area.

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

Default

default-metric 1

Parameters

metric

Specifies the metric, expressed as a decimal integer, for the default route cost to be advertised into the stub area.

Values 1 to 16777215

key-rollover-interval

Syntax

key-rollover-interval *seconds*

no key-rollover-interval

Context

config>router>ospf3>area

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the key rollover interval.

The **no** form of this command resets the configured interval to the default setting.

Default

key-rollover-interval 10

Parameters

seconds

Specifies the time, in seconds, after which a key rollover will start.

Values 10 to 300

nssa

Syntax

[no] nssa

Context

config>router>ospf>area

config>router>ospf3>area

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure an OSPF or OSPFv3 NSSA and adds or removes the NSSA designation from the area.

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is that an NSSA has the capability to flood external routes that it learns throughout its area and via an ABR to the entire OSPF or OSPFv3 domain.

Existing virtual links of a non-stub area or NSSA area are removed when the designation is changed to NSSA or stub.

An area can be designated as stub or NSSA, but never both at the same time.

By default, an area is not configured as an NSSA area.

The **no** form of this command removes the NSSA designation and configuration context from the area.

Default

no nssa

originate-default-route

Syntax

originate-default-route [type-7]

originate-default-route [type-nssa]

no originate-default-route

Context

config>router>ospf>area>nssa

config>router>ospf3>area>nssa

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the generation of a default route and its LSA type into an NSSA by an NSSA ABR or ASBR.

The functionality of the **type-7** parameter and the **type-nssa** parameter is the same. The **type-7** parameter is available in the **ospf** context; the **type-nssa** parameter is available in the **ospf3** context. Include the **type-7** or **type-nssa** parameter to inject a type 7 LSA default route instead of the type 3 LSA into the NSSA configured with no summaries.

To revert to a type 3 LSA, enter the **originate-default-route** command without the **type-7** or **type-nssa** parameter.

When configuring an NSSA with no summaries, the ABR will inject a type 3 LSA default route into the NSSA area. Some older implementations expect a type 7 LSA default route.

The **no** form of this command disables origination of a default route.

Default

no originate-default-route

Parameters

type-7 | type-nssa

Specifies that a type 7 LSA or type NSSA should be used for the default route.

Default Type 3 LSA for the default route

redistribute-external

Syntax

[no] redistribute-external

Context

config>router>ospf>area>nssa

config>router>ospf3>area>nssa

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the redistribution of external routes into the NSSA or an NSSA ABR that is exporting the routes into non-NSSA areas.

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF or OSPFv3 areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (provided it is an ASBR) throughout its area and via an Area Border Router to the entire OSPF or OSPFv3 domain.

The **no** form of this command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.

Default

redistribute-external

stub

Syntax

[no] stub

Context

```
config>router>ospf>area  
config>router>ospf3>area
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure an OSPF or OSPF3 stub area and adds or removes the stub designation from the area.

External routing information is not flooded into stub areas. All routers in the stub area must be configured with the **stub** command. An OSPF or OSPF3 area cannot be both an NSSA and a stub area.

Existing virtual links of a non-stub or NSSA area will be removed when its designation is changed to NSSA or stub.

By default, an area is not a stub area.

The **no** form of this command removes the **stub** designation and configuration context from the area.

Default

no stub

summaries

Syntax

[no] summaries

Context

```
config>router>ospf>area>nssa  
config>router>ospf>area>stub  
config>router>ospf3>area>nssa  
config>router>ospf3>area>stub
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables sending summary (type 3) advertisements into a stub area or NSSA on an ABR.

This parameter is particularly useful to reduce the size of the routing and link-state database (LSDB) tables within the stub or NSSA area.

By default, summary route advertisements are sent into the stub area or NSSA.

The **no** form of this command disables sending summary route advertisements and, for stub areas, only the default route is advertised by the ABR.

Default

summaries

3.12.2.1.4 Interface/virtual link commands

authentication

Syntax

authentication bidirectional *sa-name*

authentication [**inbound** *sa-name* **outbound** *sa-name*]

no authentication

Context

config>router>ospf3>area>interface

config>router>ospf3>area>virtual-link

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an interface with a static security association (SA) used to authenticate OSPFv3 packets.

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

Default

no authentication

Parameters

bidirectional *sa-name*

Specifies the IPsec SA name, up to 32 characters, used for transmitting and receiving OSPFv3 packets.

inbound *sa-name*

Specifies the IPsec SA name, up to 32 characters, used for receiving OSPFv3 packets.

outbound *sa-name*

Specifies the IPsec SA name, up to 32 characters, used for transmitting OSPFv3 packets.

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

config>router>ospf>area>interface

config>router>ospf3>area>virtual-link

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the password used by the OSPF or OSPFv3 interface or virtual-link to send and receive OSPF or OSPFv3 protocol packets on the interface when simple password authentication is configured.

All neighboring routers must use the same type of authentication and password for proper protocol communication. If the **authentication-type** is configured as **password**, this key must be configured.

By default, no authentication key is configured.

The **no** form of this command removes the authentication key.

Default

no authentication-key

Parameters

authentication-key

Specifies the authentication key. The key can be any combination of ASCII characters up to 8 characters (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 22 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" "). This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted hash form is assumed.

authentication-type

Syntax

authentication-type {password | message-digest}

no authentication-type

Context

```
config>router>ospf>area>interface  
config>router>ospf3>area>virtual-link
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables authentication and specifies the type of authentication to be used on the OSPF or OSPFv3 interface.

Both **password** and **message-digest** authentication are supported.

By default, authentication is not enabled on an interface.

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

Default

no authentication

Parameters

password

Keyword to enable simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple **password** authentication is enabled.

message-digest

Keyword to enable message digest MD5 authentication in accordance with RFC 1321. If this option is configured, at least one message digest key must be configured.

bfd-enable

Syntax

```
[no] bfd-enable [remain-down-on-failure]
```

Context

```
config>router>ospf>area>interface
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the use of bidirectional forwarding (BFD) to control the state of the associated OSPF interface. By enabling BFD on an OSPF interface, the state of the interface is tied to the state of the BFD session between the local node and the remote node.

The optional **remain-down-on-failure** parameter can be specified on OSPF interfaces that are enabled for BFD to keep OSPF from reaching the full state if the BFD session to that neighbor cannot be established. This option is disabled by default and should be used only if there is a chance that unicast packets might be discarded while multicast packets are forwarded.

The **no** form of this command removes BFD from the associated OSPF adjacency.

Default

no bfd-enable

dead-interval

Syntax

dead-interval *seconds*

no dead-interval

Context

config>router>ospf>area>interface

config>router>ospf>area>virtual-link

config>router>ospf3>area>interface

config>router>ospf3>area>virtual-link

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the time, in seconds, that OSPF or OSPFv3 waits before declaring a neighbor router down. If no hello packets are received from a neighbor for the duration of the dead interval, the router is assumed to be down. The minimum interval must be two times the hello interval.

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

Default

dead-interval 40

Special Cases

OSPF or OSPFv3 Interface

If the **dead-interval** configured applies to an interface, all nodes on the subnet must have the same dead interval.

Virtual Link

If the **dead-interval** configured applies to a virtual link, the interval on both termination points of the virtual link must have the same dead interval.

Parameters

seconds

Specifies the dead interval, in seconds, expressed as a decimal integer.

Values 1 to 65535

export

Syntax

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

Context

config>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures export routing policies that determine the routes exported from the routing table to OSPF.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered overrides the previous command. A maximum of five policy names can be specified.

If an **aggregate** command is also configured in the **config>router** context, the aggregation is applied before the export policy is applied.

Routing policies are created in the **config>router>policy-options** context.

The **no** form of this command removes the specified *policy-name* or all policies from the configuration if no *policy-name* is specified.

Default

no export

Parameters

policy-name

Specifies the export policy name. Up to five *policy-name* arguments can be specified.

export-limit

Syntax

export-limit *number* [**log** *percentage*]

no export-limit

Context

config>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum number of routes (prefixes) that can be exported into OSPF from the route table.

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

Default

no export-limit

Parameters

number

Specifies the maximum number of routes (prefixes) that can be exported into OSPF from the route table.

Values 1 to 4294967295

percentage

Specifies the percentage of the **export-limit**, at which a warning log message and SNMP notification would be sent.

Values 1 to 100

hello-interval

Syntax

hello-interval *seconds*

no hello-interval

Context

config>router>ospf>area>interface

config>router>ospf>area>virtual-link

config>router>ospf3>area>interface

config>router>ospf3>area>virtual-link

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the interval between OSPF or OSPFv3 hellos issued on the interface or virtual link.

The **hello-interval**, in combination with the **dead-interval**, is used to establish and maintain the adjacency. Use this parameter to edit the frequency that hello packets are sent.

Reducing the interval, in combination with an appropriate reduction in the associated **dead-interval**, allows for faster detection of link or router failures at the cost of higher processing costs.

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

Default

hello-interval 10

Special Cases

OSPF Interface

If the **hello-interval** configured applies to an interface, all nodes on the subnet must have the same hello interval.

Virtual Link

If the **hello-interval** configured applies to a virtual link, the interval on both termination points of the virtual link must have the same hello interval.

Parameters

seconds

Specifies the hello interval, in seconds, expressed as a decimal integer.

Values 1 to 65535

interface

Syntax

[no] **interface** *ip-int-name* [**secondary**]

Context

config>router>ospf>area

config>router>ospf3>area

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an OSPF or OSPFv3 interface.

By default, interfaces are not activated in any interior gateway protocol, such as OSPF or OSPFv3, unless explicitly configured.

The **no** form of this command deletes the OSPF interface configuration for this interface. The **shutdown** command in the **config>router>ospf>interface** context can be used to disable an interface without removing the configuration for the interface.

Default

no interface

Parameters

ip-int-name

Specifies the IP interface name. Interface names must be unique within the group of defined IP interfaces for the **config>router>interface** command. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

If the IP interface name does not exist or does not have an IP address configured, an error message will be returned.

If the IP interface exists in a different area it will be moved to this area.

secondary

Keyword to enable multiple secondary adjacencies to be established over a single IP interface.

interface-type

Syntax

interface-type {**broadcast** | **point-to-point**}

no interface-type

Context

config>router>ospf>area>interface

config>router>ospf3>area>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the interface type to be either broadcast or point-to-point.

Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the broadcast adjacency maintenance overhead of the Ethernet link provided the link is used as a point-to-point.

If the interface type is not known at the time the interface is added to OSPF or OSPFv3, and subsequently the IP interface is bound (or moved) to a different interface type, this command must be entered manually.

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

Default

interface-type broadcast (if the physical interface is Ethernet or unknown)

Special Cases

Virtual-Link

A virtual link is always regarded as a point-to-point interface and not configurable.

Parameters

broadcast

Keyword to configure the interface to maintain this link as a broadcast network. To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.

point-to-point

Keyword to configure the interface to maintain this link as a point-to-point link.

loopfree-alternate-exclude

Syntax

[no] **loopfree-alternate-exclude**

Context

config>router>ospf>area>interface

config>router>ospf3>area>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command instructs IGP to exclude a specific interface or all interfaces that are participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This reduces LFA SPF calculation where it is not needed.

When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the **loopfree-alternate-exclude** command can only be executed under the area in which the specified interface is primary. If the command is enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

The **no** form of this command reinstates the default value.

Default

no loopfree-alternate-exclude

message-digest-key

Syntax

message-digest-key *keyid* **md5** [*key* | *hash-key*] [**hash** | **hash2**]

no message-digest-key *keyid*

Context

config>router>ospf>area>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a message digest key when MD5 authentication is enabled on the interface. Multiple message digest keys can be configured.

The **no** form of this command removes the message digest key identified by the *key-id*.

Default

no message-digest-key

Parameters

key-id

Specifies the message digest key, expressed as a decimal integer.

Values 1 to 255

md5 key

Specifies the MD5 key. The *key* can be any alphanumeric string up to 16 characters.

md5 hash-key

Specifies the MD5 hash key. The key can be any combination of ASCII characters up to 32 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted hash form is assumed.

metric

Syntax

metric *metric*

no metric

Context

config>router>ospf>area>interface

config>router>ospf3>area>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an explicit route cost metric for the OSPF or OSPFv3 interface that overrides the metrics calculated based on the speed of the underlying link.

The **no** form of this command deletes the manually configured interface metric, so the interface uses the computed metric based on the **reference-bandwidth** command setting and the speed of the underlying link.

Default

no metric

Parameters

metric

Specifies the metric to be applied to the interface, expressed as a decimal integer.

Values 1 to 65535

mtu

Syntax

mtu *bytes*

no mtu

Context

config>router>ospf>area>interface

config>router>ospf3>area>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the OSPF or OSPFv3 interface MTU value used when negotiating an OSPF or OSPFv3 adjacency.

The operational OSPF or OSPFv3 MTU value is calculated as follows.

If this command is not configured, the OSPF or OSPFv3 interface operational MTU derives the MTU value from the IP interface MTU (which is derived from the port MTU); for example, port MTU minus 14 bytes for a null-encapsulated Ethernet port for OSPF (not OSPFv3). If the derived MTU value is less than 576 bytes, the OSPF interface operational MTU is set to 576 bytes. If a lower interface MTU is required, it must be explicitly configured using this command.

If this command is configured for OSPF (not OSPFv3):

- if the OSPF interface MTU is less than 576 bytes, it becomes the operational OSPF MTU, regardless of the port MTU value
- if the OSPF interface MTU is equal to or greater than 576 bytes, and the derived interface MTU is less than 576 bytes, the operational OSPF MTU is set to 576 bytes
- if the OSPF interface MTU is equal to or greater than 576 bytes, and the derived interface MTU is greater than 576 bytes, the operational OSPF MTU is set to the lesser of the values configured with this command and the derived MTU

The port MTU must be set to 512 bytes or higher, because OSPF cannot support port MTU values lower than 512 bytes.

If this command is configured for OSPFv3:

- the operational OSPFv3 MTU is set to the lesser of the values configured with this command and the derived MTU
- this applies only when the port MTU is set to 1280 bytes or higher, because OSPFv3 cannot support port MTU values less than 1280 bytes

To determine the actual packet size, add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF or OSPFv3 (IP) packet MTU configured with this command.

If this command is configured to a value less than the interface or port MTU value, the OSPF or OSPFv3 MTU value will be used to transmit OSPF packets.

The **no** form of this command uses the value derived from the MTU configured in the **config>port** context.

Default

no mtu

Parameters

bytes

Specifies the MTU to be used by OSPF or OSPFv3 for this logical interface, in byte.

Values OSPF: 512 to 9710 (9724 to 14) (depends on the physical media)
 OSPFv3: 1280 to 9710 (9724 to 14) (depends on the physical media)

node-sid

Syntax

node-sid *index value*

node-sid *label value*

no node-sid

Context

config>router>ospf>area>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns a node SID index or label value to the prefix representing the primary address of an IPv4 network interface of type **loopback**. Only a single node SID can be assigned to an interface. The secondary address of an IPv4 interface cannot be assigned a node SID index and does not inherit the SID of the primary IPv4 address.

This command fails if the network interface is not of type **loopback** or if the interface is defined in an IES or a VPRN context. Also, assigning the same SID index or label value to the same interface in two different IGP instances is not allowed within the same node.

The value of the label or index SID is taken from the range configured for this IGP instance. When using the global mode of operation, a new segment routing module checks that the same index or label value cannot be assigned to more than one loopback interface address. When using the per-instance mode of operation, this check is not required because the index and label ranges of the various IGP instance are not allowed to overlap.

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

Default

no node-sid

Parameters

value

Specifies the node SID index or label value.

Values 0 to 4294967295

passive

Syntax

[no] passive

Context

```
config>router>ospf>area>interface  
config>router>ospf3>area>interface
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds the passive property to the OSPF or OSPFv3 interface where passive interfaces are advertised as OSPF or OSPFv3 interfaces but do not run the OSPF or OSPFv3 protocol.

By default, only interface addresses that are configured for OSPF or OSPFv3 will be advertised as OSPF or OSPFv3 interfaces. The **passive** parameter allows an interface to be advertised as an OSPF or OSPFv3 interface without running the OSPF or OSPFv3 protocol.

While in passive mode, the interface will ignore ingress OSPF or OSPFv3 protocol packets and not transmit any OSPF or OSPFv3 protocol packets.

The **no** form of this command removes the passive property from the OSPF or OSPFv3 interface.

Default

no passive

priority

Syntax

priority *number*
no priority

Context

```
config>router>ospf>area>interface  
config>router>ospf3>area>interface
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the priority of the OSPF or OSPFv3 interface that is used in an election of the designated router on the subnet.

This parameter is only used if the interface is of type broadcast. The router with the highest priority interface becomes the designated router. A router with priority 0 is not eligible to be the designated router or backup designated touter.

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

Default

priority 1

Parameters

number

Specifies the interface priority, expressed as a decimal integer.

Values 0 to 255

retransmit-interval

Syntax

retransmit-interval *seconds*

no retransmit-interval

Context

config>router>ospf>area>interface

config>router>ospf>area>virtual-link

config>router>ospf3>area>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the length of time that OSPF or OSPFv3 will wait before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor.

The value should be longer than the expected round trip delay between any two routers on the attached network. When the retransmit interval expires and no acknowledgement has been received, the LSA will be retransmitted.

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

Default

retransmit-interval 5

Parameters

seconds

Specifies the retransmit interval, in seconds, expressed as a decimal integer.

Values 1 to 1800

transit-delay

Syntax

transit-delay *seconds*

no transit-delay

Context

config>router>ospf>area>interface

config>router>ospf>area>virtual-link

config>router>ospf3>area>interface

config>router>ospf3>area>virtual-link

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the estimated time, in seconds, that it takes to transmit an LSA on the interface or virtual link.

The **no** form of this command reverts to the default delay time

Default

transit-delay 1

Parameters

seconds

Specifies the transit delay, in seconds, expressed as a decimal integer.

Values 1 to 1800

virtual-link

Syntax

[no] **virtual-link** *router-id* **transit-area** *area-id*

Context

config>router>ospf>area

config>router>ospf3>area

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a virtual link to connect area border routers to the backbone through a virtual link.

The backbone area (area 0.0.0.0) must be contiguous and all other areas must be connected to the backbone area. If it is not practical to connect an area to the backbone, the area border routers must be connected via a virtual link. The two area border routers will form a point-to-point like adjacency across the transit area. A virtual link can only be configured while in the area 0.0.0.0 context.

The *router-id* specified in this command must be associated with the virtual neighbor. The transit area cannot be a stub area or a NSSA.

The **no** form of this command deletes the virtual link.

Default

no virtual-link

Parameters

router-id

Specifies the router ID of the virtual neighbor in IP address dotted decimal notation.

transit-area area-id

Specifies the area ID specified identifies the transit area that links the backbone area to the area that has no physical connection with the backbone, expressed in dotted decimal notation or as a 32-bit decimal integer.

Values	a.b.c.d (dotted-decimal)
	0 to 4294967295 (decimal integer)

3.12.2.2 Show commands

```
ospf
```

Syntax

```
ospf
```

Context

```
show>router
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context display OSPF information.

ospf3

Syntax

ospf

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context display OSPFv3 information.

area

Syntax

area [*area-id*] [**detail**]

Context

show>router>ospf

show>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays configuration information about all areas or the specified area. When **detail** is specified, operational and statistical information will be displayed.

Parameters

area-id

Specifies the OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

detail

Displays detailed information about the area.

Output

The following output is an example of area information, and [Table 29: Output fields: OSPF area](#) describes the output fields.

Sample output

```
A:7210# show router ospf area detail
=====
OSPF Areas (Detailed)
=====
-----
Area Id: 0.0.0.0
-----
Area Id      : 0.0.0.0          Type      : Standard
Virtual Links : 0              Total Nbrs : 2
Active IFs    : 3              Total IFs  : 3
Area Bdr Rtrs : 0              AS Bdr Rtrs : 0
SPF Runs      : 7              Last SPF Run : 10/26/2006 10:09:18
Router LSAs   : 3              Network LSAs : 3
Summary LSAs  : 0              Asbr-summ LSAs : 0
Nssa ext LSAs : 0              Area opaque LSAs : 3
Total LSAs    : 9              LSA Cksum Sum : 0x28b62
Blackhole Range : True         Unknown LSAs : 0
=====
*A:Bombadil# show router ospf area 0.0.0.0 detail
=====
OSPF Area (Detailed) : 0.0.0.0
=====
-----
Configuration
-----
Area Id      : 0.0.0.0          Type      : Standard
-----
Statistics
-----
Virtual Links : 0              Total Nbrs : 2
Active IFs    : 3              Total IFs  : 3
Area Bdr Rtrs : 0              AS Bdr Rtrs : 0
SPF Runs      : 7              Last SPF Run : 10/26/2006 10:09:18
Router LSAs   : 3              Network LSAs : 3
Summary LSAs  : 0              Asbr-summ LSAs : 0
Nssa ext LSAs : 0              Area opaque LSAs : 3
Total LSAs    : 9              LSA Cksum Sum : 0x28b62
Blackhole Range : True         Unknown LSAs : 0
=====
```

Table 29: Output fields: OSPF area

Label	Description
Area Id	Displays a 32 bit integer uniquely identifying an area
Type	NSSA — This area is configured as an NSSA area Standard — This area is configured as a standard area (not NSSA or stub) Stub — This area is configured as a stub area
SPF Runs	Displays the number of times that the intra-area route table has been calculated using this area link-state database

Label	Description
LSA Count	Displays the total number of link-state advertisements in this area link-state database, excluding AS external LSAs
LSA Cksum Sum	Displays the 32-bit unsigned sum of the link-state database advertisements LS checksums contained in this area link-state database. This checksum excludes AS external LSAs (type-5).
No. of OSPF Areas	Displays the number of areas configured on the router
Virtual Links	Displays the number of virtual links configured through this transit area
Active IFs	Displays the active number of interfaces configured in this area
Area Bdr Rtrs	Displays the total number of ABRs reachable within this area
AS Bdr Rtrs	Displays the total number of ASBRs reachable within this area
Last SPF Run	Displays the time when the last intra-area SPF was run on this area
Router LSAs	Displays the total number of router LSAs in this area
Network LSAs	Displays the total number of network LSAs in this area
Summary LSAs	Displays the summary of LSAs in this area
Asbr-summ LSAs	Displays the summary of ASBR LSAs in this area
Nssa-ext LSAs	Displays the total number of NSSA-EXT LSAs in this area
Area opaque LSAs	Displays the total number of opaque LSAs in this area
Total Nbrs	Displays the total number of neighbors in this area
Total IFs	Displays the total number of interfaces configured in this area
Total LSAs	Displays the sum of LSAs in this area excluding autonomous system external LSAs
Blackhole Range	False — No blackhole route is installed for aggregates configured in this area

database

Syntax

database [**type** {**router** | **network** | **summary** | **asbr-summary** | **external** | **nssa** | **all**}] [**area** *area-id*] [**adv-router** *router-id*] [*link-state-id*] [**detail**]

Context

```
show>router>ospf  
show>router>ospf3
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information about the OSPF or OSPFv3 LSDB.

When no command line options are specified, the command displays brief output for all database entries.

Parameters

type

Specifies to filter the OSPF or OSPFv3 LSDB information based on which type is specified of the following types: router, network, summary, asbr-summary, external, nssa, all.

type router

Displays only router (Type 1) LSAs in the LSDB.

type network

Displays only network (Type 2) LSAs in the LSDB.

type summary

Displays only summary (Type 3) LSAs in the LSDB.

type asbr-summary

Displays only ASBR summary (Type 4) LSAs in the LSDB.

type external

Displays only AS external (Type 5) LSAs in the LSDB. External LSAs are maintained globally and not per area. If the display of external links is requested, the area parameter, if present, is ignored.

type nssa

Displays only NSSA area-specific AS external (Type 7) LSAs in the LSDB.

type all

Displays all LSAs in the LSDB. The **all** keyword is intended to be used with either the **area area-id** or the **adv-router router-id [link-state-id]** parameters.

area area-id

Displays LSDB information associated with the specified OSPF area ID.

adv-router router-id [link-state-id]

Displays LSDB information associated with the specified advertising router. To further narrow the number of items displayed, the *link-state-id* can optionally be specified.

detail

Displays detailed information about the LSDB entries.

Output

The following output is an example of database information, and [Table 30: Output fields: OSPF database](#) describes the output fields.

Sample output

```
A:ALA-A# show router ospf database
=====
OSPF Link State Database (Type : All)
=====
Area Id      Type      Link State Id  Adv Rtr Id    Age  Sequence      Cksum
-----
0.0.0.0      Router    180.0.0.2      180.0.0.2     1800 0x800000b6 0xf54
0.0.0.0      Router    180.0.0.5      180.0.0.5     1902 0x8000009d 0xcb7c
0.0.0.0      Router    180.0.0.8      180.0.0.8     1815 0x8000009a 0x529b
0.0.0.0      Router    180.0.0.9      180.0.0.9     1156 0x80000085 0xd00f
0.0.0.0      Router    180.0.0.10     180.0.0.10    533  0x8000009d 0x3f1f
0.0.0.0      Router    180.0.0.11     180.0.0.11    137  0x80000086 0xc58f
0.0.0.0      Router    180.0.0.12     180.0.0.12    918  0x8000009d 0x4cf3
0.0.0.0      Router    180.0.0.13     180.0.0.13    1401 0x800000a2 0x879c
0.0.0.0      Network   180.0.53.28    180.0.0.28    149  0x80000083 0xe5cd
0.0.0.0      Network   180.0.54.28    180.0.0.28    1259 0x80000083 0xdad7
0.0.0.0      Summary   180.0.0.15     180.0.0.10    378  0x80000084 0xeba1
0.0.0.0      Summary   180.0.0.15     180.0.0.12    73   0x80000084 0xdfab
0.0.0.0      Summary   180.0.0.18     180.0.0.10    1177 0x80000083 0xcfbf
0.0.0.1      Summary   180.100.25.4   180.0.0.12    208  0x80000091 0x3049
0.0.0.1      AS Summ   180.0.0.8      180.0.0.10    824  0x80000084 0x3d07
0.0.0.1      AS Summ   180.0.0.8      180.0.0.12    1183 0x80000095 0x4bdf
0.0.0.1      AS Summ   180.0.0.9      180.0.0.10    244  0x80000082 0x73cb
n/a         AS Ext    7.1.0.0        180.0.0.23    1312 0x80000083 0x45e7
n/a         AS Ext    7.2.0.0        180.0.0.23    997  0x80000082 0x45e6
n/a         AS Ext    10.20.0.0      180.0.0.23    238  0x80000081 0x2d81
...
-----
No. of LSAs: 339
=====
A:ALA-A#

A:ALA-A# show router ospf database detail
=====
OSPF Link State Database (Type : All) (Detailed)
-----
Router LSA for Area 0.0.0.0
-----
Area Id      : 0.0.0.0          Adv Router Id : 180.0.0.2
Link State Id : 180.0.0.2      LSA Type      : Router
Sequence No   : 0x800000b7    Checksum      : 0xd55
Age           : 155         Length        : 192
Options       : E
Flags         : None
Link Type (1) : Point To Point
Nbr Rtr Id (1) : 180.0.0.13
No of TOS (1) : 0
Link Type (2) : Stub Network
Network (2)    : 180.0.22.0
No of TOS (2) : 0
Link Type (3) : Point To Point
Nbr Rtr Id (3) : 180.0.0.12
No of TOS (3) : 0
Link Type (4) : Stub Network
Network (4)    : 180.0.5.0
No of TOS (4) : 0
Link Count    : 14
I/F Address (1) : 180.0.22.2
Metric-0 (1)   : 25
Mask (2)       : 255.255.255.0
Metric-0 (2)   : 25
I/F Address (3) : 180.0.5.2
Metric-0 (3)   : 25
Mask (4)       : 255.255.255.0
Metric-0 (4)   : 25
```

```

Link Type (5)      : Point To Point
Nbr Rtr Id (5)    : 180.0.0.8      I/F Address (5) : 180.0.13.2
No of TOS (5)     : 0              Metric-0 (5)    : 6
Link Type (6)     : Stub Network
Network (6)       : 180.0.13.0     Mask (6)        : 255.255.255.0
No of TOS (6)     : 0              Metric-0 (6)    : 6
Link Type (7)     : Point To Point
Nbr Rtr Id (7)    : 180.0.0.5     I/F Address (7) : 180.0.14.2
No of TOS (7)     : 0              Metric-0 (7)    : 6
Link Type (8)     : Stub Network
Network (8)       : 180.0.14.0     Mask (8)        : 255.255.255.0
No of TOS (8)     : 0              Metric-0 (8)    : 6
Link Type (9)     : Point To Point
Nbr Rtr Id (9)    : 180.0.0.11    I/F Address (9) : 180.0.17.2
No of TOS (9)     : 0              Metric-0 (9)    : 25
Link Type (10)    : Stub Network
Network (10)      : 180.0.17.0     Mask (10)       : 255.255.255.0
No of TOS (10)    : 0              Metric-0 (10)   : 25
Link Type (11)    : Stub Network
Network (11)      : 180.0.0.2      Mask (11)       : 255.255.255.255
No of TOS (11)    : 0              Metric-0 (11)   : 1
Link Type (12)    : Stub Network
Network (12)      : 180.0.18.0     Mask (12)       : 255.255.255.0
No of TOS (12)    : 0              Metric-0 (12)   : 24
Link Type (13)    : Point To Point
Nbr Rtr Id (13)   : 180.0.0.10    I/F Address (13) : 180.0.3.2
No of TOS (13)    : 0              Metric-0 (13)   : 25
Link Type (14)    : Stub Network
Network (14)      : 180.0.3.0      Mask (14)       : 255.255.255.0
No of TOS (14)    : 0              Metric-0 (14)   : 25
-----
AS Ext LSA for Network 180.0.0.14
-----
Area Id           : N/A            Adv Router Id    : 180.0.0.10
Link State Id     : 180.0.0.14    LSA Type        : AS Ext
Sequence No      : 0x80000083     Checksum        : 0xa659
Age              : 2033           Length          : 36
Options          : E
Network Mask      : 255.255.255.255 Fwding Address  : 180.1.6.15
Metric Type      : Type 2         Metric-0         : 4
Ext Route Tag     : 0
-----
...
A:ALA-A#

```

Table 30: Output fields: OSPF database

Label	Description
Area Id	Displays the OSPF area identifier
Type	Router — router LSA type (OSPF)
LSA Type	Network — network LSA type (OSPF)
	Summary — summary LSA type (OSPF)
	ASBR Summary — ASBR summary LSA type (OSPF)
	Nssa-ext — LSA area-specific, NSSA external (OSPF)

Label	Description
	Area opaque — area opaque LSA type (OSPF)
Link State Id	Displays the link-state ID. The link-state ID is an LSA type specific field containing either a number to distinguish several LSAs from the same router, an interface ID, or a router-id; it identifies the piece of the routing domain being described by the advertisement.
Adv Rtr Id Adv Router Id	Displays the router identifier of the router advertising the LSA
Age	Displays the age of the link state advertisement in seconds
Sequence Sequence No	Displays the signed 32-bit integer sequence number
Cksum Checksum	Displays the 32-bit unsigned sum of the link-state advertisements LS checksums
No. of LSAs	Displays the number of LSAs displayed
Options	EA — External attribute LSA support DC — Demand circuit support R — If clear, a node can participates in OSPF topology distribution without being used to forward transit traffic N — Type 7 LSA support E — External routes support
Prefix Options	P — Propagate NSSA LSA
Flags	None — No flags set V — The router is an endpoint for one or more fully adjacent virtual links having the described area as the transit area E — The router is an AS boundary router B — The router is an area border router
Link Count	Displays the number of links advertised in the LSA
Link Type (<i>n</i>)	Displays the link type of the <i>n</i> th link in the LSA
Network (<i>n</i>)	Displays the network address of the <i>n</i> th link in the LSA
Metric-0 (<i>n</i>)	Displays the cost metric of the <i>n</i> th link in the LSA

interface

Syntax

```
interface [ip-int-name | ip-address] [detail]  
interface [area area-id] [detail]
```

Context

```
show>router>ospf  
show>router>ospf3
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the details of the OSPF or OSPFv3 interface, this interface can be identified by IP address or IP interface name. When neither is specified, all in-service interfaces are displayed.

The **detail** option produces a great amount of data. Nokia recommends to detail only when requesting a specific interface.

Parameters

ip-addr
Displays only the interface identified by this IP address.

Values a.b.c.d

ip-int-name
Displays only the interface identified by this interface name, up to 32 characters.

area area-id
Displays all interfaces configured in this area.

Values ip-address: a.b.c.d
 area: 0 to 4294967295

detail
Displays detailed information about the interfaced.

Output

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

- Standard output: [Sample output](#), [Table 31: Output fields: OSPF interface](#)
- Detailed output: [Sample output — detailed](#), [Table 32: Output fields: OSPF interface detail](#)

Sample output

```

*A:JC-NodeA# show router ospf interface area detail
=====
OSPF Interfaces in Area (Detailed) : 1
=====
Interface : ip-10.10.1.1
-----
IP Address       : 10.10.1.1
Area Id          : 0.0.0.1
Hello Intrlvl    : 5 sec
Retrans Intrlvl  : 5 sec
Cfg Metric       : 0
Transit Delay    : 1
Passive          : False
Admin Status     : Enabled
Designated Rtr   : 10.20.1.1
IF Type          : Broadcast
Oper MTU         : 1500
Oper Metric      : 1000
Nbr Count        : 0
Tot Rx Packets   : 0
Rx Hellos        : 0
Rx DBDs          : 0
Rx LSRs          : 0
Rx LSUs          : 0
Rx LS Acks       : 0
Retransmits      : 0
Bad Networks     : 0
Bad Areas        : 0
Bad Auth Types   : 0
Bad Neighbors    : 0
Bad Lengths      : 0
Bad Dead Int.    : 0
Bad Versions     : 0
LSA Count        : 0
TE Metric        : 678
Priority          : 1
Rtr Dead Intrlvl : 15 sec
Poll Intrlvl     : 120 sec
Advert Subnet     : True
Auth Type        : None
Cfg MTU          : 0
Oper State        : Designated Rtr
Backup Desig Rtr  : 0.0.0.0
Network Type      : Transit
Last Enabled     : 04/11/2007 16:06:27
If Events         : 5
Tot Tx Packets    : 1116
Tx Hellos         : 1116
Tx DBDs           : 0
Tx LSRs           : 0
Tx LSUs           : 0
Tx LS Acks        : 0
Discards          : 0
Bad Virt Links    : 0
Bad Dest Adrs     : 0
Auth Failures     : 0
Bad Pkt Types     : 0
Bad Hello Int.    : 0
Bad Options       : 0
Bad Checksums     : 0
LSA Checksum      : 0x0
=====
*A:JC-NodeA#

```

Table 31: Output fields: OSPF interface

Label	Description
If Name	Displays the interface name
Area Id	Displays a 32-bit integer uniquely identifying the area to which this interface is connected. Area ID 0.0.0.0 is used for the OSPF backbone.
D Rtr Id	Displays the IP interface address of the router identified as the designated router for the network in which this interface is configured. Set to 0.0.0.0 if there is no Designated router.
BD Rtr Id	The IP interface address of the router identified as the backup designated router for the network in which this interface is configured. Set to 0.0.0.0 if there is no backup designated router.
Adm	Dn — OSPF on this interface is administratively shut down

Label	Description
	Up — OSPF on this interface is administratively enabled
Opr	<p>Down — This is the initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable.</p> <p>Wait — The router is trying to determine the identity of the (Backup) designated router for the network</p> <p>PToP — The interface is operational, and connects either to a physical point-to-point network or to a virtual link</p> <p>DR — This router is the designated router for this network</p> <p>BDR — This router is the backup designated router for this network</p> <p>ODR — The interface is operational and part of a broadcast or NBMA network on which another router has been selected to be the designated router</p>
No. of OSPF Interfaces	Displays the number of interfaces listed

Sample output — detailed

```
A:SetupCLI# show router ospf interface detail
=====
OSPF Interfaces (Detailed)
-----
Interface : system
-----
IP Address       : 10.1.255.255
Area Id          : 0.0.0.0
Hello Intrvl     : 10 sec
Retrans Intrvl   : 5 sec
Cfg Metric       : 0
Transit Delay    : 1
Passive          : True Cfg MTU
Admin Status     : Enabled
Designated Rtr   : 2.2.2.2
IF Type          : Broadcast
Oper MTU         : 1500
Oper Metric      : 0
Nbr Count        : 0
Tot Rx Packets   : 0
Rx Hellos        : 0
Rx DBDs          : 0
Rx LSRs          : 0
Rx LSUs          : 0
Rx LS Acks       : 0
Retransmits      : 0
Bad Networks     : 0
Bad Areas        : 0
Bad Auth Types   : 0
Bad Neighbors    : 0
Bad Lengths      : 0
Bad Dead Int.    : 0
Bad Versions     : 0
LSA Count        : 0
Priority         : 1
Rtr Dead Intrvl  : 40 sec
Poll Intrvl      : 120 sec
Advert Subnet    : True
Auth Type        : None
                 : 0
Oper State       : Designated Rtr
Backup Desig Rtr : 0.0.0.0
Network Type     : Transit
Last Enabled     : 05/14/2006 09:16:26
If Events        : 5
Tot Tx Packets   : 0
Tx Hellos        : 0
Tx DBDs          : 0
Tx LSRs          : 0
Tx LSUs          : 0
Tx LS Acks       : 0
Discards         : 0
Bad Virt Links   : 0
Bad Dest Addrs   : 0
Auth Failures    : 0
Bad Pkt Types    : 0
Bad Hello Int.   : 0
Bad Options      : 0
Bad Checksums    : 0
LSA Checksum     : 0x0
```

```

-----
Interface : sender
-----
IP Address       : 10.1.1.1
Area Id          : 0.0.0.0
Hello Intrvl     : 10 sec
Retrans Intrvl   : 5 sec
Cfg Metric       : 0
Transit Delay    : 1
Passive          : False
Priority         : 1
Rtr Dead Intrvl  : 40 sec
Poll Intrvl      : 120 sec
Advert Subnet    : True
Auth Type        : None
Cfg MTU          : 0
=====
A:SetupCLI#

```

Table 32: Output fields: OSPF interface detail

Label	Description
Interface	Displays the IP address of this OSPF interface
IP Address	Displays the IP address and mask of this OSPF interface
Interface Name	Displays the interface name
Area Id	Displays a 32-bit integer uniquely identifying the area to which this interface is connected. Area ID 0.0.0.0 is used for the OSPF backbone.
Priority	Displays the priority of this interface. Used in multi-access networks, this field is used in the designated router election algorithm.
Hello Intrvl	Displays the length of time, in seconds, between the Hello packets that the router sends on the interface. This value must be the same for all routers attached to a common network.
Rtr Dead Intrvl	Displays the number of seconds that router Hello packets have not been seen before its neighbors declare the router down. This should be some multiple of the Hello interval. This value must be the same for all routers attached to a common network.
Retrans Intrvl	Displays the number of seconds between link-state advertisement retransmissions, for adjacencies belonging to this interface. This value is also used when retransmitting database description and link-state request packets.
Poll Intrvl	Displays the larger time interval, in seconds, between the Hello packets sent to an inactive non-broadcast multi-access neighbor
Metric	Displays the metric to be advertised for this interface
Advert Subnet	<p>False — When a point-to-point interface is configured as false, then the subnet is not advertised and the endpoints are advertised as host routes</p> <p>True — When a point-to-point interface is configured to true, then the subnet is advertised</p>

Label	Description
Transit Delay	Displays the estimated number of seconds it takes to transmit a link state update packet over this interface
Auth Type	<p>Displays the authentication procedure to be used for the packet</p> <p>None — Routing exchanges over the network/subnet are not authenticated</p> <p>Simple — A 64-bit field is configured on a per-network basis. All packets sent on a particular network must have this configured value in their OSPF header 64-bit authentication field. This essentially serves as a "clear" 64-bit password.</p> <p>MD5 — A shared secret key is configured in all routers attached to a common network/subnet. For each OSPF protocol packet, the key is used to generate/verify a "message digest" that is appended to the end of the OSPF packet.</p>
Passive	<p>False — This interfaces operates as a normal OSPF interface with regard to adjacency forming and network/link behavior</p> <p>True — No OSPF hellos will be sent out on this interface and the router advertises this interface as a stub network/link in its router LSAs</p>
MTU	Displays the desired size of the largest packet that can be sent/received on this OSPF interface, specified in octets. This size DOES include the underlying IP header length, but not the underlying layer headers/trailers.
Admin Status	<p>Disabled — OSPF on this interface is administratively shut down</p> <p>Enabled — OSPF on this interface is administratively enabled</p>
Oper State	<p>Down — This is the initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable</p> <p>Waiting — The router is trying to determine the identity of the (backup) designated router for the network</p> <p>Point To Point — The interface is operational, and connects either to a physical point-to-point network or to a virtual link</p> <p>Designated Rtr — This router is the Designated Router for this network</p> <p>Other Desig Rtr — The interface is operational and part of a broadcast or NBMA network on which another router has been selected to be the designated router</p> <p>Backup Desig Rtr — This router is the backup designated router for this network</p>
DR-Id	Displays the IP Interface address of the router identified as the designated router for the network in which this interface is configured. Set to 0.0.0.0 if there is no designated router

Label	Description
BDR-Id	The IP interface address of the router identified as the backup designated router for the network in which this interface is configured. Set to 0.0.0.0 if there is no backup designated router.
IF Type	Broadcast — LANs, such as Ethernet NBMA — X.25 and similar technologies Point-To-Point — Links that are definitively point to point
Network Type	Stub — OPSF has not established a neighbor relationship with any other OSPF router on this network as such only traffic sourced or destined to this network will be routed to this network Transit — OPSF has established at least one neighbor relationship with any other OSPF router on this network as such traffic en route to other networks may be routed via this network
Oper MTU	Displays the operational size of the largest packet which can be sent/received on this OSPF interface, specified in octets. This size DOES include the underlying IP header length, but not the underlying layer headers/trailers.
Last Enabled	Displays the time that this interface was last enabled to run OSPF on this interface
Nbr Count	Displays the number of OSPF neighbors on the network for this interface
If Events	Displays the number of times this OSPF interface has changed its state, or an error has occurred since this interface was last enabled
Tot Rx Packets	Displays the total number of OSPF packets received on this interface since this interface was last enabled
Tot Tx Packets	Displays the total number of OSPF packets transmitted on this interface since this interface was last enabled
Rx Hellos	Displays the total number of OSPF Hello packets received on this interface since this interface was last enabled
Tx Hellos	Displays the total number of OSPF Hello packets transmitted on this interface since this interface was last enabled
Rx DBDs	Displays the total number of OSPF database description packets received on this interface since this interface was last enabled
Tx DBDs	Displays the total number of OSPF database description packets transmitted on this interface since this interface was last enabled
Rx LSRs	Displays the total number of link-state requests (LSRs) received on this interface since this interface was last enabled

Label	Description
Tx LSRs	Displays the total number of LSRs transmitted on this interface since this interface was last enabled
Rx LSUs	Displays the total number of Link State Updates (LSUs) received on this interface since this interface was last enabled
Tx LSUs	Displays the total number of LSUs transmitted on this interface since this interface was last enabled
Rx LS Acks	Displays the total number of link state acknowledgments received on this interface since this interface was last enabled
Tx LS Acks	Displays the total number of link state acknowledgments transmitted on this interface since this interface was last enabled
Retransmits	Displays the total number of OSPF retransmits sent on this interface since this interface was last enabled
Discards	Displays the total number of OSPF packets discarded on this interface since this interface was last enabled
Bad Networks	Displays the total number of OSPF packets received with invalid network or mask since this interface was last enabled
Bad Virt Links	Displays the total number of OSPF packets received on this interface that are destined to a virtual link that does not exist since this interface was last enabled
Bad Areas	Displays the total number of OSPF packets received with an area mismatch since this interface was last enabled
Bad Dest Addr	Displays the total number of OSPF packets received with the incorrect IP destination address since this interface was last enabled
Bad Auth Types	Displays the total number of OSPF packets received with an invalid authorization type since this interface was last enabled
Auth Failures	Displays the total number of OSPF packets received with an invalid authorization key since this interface was last enabled
Bad Neighbors	Displays the total number of OSPF packets received where the neighbor information does not match the information this router has for the neighbor since this interface was last enabled
Bad Pkt Types	Displays the total number of OSPF packets received with an invalid OSPF packet type since this interface was last enabled
Bad Lengths	Displays the total number of OSPF packets received on this interface with a total length not equal to the length specified in the packet since this interface was last enabled

Label	Description
Bad Hello int.	Displays the total number of OSPF packets received where the hello interval specified in packet was not equal to that configured on this interface since this interface was last enabled
Bad Dead Int.	Displays the total number of OSPF packets received where the dead interval specified in the packet was not equal to that configured on this interface since this interface was last enabled
Bad Options	Displays the total number of OSPF packets received with an option that does not match those configured for this interface or area since this interface was last enabled
Bad Versions	Displays the total number of OSPF packets received with bad OSPF version numbers since this interface was last enabled
Te Metric	Displays the TE metric configured for this interface. This metric is flooded out in the TE metric sub-TLV in the OSPF TE LSAs. Depending on the configuration, either the TE metric value or the native OSPF metric value is used in CSPF computations.
Te State	Displays the MPLS interface TE status from the OSPF standpoint
Admin Groups	Displays the bit-map inherited from the MPLS interface that identifies the admin groups to which this interface belongs

neighbor

Syntax

neighbor [*ip-int-name* | *ip-address*] [**detail**]

Context

show>router>ospf
show>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays all neighbor information. To reduce the amount of output, the user can select the neighbors on a specific interface by address or name.

The **detail** option produces a large amount of data. Nokia recommends to use **detail** only when requesting a specific neighbor.

Parameters

ip-address

Displays neighbor information for the neighbor identified by the specified IP address.

Values a.b.c.d

ip-int-name

Displays neighbor information only for neighbors of the interface identified by the interface name, up to 32 characters.

Output

The following outputs are examples of OSPF neighbor information. The associated tables describe the output fields.

- Standard output: [Sample output](#), [Table 33: Output fields: OSPF neighbor](#)
- Detailed output: [Sample output — detailed](#), [Table 34: Output fields: OSPF neighbor detail](#)

Sample output

```
A:ALA-A# show router ospf neighbor
=====
OSPF Neighbors
=====
Interface-Name          Rtr Id          State    Pri  RetxQ  TTL
-----
pc157-2/1               10.13.8.158     Full     1    0      37
pc157-2/2               10.13.7.165     Full    100    0      33
pc157-2/3               10.13.6.188     Full     1    0      38
-----
No. of Neighbors: 3
=====
A:ALA-A#
```

Table 33: Output fields: OSPF neighbor

Label	Description
Nbr IP Addr	Displays the IP address this neighbor is using in its IP Source Address. Note that, on addressless links, this will not be 0.0.0.0, but the address of another of the neighbor's interfaces.
Nbr Rtr Id	Displays a 32-bit integer uniquely identifying the neighboring router in the autonomous system
Nbr State	Down — This is the initial state of a neighbor conversation. It indicates that there has been no recent information received from the neighbor. Attempt — This state is only valid for neighbors attached to NBMA networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort should be made to contact the neighbor. Init — In this state, an Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet

Label	Description
	<p>been established with the neighbor (that is, the router did not appear in the neighbor Hello packet).</p> <p>Two Way — In this state, communication between the two routers is bidirectional.</p> <p>ExchStart — This is the first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide the initial database descriptor sequence number.</p> <p>Exchange — In this state, the router is describing its entire link state database by sending database description packets to the neighbor.</p> <p>Loading — In this state, link state request packets are sent to the neighbor asking for the more recent LSAs that have been discovered (but not yet received) in the exchange state.</p> <p>Full — In this state, the neighboring routers are fully adjacent. These adjacencies will now appear in router LSAs and network LSAs.</p>
Priority	Displays the priority of this neighbor in the designated router election algorithm. The value 0 signifies that the neighbor is not eligible to become the designated router on this particular network.
RetxQ Len	Displays the current length of the retransmission queue
Dead Time	Displays the time until this neighbor is declared down; this timer is set to the dead router interval when a valid hello packet is received from the neighbor
No. of Neighbors	Displays the number of adjacent OSPF neighbors on this interface

Sample output — detailed

```

A:ALA-A# show router ospf neighbor detail
=====
OSPF Neighbors
-----
Neighbor Rtr Id   : 10.13.8.158           Interface: pc157-2/1
-----
Neighbor IP Addr  : 10.16.1.8
Local IF IP Addr  : 10.16.1.7
Area Id           : 0.0.0.0
Designated Rtr    : 0.0.0.0              Backup Desig Rtr : 0.0.0.0
Neighbor State     : Full                  Priority          : 1
Retrans Q Length  : 0                     Options           : -E--0-
Events            : 4                     Last Event Time   : 05/06/2006 00:11:16
Up Time           : 1d 18:20:20            Time Before Dead  : 38 sec
GR Helper         : Not Helping            GR Helper Age     : 0 sec
GR Exit Reason    : None                   GR Restart Reason : Unknown
Bad Nbr States    : 1                     LSA Inst fails    : 0

```

```

Bad Seq Nums      : 0
Bad Packets       : 0
Option Mismatches: 0
Num Restarts      : 0
Bad MTUs          : 0
LSA not in LSDB   : 0
Nbr Duplicates    : 0
Last Restart at   : Never
-----
Neighbor Rtr Id   : 10.13.7.165      Interface: pc157-2/2
-----
Neighbor IP Addr  : 10.12.1.3
Local IF IP Addr  : 10.12.1.7
Area Id           : 0.0.0.0
Designated Rtr    : 10.13.9.157      Backup Desig Rtr : 10.13.7.165
Neighbor State     : Full              Priority           : 100
Retrans Q Length   : 0                Options            : -E--0-
Events             : 4                Last Event Time    : 05/05/2006 01:39:13
Up Time            : 0d 16:52:27       Time Before Dead   : 33 sec
GR Helper          : Not Helping       GR Helper Age      : 0 sec
GR Exit Reason     : None              GR Restart Reason  : Unknown
Bad Nbr States     : 0                LSA Inst fails     : 0
Bad Seq Nums       : 0                Bad MTUs           : 0
Bad Packets        : 0                LSA not in LSDB    : 0
Option Mismatches  : 0                Nbr Duplicates     : 0
Num Restarts       : 0                Last Restart at    : Never
-----
Neighbor Rtr Id   : 10.13.6.188      Interface: pc157-2/3
-----
Neighbor IP Addr  : 10.14.1.4
Local IF IP Addr  : 10.14.1.7
Area Id           : 0.0.0.0
Designated Rtr    : 10.13.9.157      Backup Desig Rtr : 10.13.6.188
Neighbor State     : Full              Priority           : 1
Retrans Q Length   : 0                Options            : -E--0-
Events             : 4                Last Event Time    : 05/05/2006 08:35:18
Up Time            : 0d 09:56:25       Time Before Dead   : 38 sec
GR Helper          : Not Helping       GR Helper Age      : 0 sec
GR Exit Reason     : None              GR Restart Reason  : Unknown
Bad Nbr States     : 1                LSA Inst fails     : 0
Bad Seq Nums       : 0                Bad MTUs           : 0
Bad Packets        : 0                LSA not in LSDB    : 0
Option Mismatches  : 0                Nbr Duplicates     : 0
Num Restarts       : 0                Last Restart at    : Never
=====
A:ALA-A#

```

Table 34: Output fields: OSPF neighbor detail

Label	Description
Neighbor IP Addr	Displays the IP address this neighbor is using in its IP source address. Note that, on links with no address, this will not be 0.0.0.0, but the address of another of the neighbor interfaces.
Local IF IP Addr	Displays the IP address of this OSPF interface
Area Id	Displays a 32-bit integer uniquely identifying the area to which this interface is connected. Area ID 0.0.0.0 is used for the OSPF backbone
Designated Rtr	Displays the IP interface address of the router identified as the designated router for the network in which this interface is configured; set to 0.0.0.0 if there is no designated router

Label	Description
Neighbor Rtr Id	Displays a 32-bit integer uniquely identifying the neighboring router in the AS
Neighbor State	<p>Down — This is the initial state of a neighbor conversation. It indicates that there has been no recent information received from the neighbor</p> <p>Attempt — This state is only valid for neighbors attached to NBMA networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort should be made to contact the neighbor.</p> <p>Init — In this state, an Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor (that is, the router did not appear in the neighbor Hello packet).</p> <p>Two Way — In this state, communication between the two routers is bidirectional.</p> <p>Exchange start — This is the first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial database descriptor sequence number.</p> <p>Exchange — In this state the router is describing its entire link state database by sending database description packets to the neighbor</p> <p>Loading — In this state, link state request packets are sent to the neighbor asking for the more recent LSAs that have been discovered (but not yet received) in the exchange state.</p> <p>Full — In this state, the neighboring routers are fully adjacent. These adjacencies will now appear in router-LSAs and network-LSAs.</p>
Priority	Displays the priority of this neighbor in the designated router election algorithm. The value 0 signifies that the neighbor is not eligible to become the designated router on this particular network.
Retrans Q Length	Displays the current length of the retransmission queue
Options	<p>E — External routes support</p> <p>N/P — Type 7 LSA support</p> <p>EA — External attribute LSA support</p> <p>DC — Demand circuit support</p> <p>O — Opaque LSA support</p>
Backup Desig Rtr	Displays the IP Interface address of the router identified as the backup designated router for the network in which this interface

Label	Description
	is configured; set to 0.0.0.0 if there is no backup designated router
Events	Displays the number of times this neighbor relationship has changed state, or an error has occurred
Last Event Time	Displays the time when the last event occurred that affected the adjacency to the neighbor
Up Time	Displays the uninterrupted time, in hundredths of seconds, the adjacency to this neighbor has been up. To evaluate when the last state change occurred, see Last Event Time.
Time Before Dead	Displays the time until this neighbor is declared down; this timer is set to the dead router interval when a valid hello packet is received from the neighbor
Bad Nbr States	Displays the total number of OSPF packets received when the neighbor state was not expecting to receive this packet type since this interface was last enabled
LSA Inst fails	Displays the total number of times an LSA could not be installed into the LSDB due to a resource allocation issue since this interface was last enabled
Bad Seq Nums	Displays the total number of times when a database description packet was received with a sequence number mismatch since this interface was last enabled
Bad MTUs	Displays the total number of times when the MTU in a received database description packet was larger than the MTU of the receiving interface since this interface was last enabled
Bad Packets	Displays the total number of times when an LS update was received with an illegal LS type or an option mismatch since this interface was last enabled
LSA not in LSDB	Displays the total number of times when an LS request was received for an LSA not installed in the LSDB of this router since this interface was last enabled
Option Mismatches	Displays the total number of times when an LS update was received with an option mismatch since this interface was last enabled.
Nbr Duplicates	Displays the total number of times when a duplicate database description packet was received during the exchange state since this interface was last enabled

prefix-sids

Syntax

prefix-sids [*ip-prefix[/prefix-length]*] [*sid sid*] [*adv-router router-id*]

Context

show>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays OSPF prefix SIDs.

Parameters

ip-prefix[/prefix-length]
Displays information about the specified IP prefix and length, up to 64 characters.

sid
Displays information for the specific segment identifier.

Values 0 to 524287

router-id
Displays information for the specific advertising router identified by its router ID.

Output

The following output is an example of OSPF prefix SID information, and [Table 35: Output fields: prefix SIDs](#) describes the output fields.

Sample output

```
*A:Dut-C# show router ospf prefix-sids
=====
Rtr Base OSPFv2 Instance 0 Prefix-Sids
=====
```

Prefix	Area Adv-Rtr	RtType Active	SID Flags
-----	-----	-----	-----
10.20.1.1/32	0.0.0.0	INTER-AREA	11
	10.20.1.2	N	NnP
10.20.1.1/32	0.0.0.1	INTRA-AREA	11
	10.20.1.1	Y	NnP
10.20.1.2/32	0.0.0.0	INTRA-AREA	22
	10.20.1.2	Y	NnP
10.20.1.2/32	0.0.0.1	INTER-AREA	22
	10.20.1.2	N	NnP
10.20.1.3/32	0.0.0.0	INTRA-AREA	33
	10.20.1.3	Y	NnP
10.20.1.3/32	0.0.0.1	INTER-AREA	33
	10.20.1.2	N	NnP

```

10.20.1.4/32          0.0.0.0      INTRA-AREA  44
                    10.20.1.4      Y          NnP
10.20.1.4/32          0.0.0.1      INTER-AREA  44
                    10.20.1.2      N          NnP
10.20.1.5/32          0.0.0.0      INTRA-AREA  55
                    10.20.1.5      Y          NnP
10.20.1.5/32          0.0.0.1      INTER-AREA  55
                    10.20.1.2      N          NnP
10.20.1.6/32          0.0.0.0      INTER-AREA  66
                    10.20.1.4      N          NnP
10.20.1.6/32          0.0.0.0      INTER-AREA  66
                    10.20.1.5      Y          NnP
10.20.1.6/32          0.0.0.1      INTER-AREA  66
                    10.20.1.2      N          NnP
-----
No. of Prefix/SIDs: 13
Flags:  N = Node-SID
        nP = no penultimate hop POP
        M = Mapping server
        E = Explicit-Null
        V = Prefix-SID carries a value
        L = value/index has local significance
        I = Inter Area flag
        A = Attached flag
=====

*A:Dut-C# show router ospf prefix-sids sid 66
=====
Rtr Base OSPFv2 Instance 0 Prefix-Sids
=====
Prefix          Area      RtType      SID
                Adv-Rtr    Active      Flags
-----
10.20.1.6/32    0.0.0.0      INTER-AREA  66
                10.20.1.4      N          NnP
10.20.1.6/32    0.0.0.0      INTER-AREA  66
                10.20.1.5      Y          NnP
10.20.1.6/32    0.0.0.1      INTER-AREA  66
                10.20.1.2      N          NnP
-----
No. of Prefix/SIDs: 3
Flags:  N = Node-SID
        nP = no penultimate hop POP
        M = Mapping server
        E = Explicit-Null
        V = Prefix-SID carries a value
        L = value/index has local significance
        I = Inter Area flag
        A = Attached flag
=====

*A:Dut-C#

```

Table 35: Output fields: prefix SIDs

Label	Description
Prefix	Displays the IP prefix for the SID
Area	Displays the OSPF area
Adv-Rtr	Displays the advertised router IP address

Label	Description
RtType	Displays the type of route
Active	Displays the status of the route: active (Y) or inactive (N)
SID	Displays the segment routing identifier (SID)
Flags	Displays the flags related to the advertised router: R = Re-advertisement N = Node SID nP = No penultimate hop POP E = Explicit null V = Prefix-SID carries a value L = Value/index has local significance

range

Syntax

range [*area-id*]

Context

show>router>ospf
show>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays ranges of addresses on an area border router (ABR) for the purpose of route summarization or suppression.

Parameters

area-id
Displays the configured ranges for the specified area.

Output

The following output is an example of OSPF range information, and [Table 36: Output fields: OSPF range](#) describes the output fields.

Sample output

```
A:ALA-A# show router ospf range
=====
OSPF Ranges
```



```
=====
Area Id      Address/Mask      Advertise  LSDB Type
-----
No. of Ranges: 0
=====
A:ALA-A#

A:ALA-A# show router ospf range 180.0.7.9
=====
OSPF Ranges for Area Id : 180.0.7.9
=====
Area Id      Address/Mask      Advertise  LSDB Type
-----
No. of Ranges: 0
=====
A:ALA-A#
```

Table 36: Output fields: OSPF range

Label	Description
Area Id	Displays a 32-bit integer uniquely identifying an area. Area ID 0.0.0.0 is used for the OSPF backbone.
Address/Mask	Displays the mask for the range expressed as a decimal integer mask length or in dotted decimal notation
Advertise	False — The specified address/mask is not advertised outside the area True — The specified address/mask is advertised outside the area
LSDB Type	NSSA — This range was specified in the NSSA context, and specifies that the range applies to external routes (via type-7 LSAs) learned within the NSSA when the routes are advertised to other areas as type-5 LSAs. Summary — This range was not specified in the NSSA context, the range applies to summary LSAs even if the area is an NSSA.

routes

Syntax

routes [*ip-prefix[/prefix-length]*] [**type**] [**detail**] [**alternative**] [**summary**] [**exclude-shortcut**]

Context

show>router>ospf
show>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information about OSPF or OSPFv3 routes.

Parameters

- ip-prefix/prefix-length*

Specifies the IP address.
- type*

Displays information about the specified type.

Valuesintra-area, inter-area, external-1, external-2, nssa-1, nssa-2
- detail*

Displays detailed information about the routes.
- alternative*

Displays the level of protection per prefix.
- summary*

Displays summary information about the routes.
- exclude-shortcut*

Displays routes without shortcuts.

Output

The following output is an example of OSPF route information.

Sample output

```
A:ALU-A# show router ospf routes
=====
OSPFv2 (0) Routing Table
=====
Destination Type(Dest) Stat SID SIDflgs
NHIP NHIF Cost[E2]
-----
1.1.1.1/32 IA (HOST) N (R)
1.1.3.1 3 1000
1.1.2.0/24 IA (NET) N (R)
1.1.3.1 3 2000
1.2.3.2 4 2000
1.1.3.0/24 IA (NET) D (F)
DIRECT 3 1000
1.2.3.0/24 IA (NET) D (F)
DIRECT 4 1000
1.2.4.0/24 IA (NET) N (R)
2.2.3.2 5 2000
1.3.5.0/24 IA (NET) D (F)
DIRECT 6 1000
1.4.5.0/24 IA (NET) N (R)
1.3.5.5 6 2000
1.4.6.0/24 IE (NET) N (R)
2.2.3.2 5 3000
```

```

1.3.5.5 6 3000
1.5.6.0/24 IE (NET) N (R)
1.3.5.5 6 2000
2.2.2.2/32 IA (HOST) N (R)
2.2.3.2 5 1000
2.2.3.0/24 IA (NET) D (F)
DIRECT 5 1000
3.3.3.3/32 IA (HOST) D (F)
DIRECT 2 0
4.4.4.4/32 IA (HOST) N (R)
2.2.3.2 5 2000
1.3.5.5 6 2000
5.5.5.5/32 IA (HOST) N (R)
1.3.5.5 6 1000
6.6.6.6/32 IE (HOST) N (R)
1.3.5.5 6 2000
10.20.1.1/32 IA (HOST) N (R) 11 NnP
1.1.3.1 3 1000
10.20.1.2/32 IA (HOST) N (R) 22 NnP
2.2.3.2 5 1000
10.20.1.3/32 IA (HOST) D (F) 33 NnP
DIRECT 1 0
10.20.1.4/32 IA (HOST) N (R) 44 NnP
2.2.3.2 5 2000
1.3.5.5 6 2000
10.20.1.5/32 IA (HOST) N (R) 55 NnP
1.3.5.5 6 1000
10.20.1.6/32 IE (HOST) N (R) 66 NnP
1.3.5.5 6 2000
10.20.1.1/0 IA (RTR) N (N)
1.1.3.1 3 1000
10.20.1.2/0 IA (AB-AS) N (N)
2.2.3.2 5 1000
10.20.1.2/0 IA (AB-AS) N (N)
1.2.3.2 4 1000
10.20.1.4/0 IA (AB-AS) N (N)
2.2.3.2 5 2000
1.3.5.5 6 2000
10.20.1.5/0 IA (AB-AS) N (N)
1.3.5.5 6 1000

```

```

-----
No. of routes found: 26 (31 paths)
Stat: D = direct N = not direct
(RTM stat):(R) = added (F) = add failed
(N) = not added (D) = policy discarded
SID Flags : N = Node-SID
nP = no penultimate hop POP
M = Mapping server
E = Explicit-Null
V = Prefix-SID carries a value
L = value/index has local significance
I = Inter Area flag
A = Attached flag

```

```

=====
A:ALU-A#

```

```

A:ALU-A# show router ospf routes alternative detail

```

```

=====
OSPFv2 (0) Routing Table (detailed)
=====

```

```

Destination Type(Dest) Stat
NHIP NHIF Cost[E2] Area Tunnel-Information
A-NHIP(L) A-NHIF A-Cost[E2] A-Type PGID
-----

```

```
1.1.2.0/24 IA (NET) D (F)
DIRECT 2 10 0.0.0.0
1.1.3.0/24 IA (NET) D (F)
DIRECT 3 10 0.0.0.0
1.2.3.0/24 IA (NET) N (R)
1.1.2.2 2 20 0.0.0.0
1.1.3.3 3 20 0.0.0.0
1.2.4.0/24 IA (NET) N (R)
1.1.2.2 2 20 0.0.0.0
1.1.3.3(L) 3 30 LINK 0x130015
1.3.5.0/24 IA (NET) N (R)
1.1.3.3 3 20 0.0.0.0
1.1.2.2(L) 2 30 LINK 0x130016
1.4.5.0/24 IA (NET) N (R)
1.1.2.2 2 30 0.0.0.0
1.1.3.3 3 30 0.0.0.0
1.4.6.0/24 IA (NET) N (R)
1.1.2.2 2 30 0.0.0.0
1.1.3.3(L) 3 40 LINK 0x130015
1.5.6.0/24 IA (NET) N (R)
1.1.3.3 3 30 0.0.0.0
1.1.2.2(L) 2 40 LINK 0x130016
10.20.1.1/32 IA (HOST) D (F)
DIRECT 1 0 0.0.0.0
10.20.1.2/32 IA (HOST) N (R)
1.1.2.2 2 10 0.0.0.0
1.1.3.3(L) 3 20 LINK 0x130015
10.20.1.3/32 IA (HOST) N (R)
1.1.3.3 3 10 0.0.0.0
1.1.2.2(L) 2 20 LINK 0x130016
10.20.1.4/32 IA (HOST) N (R)
1.1.2.2 2 20 0.0.0.0
1.1.3.3(L) 3 30 LINK 0x130015
10.20.1.5/32 IA (HOST) N (R)
1.1.3.3 3 20 0.0.0.0
1.1.2.2(L) 2 30 LINK 0x130016
10.20.1.6/32 IA (HOST) N (R)
1.1.3.3 3 30 0.0.0.0
1.1.2.2 2 30 0.0.0.0
10.20.1.2/0 IA (RTR) N (N)
1.1.2.2 2 10 0.0.0.0
10.20.1.3/0 IA (RTR) N (N)
1.1.3.3 3 10 0.0.0.0
10.20.1.4/0 IA (RTR) N (N)
1.1.2.2 2 20 0.0.0.0
10.20.1.5/0 IA (RTR) N (N)
1.1.3.3 3 20 0.0.0.0
10.20.1.6/0 IA (RTR) N (N)
1.1.3.3 3 30 0.0.0.0
1.1.2.2 2 30 0.0.0.0
```

```
-----
19 OSPFv2 routes found (23 paths)
Flags: L = Loop-Free Alternate nexthop
Stat: D = direct N = not direct
(RTM stat):(R) = added (F) = add failed
(N) = not added (D) = policy discarded
=====
```

```
A:ALU-A#
```

spf

Syntax
spf

Context
show>router>ospf
show>router>ospf3

Platforms
Supported on all 7210 SAS platforms as described in this document

Description
This command displays statistics of shortest-path-first (SPF) calculations.

Output
The following output is an example of SPF information, and [Table 37: Output fields: OSPF SFP](#) describes the output fields.

Sample output

```
A:ALA-A# show router ospf spf
=====
OSPF SPF Statistics
=====
Total SPF Runs           : 109
Last Full SPF run @      : 11/07/2006 18:43:07
Last Full SPF Time       : < 0.01 secs
  Intra SPF Time         : < 0.01 secs
  Inter SPF Time         : < 0.01 secs
  Extern SPF Time        : < 0.01 secs
  RTM Updt Time          : < 0.01 secs

Min/Avg/Max Full SPF Times : 0.02/0.00/0.06 secs
Min/Avg/Max RTM Updt Times : 0.02/0.00/0.06 secs

Total Sum Incr SPF Runs : 333
Last Sum Incr SPF run @ : 11/07/2006 18:43:09
Last Sum Incr Calc Time : < 0.01 secs

Total Ext Incr SPF Runs : 0
=====
A:ALA-A#
```

Table 37: Output fields: OSPF SFP

Label	Description
Total SPF Runs	Displays the total number of incremental SPF runs triggered by new or updated LSAs

Label	Description
Last Full SPF run @	Displays the date and time when the external OSPF Dijkstra (SPF) was last run
Last Full SPF Time	Displays the length of time, in seconds, when the last full SPF was run
Intra SPF Time	Displays the time when intra-area SPF was last run on this area
Inter SPF Time	Displays the total number of incremental SPF runs triggered by new or updated type-3 and type-4 summary LSAs
Extern SPF Time	Displays the total number of incremental SPF runs triggered by new or updated type-5 external LSAs
RTM Updt Time	Displays the time, in hundredths of seconds, used to perform a total SPF calculation
Min/Avg/Max Full SPF Time	Min — The minimum time, in hundredths of seconds, used to perform a total SPF calculation Avg — The average time, in hundredths of seconds, of all the total SPF calculations performed by this OSPF router Max — The maximum time, in hundredths of seconds, used to perform a total SPF calculation
Total Sum Incr SPF Runs	Displays the total number of incremental SPF runs triggered by new or updated type-3 and type-4 summary LSAs
Total Ext Incr SPF Runs	Displays the total number of incremental SPF runs triggered by new or updated type-5 external LSAs

statistics

Syntax

statistics

Context

show>router>ospf
show>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the global OSPF or OSPFv3 statistics.

Output

The following output is an example of OSPF statistics information, and [Table 38: Output fields: OSPF statistics](#) describes the output fields.

Sample output

```
A:ALA-A# show router ospf statistics
=====
OSPF Statistics
=====
Rx Packets      : 308462      Tx Packets      : 246800
Rx Hellos       : 173796      Tx Hellos       : 149062
Rx DBDs         : 67          Tx DBDs         : 48
Rx LSRs         : 21          Tx LSRs         : 19
Rx LSUs         : 105672      Tx LSUs         : 65530
Rx LS Acks      : 28906      Tx LS Acks      : 32141
New LSAs Recvd  : 38113      New LSAs Orig   : 21067
Ext LSAs Count  : 17          No of Areas     : 3
Total SPF Runs  : 327         Ext SPF Runs    : 0
Retransmits     : 46          Discards        : 0
Bad Networks    : 0           Bad Virt Links  : 0
Bad Areas       : 0           Bad Dest Addrs  : 0
Bad Auth Types  : 0           Auth Failures   : 0
Bad Neighbors   : 0           Bad Pkt Types   : 0
Bad Lengths     : 0           Bad Hello Int.  : 0
Bad Dead Int.   : 0           Bad Options     : 0
Bad Versions    : 0           Bad Checksums   : 0
Failed SPF Attempts: 0
CSPF Requests   : 0           CSPF Request Drops : 0
CSPF Path Found : 0           CSPF Path Not Found: 0
=====
A:ALA-A#
```

Table 38: Output fields: OSPF statistics

Label	Description
Rx Packets	Displays the total number of OSPF packets received on all OSPF enabled interfaces
Tx Packets	Displays the total number of OSPF packets transmitted on all OSPF enabled interfaces
Rx Hellos	Displays the total number of OSPF Hello packets received on all OSPF enabled interfaces
Tx Hellos	Displays the total number of OSPF Hello packets transmitted on all OSPF enabled interfaces
Rx DBDs	Displays the total number of OSPF database description packets received on all OSPF enabled interfaces
Tx DBDs	Displays the total number of OSPF database description packets transmitted on all OSPF enabled interfaces

Label	Description
Rx LSRs	Displays the total number of OSPF link state requests (LSRs) received on all OSPF enabled interfaces
Tx LSRs	Displays the total number of OSPF link state requests (LSRs) transmitted on all OSPF enabled interfaces
Rx LSUs	Displays the total number of OSPF link state update (LSUs) received on all OSPF enabled interfaces
Tx LSUs	Displays the total number of OSPF link state update (LSUs) transmitted on all OSPF enabled interfaces
Rx LS Acks	Displays the total number of OSPF link state acknowledgments (LSAs) received on all OSPF enabled interfaces
New LSAs Recvd	Displays the total number of new OSPF link state advertisements received on all OSPF enabled interfaces
New LSAs Orig	Displays the total number of new OSPF link state advertisements originated on all OSPF enabled interfaces
Ext LSAs Count	Displays the total number of OSPF external link state advertisements
No of Areas	Displays the number of areas configured for this OSPF instance
Total SPF Runs	Displays the total number of incremental SPF runs triggered by new or updated LSAs
Ext SPF Runs	Displays the total number of incremental SPF runs triggered by new or updated type-5 external LSAs
Retransmits	Displays the total number of OSPF retransmits transmitted on all OSPF enabled interfaces
Discards	Displays the total number of OSPF packets discarded on all OSPF enabled interfaces
Bad Networks	Displays the total number of OSPF packets received on all OSPF enabled interfaces with invalid network or mask
Bad Virt Links	Displays the total number of OSPF packets received on all OSPF enabled interfaces that are destined to a virtual link that does not exist
Bad Areas	Displays the total number of OSPF packets received on all OSPF enabled interfaces with an area mismatch
Bad Dest Addr	Displays the total number of OSPF packets received on all OSPF enabled interfaces with the incorrect IP destination address

Label	Description
Bad Auth Types	Displays the total number of OSPF packets received on all OSPF enabled interfaces with an invalid authorization type
Auth Failures	Displays the total number of OSPF packets received on all OSPF enabled interfaces with an invalid authorization key
Bad Neighbors	Displays the total number of OSPF packets received on all OSPF enabled interfaces where the neighbor information does not match the information this router has for the neighbor
Bad Pkt Types	Displays the total number of OSPF packets received on all OSPF enabled interfaces with an invalid OSPF packet type
Bad Lengths	Displays the total number of OSPF packets received on all OSPF enabled interfaces with a total length not equal to the length specified in the packet
Bad Hello Int.	Displays the total number of OSPF packets received on all OSPF enabled interfaces where the hello interval specified in the packet was not equal to that configured for the respective interface
Bad Dead Int.	Displays the total number of OSPF packets received on all OSPF enabled interfaces where the dead interval specified in the packet was not equal to that configured for the respective interface
Bad Options	Displays the total number of OSPF packets received on all OSPF enabled interfaces with an option that does not match those configured for the respective interface or area
Bad Versions	Displays the total number of OSPF packets received on all OSPF enabled interfaces with bad OSPF version numbers

status

Syntax

status

Context

show>router>ospf
show>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the general status of OSPF or OSPFv3.

Output

The following output is an example of OSPF status information, and [Table 39: Output fields: OSPF status](#) describes the output fields.

Sample output

```
A:ALA-A# show router ospf status
=====
OSPF Status
=====
OSPF Router Id           : 10.13.7.165
OSPF Version             : 2
OSPF Admin Status        : Enabled
OSPF Oper Status         : Enabled
Graceful Restart         : Enabled
GR Helper Mode           : Disabled
Preference                : 10
External Preference      : 150
Backbone Router          : True
Area Border Router       : True
AS Border Router         : True
Opaque LSA Support       : True
Traffic Engineering Support : True
RFC 1583 Compatible      : True
TOS Routing Support      : False
Demand Exts Support      : False
In Overload State        : False
In External Overflow State : False
Exit Overflow Interval   : 0
Last Overflow Entered     : Never
Last Overflow Exit       : Never
External LSA Limit       : -1
Reference Bandwidth      : 100,000,000 Kbps
Init SPF Delay           : 500 msec
Sec SPF Delay            : 2000 msec
Max SPF Delay            : 15000 msec
Min LS Arrival Interval  : 500 msec
Max LSA Gen Delay        : 5000 msec
Last Ext SPF Run         : Never
Ext LSA Cksum Sum        : 0x2afce
OSPF Last Enabled        : 05/23/2006 23:34:36
Export Policies          : export-static
Import Policies          : None
Lfa Policies             : pol1
                        : pol2
                        : pol3
                        : pol4
                        : pol5
=====
A:ALA-A#
```

Table 39: Output fields: OSPF status

Label	Description
OSPF Router Id	Displays a 32-bit integer uniquely identifying the router in the autonomous system. The default is the system IP address, or if not configured, the 32 least significant bits of the system MAC address.
OSPF Version	Displays the current version number of the OSPF protocol is 2
OSPF Admin Status	Disabled — Specifies that the OSPF process is disabled on all interfaces. Enabled — Specifies that the OSPF process is active on at least one interface
OSPF Oper Status	Disabled — Specifies that the OSPF process is not operational on all interfaces. Enabled — Specifies that the OSPF process is operational on at least one interface
Preference	Displays the route preference for OSPF internal routes
External Preference	Displays the route preference for OSPF external routes
Backbone Router	False — This variable indicates that this router is not configured as an OSPF back bone router True — This variable indicates that this router is configured as an OSPF back bone router
Area Border Router	False — This router is not an area border router True — This router is an area border router
AS Border Router	False — This router is not configured as an autonomous system border router True — This router is configured as an autonomous system border router
OSPF Ldp Sync Admin Status	Indicates whether the IGP-LDP synchronization feature is enabled or disabled on all interfaces participating in the OSPF routing protocol.
Export Policies	Displays the export policies currently in use
Import Policies	Displays the import policies currently in use
LFA Policies	Displays the LFA policies currently in use

virtual-link

Syntax

virtual-link [detail]

Context

show>router>ospf
show>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information for OSPF or OSPFv3 virtual links.

Parameters

detail
Displays operational and statistical information about virtual links associated with this router.

Output

The following output is an example of OSPF virtual link information, and [Table 40: Output fields: OSPF virtual link](#) describes the output fields.

Sample output

```
A:ALA-A# show router ospf virtual-link
=====
OSPF Virtual Links
=====
Nbr Rtr Id      Area Id      Local Interface  Metric State
-----
180.0.0.10     0.0.0.1     180.1.7.12      300    PToP
180.0.0.10     0.0.0.2     180.2.7.12      300    PToP
-----
No. of OSPF Virtual Links: 2
=====
A:ALA-A#

A:ALA-A# show router ospf virtual-link detail
=====
OSPF Virtual Links (detailed)
=====
Neighbor Router Id : 180.0.0.10
-----
Nbr Router Id   : 180.0.0.10      Area Id       : 0.0.0.1
Local Interface: 180.1.7.12      Metric        : 300
State           : Point To Point  Admin State   : Up
Hello Intrvl    : 10 sec          Rtr Dead Intrvl: 60 sec
Tot Rx Packets  : 43022           Tot Tx Packets : 42964
Rx Hellos       : 24834           Tx Hellos     : 24853
```

```

Rx DBDs      : 3          Tx DBDs      : 2
Rx LSRs      : 0          Tx LSRs      : 0
Rx LSUs      : 15966      Tx LSUs      : 16352
Rx LS Acks   : 2219       Tx LS Acks   : 1757
Retransmits  : 0          Discards    : 0
Bad Networks : 0          Bad Versions : 0
Bad Areas    : 0          Bad Dest Adrs : 0
Bad Auth Types : 0        Auth Failures : 0
Bad Neighbors : 0         Bad Pkt Types : 0
Bad Lengths  : 0          Bad Hello Int. : 0
Bad Dead Int. : 0          Bad Options  : 0
Retrans Intrvl : 5 sec     Transit Delay : 1 sec
Last Event    : 11/07/2006 17:11:56 Authentication : None
-----
Neighbor Router Id : 180.0.0.10
-----
Nbr Router Id : 180.0.0.10      Area Id      : 0.0.0.2
Local Interface: 180.2.7.12      Metric      : 300
State          : Point To Point  Admin State  : Up
Hello Intrvl   : 10 sec          Rtr Dead Intrvl: 60 sec
Tot Rx Packets : 43073           Tot Tx Packets : 43034
Rx Hellos      : 24851           Tx Hellos      : 24844
Rx DBDs        : 3              Tx DBDs        : 2
Rx LSRs        : 1              Tx LSRs        : 1
Rx LSUs        : 18071          Tx LSUs        : 17853
Rx LS Acks     : 147            Tx LS Acks     : 334
Retransmits    : 0              Discards      : 0
Bad Networks   : 0              Bad Versions   : 0
Bad Areas      : 0              Bad Dest Adrs  : 0
Bad Auth Types : 0              Auth Failures  : 0
Bad Neighbors  : 0              Bad Pkt Types  : 0
Bad Lengths    : 0              Bad Hello Int. : 0
Bad Dead Int.  : 0              Bad Options    : 0
Retrans Intrvl : 5 sec          Transit Delay   : 1 sec
Last Event     : 11/07/2006 17:12:00 Authentication : MD5
=====
A:ALA-A#

```

Table 40: Output fields: OSPF virtual link

Label	Description
Nbr Rtr ID	Displays the router IDs of neighboring routers
Area Id	Displays a 32-bit integer that identifies an area
Local Interface	Displays the IP address of the local egress interface used to maintain the adjacency to reach this virtual neighbor
Metric	Displays the metric value associated with the route. This value is used when importing this static route into other protocols. When the metric is configured as zero, the metric configured in OSPF, default-import-metric, applies. This value is also used to determine which static route to install in the forwarding table.
State	Displays the operational state of the virtual link to the neighboring router

Label	Description
Authentication	Specifies whether authentication is enabled for the interface or virtual link
Hello Intrval	Displays the length of time, in seconds, between the Hello packets that the router sends on the interface.
Rtr Dead Intrvl	Displays the total number of OSPF packets received where the dead interval specified in the packet was not equal to that configured on this interface since the OSPF admin status was enabled
Tot Rx Packets	Displays the total number of OSPF packets received on this interface since the OSPF admin status was enabled
Rx Hellos	Displays the total number of OSPF Hello packets received on this interface since the OSPF admin status was enabled
Rx DBDs	Displays the total number of OSPF database description packets received on this interface since the OSPF administrative status was enabled
Rx LSRs	Displays the total number of link state requests (LSRs) received on this interface since the OSPF admin status was enabled
Rx LSUs	Displays the total number of link state updates (LSUs) received on this interface since the OSPF admin status was enabled
Rx LS Acks	Displays the total number of link state acknowledgments received on this interface since the OSPF admin status was enabled
Tot Tx Packets	Displays the total number of OSPF packets transmitted on this virtual interface since it was created
Tx Hellos	Displays the total number of OSPF Hello packets transmitted on this virtual interface since it was created
Tx DBDs	Displays the total number of OSPF database description packets transmitted on this virtual interface
Tx LSRs	Displays the total number of OSPF link state requests (LSRs) transmitted on this virtual interface
Tx LSUs	Displays the total number of OSPF Hello packets transmitted on this interface since the OSPF admin status was enabled
Tx LS Acks	Displays the total number of OSPF link state acknowledgments (LSA) transmitted on this virtual interface
Retransmits	Displays the total number of OSPF retransmits sent on this interface since the OSPF admin status was last enabled

Label	Description
Discards	Displays the total number of OSPF packets discarded on this interface since the OSPF admin status was last enabled
Bad Networks	Displays the total number of OSPF packets received with invalid network or mask since the OSPF admin status was last enabled
Bad Versions	Displays the total number of OSPF packets received with bad OSPF version numbers since the OSPF admin status was last enabled
Bad Areas	Displays the total number of OSPF packets received with an area mismatch since the OSPF admin status was last enabled
Bad Dest Addr	Displays the total number of OSPF packets received with the incorrect IP destination address since the OSPF admin status was last enabled
Bad Auth Types	Displays the total number of OSPF packets received with an invalid authorization type since the OSPF admin status was last enabled
Auth Failures	Displays the total number of OSPF packets received with an invalid authorization key since the OSPF admin status was last enabled
Bad Neighbors	Displays the total number of OSPF packets received where the neighbor information does not match the information this router has for the neighbor since the OSPF admin status was last enabled
Bad Pkt Types	Displays the total number of OSPF packets received with an invalid OSPF packet type since the OSPF admin status was last enabled
Bad Lengths	Displays the total number of OSPF packets received on this interface with a total length not equal to the length specified in the packet since the OSPF admin status was last enabled
Bad Hello Int.	Displays the total number of OSPF packets received where the hello interval specified in packet was not equal to that configured on this interface since the OSPF admin status was last enabled
Bad Dead Int.	Displays the total number of OSPF packets received where the dead interval specified in the packet was not equal to that configured on this interface since the OSPF admin status was last enabled
Bad Options	Displays the total number of OSPF packets received with an option that does not match those configured for this interface or area since the OSPF admin status was last enabled

Label	Description
Retrans Intrvl	Displays the length of time, in seconds, that OSPF waits before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor
Transit Delay	Displays the time, in seconds, that it takes to transmit a link state advertisement (LSA) on the interface or virtual link
Last Event	Displays the date and time when an event was last associated with this OSPF interface

virtual-neighbor

Syntax

virtual-neighbor [*remote router-id*] [*detail*]

Context

show>router>ospf

show>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays virtual neighbor information.

Parameters

remote *router-id*

Displays the specified router ID. This reduces the amount of output displayed.

detail

Produces detailed information about the virtual neighbor. This option produces a large amount of data. Nokia recommends to use **detail** only when requesting information for a specific neighbor.

Output

The following output is an example of OSPF virtual neighbor information, and [Table 41: Output fields: OSPF virtual neighbor](#) describes the output fields.

Sample output

```
A:ALA-A# show router ospf virtual-neighbor
=====
OSPF Virtual Neighbors
=====
Nbr IP Addr      Nbr Rtr Id      Nbr State Transit Area  RetxQ Len  Dead Time
-----
```



```

180.1.6.10      180.0.0.10      Full      0.0.0.1      0      58
180.2.9.10      180.0.0.10      Full      0.0.0.2      0      52
-----
No. of Neighbors: 2
=====
A:ALA-A#

A:ALA-A# show router ospf virtual-neighbor detail
=====
OSPF Virtual Neighbors
=====
Virtual Neighbor Router Id : 180.0.0.10
-----
Neighbor IP Addr : 180.1.6.10      Neighbor Rtr Id : 180.0.0.10
Neighbor State   : Full             Transit Area    : 0.0.0.1
Retrans Q Length : 0                Options         : -E--
Events           : 4                Last Event Time : 11/07/2006 17:11:56
Up Time          : 2d 17:47:17      Time Before Dead : 57 sec
Bad Nbr States   : 1                LSA Inst fails  : 0
Bad Seq Nums     : 0                Bad MTUs        : 0
Bad Packets      : 0                LSA not in LSDB : 0
Option Mismatches: 0                Nbr Duplicates  : 0
-----
Virtual Neighbor Router Id : 180.0.0.10
-----
Neighbor IP Addr : 180.2.9.10      Neighbor Rtr Id : 180.0.0.10
Neighbor State   : Full             Transit Area    : 0.0.0.2
Retrans Q Length : 0                Options         : -E--
Events           : 4                Last Event Time : 11/07/2006 17:11:59
Up Time          : 2d 17:47:14      Time Before Dead : 59 sec
Bad Nbr States   : 1                LSA Inst fails  : 0
Bad Seq Nums     : 0                Bad MTUs        : 0
Bad Packets      : 0                LSA not in LSDB : 0
Option Mismatches: 0                Nbr Duplicates  : 0
=====
A:ALA-A#

```

Table 41: Output fields: OSPF virtual neighbor

Label	Description
Nbr IP Addr	Displays the IP address this neighbor is using in its IP source address. Note that, on links with no address, this will not be 0.0.0.0, but the address of another of neighbor interface.
Nbr Rtr ID	Displays the router IDs of neighboring routers
Transit Area	Displays the transit area ID that links the backbone area with the area that has no physical connection with the backbone
Retrans Q Length	Displays the current length of the retransmission queue
No. of Neighbors	Displays the total number of OSPF neighbors adjacent on this interface, in a state of INIT or greater, since the OSPF admin status was enabled
Nbr State	Displays the operational state of the virtual link to the neighboring router

Label	Description
Options	Displays the total number of OSPF packets received with an option that does not match those configured for this virtual interface or transit area since the OSPF admin status was enabled
Events	Displays the total number of events that have occurred since the OSPF admin status was enabled
Last Event Time	Displays the date and time when an event was last associated with this OSPF interface
Up Time	Displays the uninterrupted time, in hundredths of seconds, the adjacency to this neighbor has been up
Time Before Dead	Displays the amount of time, in seconds, until the dead router interval expires
Bad Nbr States	Displays the total number of OSPF packets received where the neighbor information does not match the information this router has for the neighbor since the OSPF admin status was last enabled
LSA Inst fails	Displays the total number of times an LSA could not be installed into the LSDB due to a resource allocation issue since the OSPF admin status was last enabled
Bad Seq Nums	Displays the total number of times when a database description packet was received with a sequence number mismatch since the OSPF admin status was last enabled
Bad MTUs	Displays the total number of times when the MTU in a received database description packet was larger than the MTU of the receiving interface since the OSPF admin status was enabled.
Bad Packets	Displays the total number of times when an LS update was received with an illegal LS type or an option mismatch since the OSPF admin status was enabled
LSA not in LSDB	Displays the total number of times when an LS request was received for an LSA not installed in the LSDB of this router since the OSPF admin status was enabled
Option Mismatches	Displays the total number of times when a LS update was received with an option mismatch since the OSPF admin status was enabled
Nbr Duplicates	Displays the total number of times when a duplicate database description packet was received during the exchange state since the OSPF admin status was enabled

3.12.2.3 Clear commands

ospf

Syntax

ospf

Context

clear>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears and resets OSPF protocol entities.

ospf3

Syntax

ospf3

Context

clear>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears and resets OSPFv3 protocol entities.

database

Syntax

database [purge]

Context

clear>router>ospf

clear>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears all LSAs received from other nodes, sets all adjacencies better than two way to one way, and refreshes all self originated LSAs.

Parameters

purge

Clears all self-originated LSAs and reoriginates all self-originated LSAs.

export

Syntax

export

Context

clear>router>ospf

clear>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command reevaluates all effective export route policies.

neighbor

Syntax

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

Context

clear>router>ospf

clear>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command marks the neighbor as dead and reinitiates the affected adjacencies.

Parameters

ip-int-name

Clears all neighbors for the interface specified by this interface name.

ip-address

Clears all neighbors for the interface specified by this IP address.

statistics

Syntax

statistics

Context

clear>router>ospf

clear>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears all neighbors, routers, interfaces, SPF, and global statistics for OSPF or OSPFv3.

3.12.2.4 Debug commands

ospf

Syntax

ospf *ospf-instance*

Context

debug>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the context for OSPF debugging purposes.

Parameters

ospf-instance

Specifies the OSPF instance.

Values 0 to 31

ospf3

Syntax

ospf3

Context

debug>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the context for OSPFv3 debugging.

area

Syntax

area [*area-id*]

no area

Context

debug>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for an OSPF area.

area-range

Syntax

area-range *ip-address*

no area-range

Context

debug>router>ospf

debug>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for an OSPF or OSPFv3 area range.

Parameters

ip-address

Specifies the IP address for the range used by the ABR to advertise the area into another area.

```
cspf
```

Syntax

```
cspf [ip-address]
```

```
no cspf
```

Context

```
debug>router>ospf
```

```
debug>router>ospf3
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for an OSPF or OSPFv3 constraint-based shortest path first (CSPF).

Parameters

ip-address

Specifies the IP address for the range used for CSPF.

```
graceful-restart
```

Syntax

```
[no] graceful-restart
```

Context

```
debug>router>ospf
```

```
debug>router>ospf3
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for an OSPF or OSPFv3 graceful-restart.

interface

Syntax

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

no interface

Context

debug>router>ospf

debug>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for an OSPF or OSPFv3 interface.

Parameters

ip-int-name

The IP interface name. An interface name cannot be in the form of an IP address.

Interface names can be any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, spaces), the entire string must be enclosed within double quotes.

ip-address

Specifies the interface IP address.

leak

Syntax

leak *ip-address*

no leak

Context

debug>router>ospf

debug>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for OSPF leaks.

Parameters

ip-address

Specifies the IP address to debug OSPF leaks.

lsdb

Syntax

lsdb [**type**] [*ls-id*] [*adv-rtr-id*] [**area** *area-id*]

no lsdb

Context

debug>router>ospf

debug>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for an OSPF or OSPFv3 link-state database (LSDB).

Parameters

type

Specifies the OSPF or OSPFv3 link-state database (LSDB) type.

Values router, network, summary, asbr, extern, nssa, area-opaque, as-opaque, link-opaque

ls-id

Specifies an LSA type specific field containing either a router ID or an IP address. It identifies the piece of the routing domain being described by the advertisement.

adv-rtr-id

Specifies the router identifier of the router advertising the LSA.

area-id

Specifies the OSPF area ID expressed in dotted-decimal notation or as a 32-bit decimal integer.

Values 0 to 4294967295

misc

Syntax

[no] misc

Context

debug>router>ospf

debug>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for miscellaneous OSPF or OSPFv3 events.

neighbor

Syntax

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

no neighbor

Context

debug>router>ospf

debug>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for an OSPF or OSPFv3 neighbor.

Parameters

ip-int-name

Specifies the neighbor interface name.

ip-address

Neighbor information for the neighbor identified by the specified router ID.

nssa-range

Syntax

nssa-range *ip-address*

no nssa-range

Context

debug>router>ospf

debug>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for an NSSA range.

Parameters

ip-address

Specifies the IP address range to debug.

packet

Syntax

packet [*packet-type*] [*interface-name*] [**ingress** | **egress**] [**detail**]

no packet

Context

debug>router>ospf

debug>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for OSPF or OSPFv3 packets.

Parameters

detail

Keyword to specify detailed packet information.

egress

Keyword to specify an egress packet.

ingress

Keyword to specify ingress packet.

interface-name

Specifies the interface name to debug, up to 32 characters.

packet-type

Specifies the OSPF packet type to debug.

Values hello, dbdescr, lsrequest, lsupdate, lsack

rtm

Syntax

rtm [*ip-address*]

no rtm

Context

debug>router>ospf

debug>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for OSPF or OSPFv3 RTM.

Parameters

ip-address

Specifies the IP address to debug.

spf

Syntax

spf [*type*] [*dest-addr*]

no spf

Context

debug>router>ospf

debug>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for OSPF or OSPFv3 SPF. Information regarding overall SPF start and stop times will be shown. To see detailed information regarding the SPF calculation of a specific route, the route must be specified as an optional argument.

Parameters

type

Specifies the area to debug.

Values intra-area, inter-area, external

dest-addr

Specifies the destination IP address to debug.

virtual-neighbor

Syntax

virtual-neighbor [*ip-address*]

no virtual-neighbor

Context

debug>router>ospf

debug>router>ospf3

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for an OSPF or OSPFv3 virtual neighbor.

Parameters

ip-address

Specifies the IP address of the virtual neighbor.

4 IS-IS

This chapter provides information to configure Intermediate System to Intermediate System (IS-IS).

4.1 Configuring IS-IS

Intermediate-system-to-intermediate-system (IS-IS) is a link-state interior gateway protocol (IGP) which uses the Shortest Path First (SPF) algorithm to determine routes. Routing decisions are made using the link-state information. IS-IS evaluates topology changes and, if necessary, performs SPF recalculations.


Entities within IS-IS include networks, intermediate systems, and end systems. In IS-IS, a network is an autonomous system (AS), or routing domain, with end systems and intermediate systems. A router is an intermediate system. End systems are network devices which send and receive protocol data units (PDUs), the OSI term for packets. Intermediate systems send, receive, and forward PDUs.

End system and intermediate system protocols allow routers and nodes to identify each other. IS-IS sends out link-state updates periodically throughout the network, so each router can maintain current network topology information.

IS-IS supports large ASs by using a two-level hierarchy. A large AS can be administratively divided into smaller, more manageable areas. A system logically belongs to one area. Level 1 routing is performed within an area. Level 2 routing is performed between areas. Routers can be configured as Level 1, Level 2, or both Level 1/2.

The following figure shows an example of an IS-IS routing domain.

Figure 11: IS-IS routing domain



OSRG033

4.1.1 Routing

OSI IS-IS routing uses two-level hierarchical routing. A routing domain can be partitioned into areas. Level 1 routers know the topology in their area, including all routers and end systems in their area but do not know the identity of routers or destinations outside of their area. Level 1 routers forward traffic with destinations outside of their area to a Level 2 router in their area.

Level 2 routers know the Level 2 topology, and know which addresses are reachable by each Level 2 router. Level 2 routers do not need to know the topology within any Level 1 area, except to the extent that a Level 2 router can also be a Level 1 router within a single area. By default, only Level 2 routers can exchange PDUs or routing information directly with external routers located outside the routing domain.

In IS-IS, there are the following types of routers:

- **Level 1 intermediate systems**

Routing is performed based on the area ID portion of the ISO address called the Network Entity Title (NET). Level 1 systems route within an area. They recognize, based on the destination address, whether the destination is within the area. If so, they route toward the destination. If not, they route to the nearest Level 2 router.

- **Level 2 intermediate systems**

Routing is performed based on the area address. They route toward other areas, disregarding other area's internal structure. A Level 2 intermediate system can also be configured as a Level 1 intermediate system in the same area.

The Level 1 router's area address portion is manually configured (see [ISO network addressing](#)). A Level 1 router will not become a neighbor with a node that does not have a common area address. However,

if a Level 1 router has area addresses A, B, and C, and a neighbor has area addresses B and D, then the Level 1 router will accept the other node as a neighbor, as address B is common to both routers. Level 2 adjacencies are formed with other Level 2 nodes whose area addresses do not overlap. If the area addresses do not overlap, the link is considered by both routers to be Level 2 only and only Level 2 LSPDs flow on the link.

Within an area, Level 1 routers exchange LSPs which identify the IP addresses reachable by each router. Specifically, zero or more IP address, subnet mask, and metric combinations can be included in each LSP. Each Level 1 router is manually configured with the IP address, subnet mask, and metric combinations, which are reachable on each interface. A Level 1 router routes as follows:

- If a specified destination address matches an IP address, subnet mask, or metric reachable within the area, the PDU is routed via Level 1 routing.
- If a specified destination address does not match any IP address, subnet mask, or metric combinations listed as reachable within the area, the PDU is routed toward the nearest Level 2 router.

Level 2 routers include in their LSPs, a complete list of IP address, subnet mask, and metrics specifying all the IP addresses which reachable in their area. This information can be obtained from a combination of the Level 1 LSPs (by Level 1 routers in the same area). Level 2 routers can also report external reachable information, corresponding to addresses reachable by routers in other routing domains or autonomous systems.

4.1.2 IS-IS frequently used terms

The following are frequently used terms for IS-IS:

- **area**
An area is a routing sub-domain which maintains detailed routing information about its own internal composition, and also maintains routing information which allows it to reach other routing sub-domains. Areas correspond to the Level 1 sub-domain.
- **end system**
End systems send NPDUs to other systems and receive NPDUs from other systems, but do not relay NPDUs. This International Standard does not specify any additional end system functions beyond those supplied by ISO 8473 and ISO 9542.
- **neighbor**
A neighbor is an adjacent system reachable by traversing a single sub-network by a PDU.
- **adjacency**
An adjacency is a portion of the local routing information which pertains to the reachability of a single neighboring end or intermediate system over a single circuit. Adjacencies are used as input to the decision process to form paths through the routing domain. A separate adjacency is created for each neighbor on a circuit and for each level of routing (Level 1 and Level 2) on a broadcast circuit.
- **circuit**
A circuit is the subset of the local routing information base pertinent to a single local Subnetwork Point of Attachments (SNPAs).
- **link**
A link is the communication path between two neighbors. A link is up when communication is possible between the two SNPAs.

- **designated IS**

A designated IS is the intermediate system on a LAN which is designated to perform additional duties. In particular, the designated IS generates link-state PDUs on behalf of the LAN, treating the LAN as a pseudonode.

- **pseudonode**

Where a broadcast sub-network has n connected intermediate systems, the broadcast sub-network is considered to be a pseudonode. The pseudonode has links to each of the n intermediate systems and each of the ISs has a single link to the pseudonode (rather than $n-1$ links to each of the other intermediate systems). Link-state PDUs are generated on behalf of the pseudonode by the designated IS.

- **broadcast sub-network**

A broadcast sub-network is a multi-access subnetwork that supports the capability of addressing a group of attached systems with a single PDU.

- **general topology sub-network**

A general topology sub-network is a topology that is modeled as a set of point-to-point links, each of which connects two systems. There are several generic types of general topology subnetworks, multipoint links, permanent point-to-point links, dynamic and static point-to-point links.

- **routing sub-domain**

A routing sub-domain consists of a set of intermediate systems and end systems located within the same routing domain.

- **level 2 sub-domain**

A level 2 sub-domain is the set of all Level 2 intermediate systems in a routing domain.

4.1.3 ISO network addressing

IS-IS uses ISO network addresses. Each address identifies a point of connection to the network, such as a router interface, and is called a Network Service Access Point (NSAP).

An end system can have multiple NSAP addresses, in which case the addresses differ only by the last byte (called the *n-selector*). Each NSAP represents a service that is available at that node. In addition to having multiple services, a single node can belong to multiple areas.

Each network entity has a special network address called a Network Entity Title (NET). Structurally, an NET is identical to an NSAP address but has an n-selector of 00. Most end systems have one NET. Intermediate systems can have up to three area IDs (area addresses).

NSAP addresses are divided into three parts. Only the area ID portion is configurable:

- **area ID**

An area ID is a variable length field between 1 and 13 bytes. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.

- **system ID**

A system ID is a six-byte system identification. This value is not configurable. The system ID is derived from the system or router ID.


- **selector ID**

A selector ID is a one-byte selector identification that must contain zeros when configuring a NET. This value is not configurable. The selector ID is always 00.

Of the total 20 bytes comprising the NET, only the first 13 bytes, the area ID portion, can be manually configured. As few as one byte can be entered or, at most, 13 bytes. If less than 13 bytes are entered, the rest is padded with zeros.

Routers with common area addresses form Level 1 adjacencies. Routers with no common NET addresses form Level 2 adjacencies, if they are capable (see the following figure).

Figure 12: Using area addresses to form adjacencies



4.1.4 IS-IS PDU configuration

The following PDUs are used by IS-IS to exchange protocol information:

- **IS-IS hello PDU**

Routers with IS-IS enabled send hello PDUs to IS-IS-enabled interfaces to discover neighbors and establish adjacencies.

- **Link-state PDUs**

Contain information about the state of adjacencies to neighboring IS-IS systems. LSPs are flooded periodically throughout an area.

- **Complete sequence number PDUs**

In order for all routers to maintain the same information, CSNPs inform other routers that some LSPs can be outdated or missing from their database. CSNPs contain a complete list of all LSPs in the current IS-IS database.

- **Partial sequence number PDUs (PSNPs)**

PSNPs are used to request missing LSPs and acknowledge that an LSP was received.

4.1.5 IS-IS operations

Routers perform IS-IS routing as follows:

- Hello PDUs are sent to the IS-IS-enabled interfaces to discover neighbors and establish adjacencies.
- IS-IS neighbor relationships are formed if the hello PDUs contain information that meets the criteria for forming an adjacency.
- Routers can build a link-state PDU based upon their local interfaces that are configured for IS-IS and prefixes learned from other adjacent routers.
- Routers flood LSPs to the adjacent neighbors except the neighbor from which they received the same LSP. The link-state database is constructed from these LSPs.
- A Shortest Path Tree (SPT) is calculated by each IS, and from this SPT the routing table is built.

4.1.6 IS-IS route summarization

IS-IS IPv4 route summarization allows users to create aggregate IPv4 addresses that include multiple groups of IPv4 addresses for a specific IS-IS level. IPv4 Routes redistributed from other routing protocols also can be summarized. It is similar to the OSPF area-range command. IS-IS IPv4 route summarization helps to reduce the size of the LSDB and the IPv4 routing table, and it also helps to reduce the chance of route flapping.

IPv4 route summarization supports:

- Level 1, Level 1-2, and Level 2
- route summarization for the IPv4 routes redistributed from other protocols
- metric used to advertise the summary address will be the smallest metric of all the more specific IPv4 routes

4.1.7 IS-IS multi-topology for IPv6

IS-IS IPv6 Type-Length-Value (TLV) encoding for IPv6 routing is supported in 7210 SAS. This type of routing is considered native IPv6 routing with IS-IS. It has a limitation that IPv4 and IPv6 topologies must be congruent, otherwise traffic may be blackholed. Service providers should ensure that the IPv4 topology and IPv6 topology are the same. With IS-IS multi-topology, service providers can use different topologies for IPv4 and IPv6.

The implementation is compliant with *draft-ietf-isis-wg-multi-topology-xx.txt*, *M-ISIS: Multi Topology (MT) Routing in IS-IS*.

The following MT topologies are supported:

- MT ID #0 - equivalent to the standard IS-IS topology
- MT ID #2 - reserved for IPv6 routing topology

4.1.8 IS-IS administrative tags

IS-IS admin tags enable a network administrator to configure route tags to tag IS-IS route prefixes. These tags can subsequently be used to control Intermediate System-to-Intermediate System (IS-IS) route redistribution or route leaking.

The IS-IS support for route tags allows the tagging of IP addresses of an interface and use the tag to apply administrative policy with a route map. A network administrator can also tag a summary route and then use a route policy to match the tag and set one or more attributes for the route.

Using these administrative policies allow the operator to control how a router handles the routes it receives from and sends to its IS-IS neighboring routers. Administrative policies are also used to govern the installation of routes in the routing table.

Route tags allow:

- policies to redistribute routes received from other protocols in the routing table to IS-IS
- policies to redistribute routes between levels in an IS-IS routing hierarchy
- policies to summarize routes redistributed into IS-IS or within IS-IS by creating aggregate (summary) addresses

4.1.8.1 Setting route tags

IS-IS route tags are configurable in the following ways:

- setting a route tag for an IS-IS interface
- setting a route tag on an IS-IS passive interface
- setting a route tag for a route redistributed from another protocol to IS-IS
- setting a route tag for a route redistributed from one IS-IS level to another IS-IS level
- setting a route tag for an IS-IS default route
- setting a route tag for an IS-IS summary address

4.1.8.2 Using route tags

The IS-IS administrative tags configured on an IS-IS router (or neighbor) will not have any effect until policies are configured to process the specified tag value.

Policies can process route tags that specify ISIS as either the origin or destination protocol, or as both origin and destination protocol.

```
config>router>policy-options>policy-statement>entry>from  
config>router>policy-options>policy-statement>entry>action tag tag-value  
config>router>policy-options>policy-statement# default-action tag tag-value
```

4.1.9 Segment routing in shortest path forwarding

Segment routing (SR) adds to IS-IS and OSPF routing protocols the ability to perform shortest path routing and source routing using the concept of abstract segment. A segment can represent a local prefix of a

node, a specific adjacency of the node (interface or next-hop), a service context, or a specific path over the network. For each segment, the IGP advertises an identifier referred to as a segment ID (SID).

When SR is used in combination with the MPLS data plane, the SID acts as a standard MPLS label. A router forwarding a packet using SR therefore pushes one or more MPLS labels. This section describes the SR MPLS feature.

Both shortest path routing and traffic engineering applications can leverage SR MPLS, which encodes a segment as an MPLS label. This section describes the shortest path forwarding applications.

When a received IPv4 prefix SID is resolved, the SR module programs the ILM with a swap operation and the LTN with a push operation, both pointing to the primary/LFA NHLFE. An IPv4 SR tunnel to the prefix destination is also added to the TTM and is available for use by L2 and L3 services.

The SR tunnel in the TTM is available in the following contexts:

- IPv4 BGP route label
- VLL, LDP VPLS, and RVPLS
- BGP-AD VPLS when the **use-provisioned-sdp** option is enabled in the binding to the PW template
- intra-AS BGP VPRN for VPN-IPv4 and VPN-IPv6 prefixes, both auto-bind and explicit SDP

The remote LFA feature included in SR expands the coverage of the LFA by computing and automatically programming the SR tunnels that are used as backup next-hops. The SR shortcut tunnels terminate on a remote alternate node, which provides loop-free forwarding for packets of the resolved prefixes. When the **loopfree-alternate** option is enabled in an IS-IS or OSPF instance, SR tunnels are protected with an LFA backup next-hop. If the prefix of a specific SR tunnel is not protected by the base LFA, the remote LFA automatically computes a backup next-hop using an SR tunnel if the **remote-lfa** option is also enabled in the IGP instance.



Note:

On the 7210 SAS-K, the maximum label push depth is four MPLS labels and the maximum label pop depth is four MPLS labels (both push and pop exclude the pseudowire hash label).

4.1.9.1 Segment routing operational procedures

4.1.9.1.1 Prefix advertisement and resolution

When segment routing is enabled in the IS-IS or OSPF instance, the router performs the following operations. See [Control protocol changes](#) for detailed information about the TLVs and sub-TLVs for both IS-IS and OSPF protocols.

1. Advertises the segment routing capability sub-TLV to routers in all areas or levels of this IGP instance. However, only neighbors with which it established an adjacency interprets the SID or label range information and use it for calculating the label to swap to or push for a resolved prefix SID.
2. Advertises the assigned index for each configured node SID in the new prefix SID sub-TLV with the N-flag (node-SID flag) set. The segment routing module programs the incoming label map (ILM) with a pop operation for each local node SID in the data path.
3. Automatically assigns and advertises an adjacency SID label for each formed adjacency over a network IP interface in the new adjacency SID sub-TLV. The following points should be considered:
 - The adjacency SID is advertised for both numbered and unnumbered network IP interfaces.

- The adjacency SID for parallel adjacencies between two IGP neighbors is not supported.
 - The adjacency SID is not advertised for an IES interface because access interfaces do not support MPLS.
 - The adjacency SID must be unique for each instance and for each adjacency. Also, ISIS MT=0 can establish an adjacency for the IPv4 address family over the same link, and in such a case a different adjacency SID is assigned to each next-hop. However, the existing IS-IS implementation assigns a single protect-group ID (PG-ID) to the adjacency, and when the state machine of a BFD session tracking the IPv4 next-hop times out, an action is triggered for the prefixes of the IPv4 address family over that adjacency.
 - The segment routing module programs the ILM with a swap to an implicit null label operation for each advertised adjacency SID.
4. Resolve received prefixes and, if a prefix SID sub-TLV exists, the segment routing module programs the ILM with a swap operation and an LTN with a push operation, both pointing to the primary/LFA NHLFE. An SR tunnel is also added to the TTM. If a node SID resolves over an IES interface, the data path is not programmed and a trap is raised. Therefore, only next-hops of an ECMP set corresponding to network IP interfaces are programmed in the data path; next-hops corresponding to IES interfaces are not programmed. If, however, the user configures the interface as a network on one side and IES on the other side, MPLS packets for the SR tunnel received on the access side are dropped.



Note:

LSA filtering causes SIDs not to be sent in one direction, which means that some node SIDs are resolved in parts of the network upstream of the advertisement suppression.

When the user enables segment routing in an IGP instance, the main SPF and LFA SPF are computed normally and the primary next-hop and LFA backup next-hop for a received prefix are added to RTM without the label information advertised in the prefix SID sub-TLV. In all cases, the segment routing tunnel is not added into the RTM.

4.1.9.1.2 Error and resource exhaustion handling

When the prefix corresponding to a node SID is being resolved, the following procedures are followed.

4.1.9.1.2.1 Procedure 1: Resolving received SID indexes or labels to different routes of the same prefix within the same IGP instance

Two variations of this procedure can occur:

1. When the 7210 SAS does not allow assigning the same SID index or label to different routes of the same prefix within the same IGP instance, it resolves only one of them if they are received from another SR implementation and they are based on the RTM active route selection.
2. When the 7210 SAS does not allow assigning different SID indexes or labels to different routes of the same prefix within the same IGP instance, it resolves only one of them if they are received from another SR implementation and they are based on the RTM active route selection.

The selected SID is used for ECMP resolution to all neighbors. If the route is inter-area and the conflicting SIDs are advertised by different ABRs, ECMP towards all ABRs uses the selected SID.

4.1.9.1.2.2 Procedure 2: Checking for SID error prior to programming ILM and NHLFE

If any of the following conditions are true, the router logs a trap and generates a syslog error message, and it does not program the ILM and NHLFE for the prefix SID:

- The received prefix SID index falls outside the locally configured SID range.
- One or more resolved ECMP next-hops for a received prefix SID did not advertise SR capability sub-TLV.
- The received prefix SID index falls outside the advertised SID range of one or more resolved ECMP next-hops.

4.1.9.1.2.3 Procedure 3: Programming ILM/NHLFE for duplicate prefix-SID indexes/labels for different prefixes

Two variations of this procedure can occur:

1. For received duplicate prefix-SID indexes or labels for different prefixes within the same IGP instance, the router does the following:
 - programs ILM/NHLFE for the first prefix
 - logs a trap and a syslog error message
 - does not program the subsequent prefix in the data path
2. For received duplicate prefix-SID indexes for different prefixes across IGP instances, there are two options:
 - In the global SID index range mode of operation, the resulting ILM label values are the same across the IGP instances. The router does the following:
 - programs ILM/NHLFE for the prefix of the winning IGP instance based on the RTM route type preference
 - logs a trap and a syslog error message
 - does not program the subsequent prefix SIDs in the data path
 - In the per-instance SID index range mode of operation, the resulting ILM label will have different values across the IGP instances. The router programs ILM/NHLFE for each prefix as expected.


4.1.9.1.2.4 Procedure 4: Programming ILM/NHLFE for the same prefix across IGP instances

In the global SID index range mode of operation, the resulting ILM label values are the same across the IGP instances. The router programs ILM/NHLFE for the prefix of the winning IGP instance based on the RTM route type preference. The router logs a trap and a syslog error message, and does not program the other prefix SIDs in data path.

In the per-instance SID index range mode of operation, the resulting ILM label has different values across the IGP instances. The router programs ILM/NHLFE for each prefix as expected.

The following figure shows an IS-IS example of the behavior in the case of a global SID index range.

Figure 13: Handling of the same prefix and SID in different IS-IS instances



Assume that the following route type preference in the RTM and tunnel type preference in the TTM are configured:

- ROUTE_PREF_ISIS_L1_INTER (RTM) 15
- ROUTE_PREF_ISIS_L2_INTER (RTM) 18
- ROUTE_PREF_ISIS_TTM 11



Note:

The TTM tunnel type preference is not used by the SR module. It is put in the TTM and is used by other applications, such as a VPRN auto-bind, to select a TTM tunnel.

1. Router A performs the following resolution within the single IS-IS instance 1, level 2. All metrics are the same, and ECMP = 2:
 - For prefix N, the RTM entry is the following:
 - prefix N
 - nhop1 = B
 - nhop2 = C
 - preference 18
 - For prefix N, the SR tunnel TTM entry is the following:
 - tunnel-id 1: prefix N-SIDx
 - nhop1 = B

- nhop2 = C
 - tunl-pref 11 (tunl-pref 10 for OSPF)
2. Add IS-IS instance 2 (level 1) in the same setup, but in routers A, B, and E only.
- For prefix N, the RTM entry is the following:
 - prefix N
 - nhop1 = B
 - preference 15RTM prefers the level 1 route over the level 2 route.
 - For prefix N, there is one SR tunnel entry for level 2 in the TTM:
 - tunnel-id 1: prefix N-SIDx
 - nhop1 = B
 - nhop2 = C
 - tunl-pref 11 (tunl-pref 10 for OSPF)

4.1.9.1.2.5 Procedure 5: Handling ILM resource exhaustion while assigning an SID index/label

If the system exhausts an ILM resource while assigning an SID index or label to a local loopback interface, index allocation fails and an error is returned in the CLI. In addition, the router logs a trap and generates a syslog error message.

4.1.9.1.2.6 Procedure 6: Handling ILM/NHLFE/other IOM or CPM resource exhaustion while resolving or programming an SID index/label

If the system exhausts an ILM, NHLFE, or any other IOM or CPM resource while resolving and programming a received prefix SID or programming a local adjacency SID, the following occurs:

- The IGP instance goes into overload and a trap and syslog error message are generated.
- The segment routing module deletes the tunnel.

The user must manually clear the IGP overload condition after freeing resources. After the IGP is brought back up, it attempts to program at the next SPF all tunnels which previously failed the programming operation.

4.1.9.2 Segment routing tunnel management

The segment routing module adds to the TTM a shortest path SR tunnel entry for each resolved remote node SID prefix and programs the data path with the corresponding LTN with the push operation pointing to the primary and LFA backup NHLFEs. The LFA backup next-hop for a prefix that was advertised with a node SID will be computed only if the **loopfree-alternate** option is enabled in the IS-IS or OSPF instance. The resulting SR tunnel that is populated in the TTM is automatically protected with FRR when an LFA backup next-hop exists for the prefix of the node SID.

With ECMP, a maximum number of primary next-hops (NHLFEs) are programmed for the same tunnel destination per IGP instance. ECMP and LFA next-hops are mutually exclusive as per the existing implementation.

The default preference for shortest path SR tunnels in the TTM is set lower than LDP tunnels but higher than BGP tunnels to allow controlled migration of customers without disrupting their current deployment when they enable segment routing. The following list presents the default preference of the various tunnel types. This includes the preference of both SR tunnels based on shortest path (referred to as SR-ISIS and SR-OSPF).

The global default TTM preference for the tunnel types is as follows:

- ROUTE_PREF_RSVP 7
- ROUTE_PREF_SR_TE 8
- ROUTE_PREF_LDP 9
- ROUTE_PREF_OSPF_TTM 10
- ROUTE_PREF_ISIS_TTM 11
- ROUTE_PREF_BGP_TTM 12
- ROUTE_PREF_GRE 255

The default value for SR-ISIS or SR-OSPF is the same regardless of whether one or more IS-IS or OSPF instances are programming a tunnel for the same prefix. The selection of an SR tunnel in this case is based on the lowest IGP instance.

The TTM preference is used in the case of VPRN auto-bind or BGP transport tunnels when the tunnel binding commands are configured to the **any** value, which parses the TTM for tunnels in the protocol preference order. The user can choose to either use the global TTM preference or list explicitly the tunnel types to be used. When the tunnel types are listed, the TTM preference is still used to select one type over the other. In both cases, a fallback to the next preferred tunnel type is performed if the selected one fails. Also, a reversion to a preferred tunnel type is performed as soon as one is available. See [BGP label route resolution using segment routing tunnels](#) and [Service packet forwarding with segment routing](#) for the detailed service and shortcut binding CLI.

For SR-ISIS and SR-OSPF, the user can configure the preference of each IGP instance in addition to the preceding default values using the following CLI syntax.

```
configure>router>isis>segment-routing>tunnel-table-pref preference <1..255>  
configure>router>ospf>segment-routing>tunnel-table-pref preference <1..255>
```

SR tunnels in the TTM are available to BGP routes, VPRN auto-bind and explicit SDP binding, and L2 services with PW template auto-bind and explicit SDP binding.

Local adjacency SIDs are not programmed into the TTM, but the remote adjacency SIDs can be used together with a node SID in a tunnel configuration in a directed LFA.

4.1.9.2.1 Tunnel MTU determination

The MTU of an SR tunnel populated into the TTM is determined the same way as it is for an IGP tunnel; for example, LDP LSP, based on the outgoing interface MTU minus the label stack size. Segment routing, however, supports remote LFA, which programs an LFA backup next-hop adding another label to the tunnel for a total of two labels.

The user must configure the MTU of all SR tunnels within each IGP instance using the following CLI syntax:

```
configure>router>isis>segment-routing>tunnel-mtu bytes
configure>router>ospf>segment-routing>tunnel-mtu bytes
```

There is no default value for this new command. If the user does not configure an SR tunnel MTU, the MTU is fully determined by IGP.

The MTU of the SR tunnel, in bytes, is determined as follows:

$$SR_Tunnel_MTU = MIN \{ Cfg_SR_MTU, IGP_Tunnel_MTU - (1 + frr-overhead) * 4 \}$$

Where:

- *Cfg_SR_MTU* is the MTU configured by the user for all SR tunnels within a specific IGP instance using the preceding commands. If no value was configured by the user, the SR tunnel MTU is determined by the IGP interface calculation.
- *IGP_Tunnel_MTU* is the minimum of the IS-IS or OSPF interface MTU among all the ECMP paths or among the primary and LFA backup paths of this SR tunnel.
- *frr-overhead* is set to 1 if **segment-routing** and **remote-lfa** options are enabled in the IGP instance. Otherwise, it is set to 0.

The SR tunnel MTU is dynamically updated whenever any of the preceding parameters changes. This includes when the set of the tunnel next-hops changes, or the user changes the configured SR MTU or interface MTU value.



Note:

For the purpose of fragmentation of IP packets forwarded in GRT or in a VPRN over an SR shortest path tunnel, the IOM always deducts the worst case MTU (5 labels or 6 labels if the hash label feature is enabled) from the outgoing interface MTU for the decision to fragment the packet or not. In this case, the preceding formula is not used.

4.1.9.3 Remote LFA with segment routing

The remote LFA next-hop calculation by the IGP LFA SPF is enabled by appending the **remote-lfa** option to the **loopfree-alternate** command:

```
configure>router>isis>loopfree-alternate remote-lfa
configure>router>ospf>loopfree-alternate remote-lfa
```

The SPF calculates the remote LFA after the regular LFA next-hop calculation when the following conditions are met:

- The **remote-lfa** option is enabled in an IGP instance.
- The LFA next-hop calculation did not result in protection for one or more prefixes resolved to a specific interface.

Remote LFA extends the loop-free alternate fast reroute (LFA FRR) protection coverage to any topology by automatically computing and establishing or tearing down shortcut tunnels (repair tunnels) to a remote LFA node that puts the packets back on the shortest path without looping them back to the node that forwarded them over the repair tunnel. A repair tunnel can be an RSVP LSP, an LDP-in-LDP tunnel, or an SR tunnel. This feature is restricted to using an SR repair tunnel to the remote LFA node.




Note:

The remote LFA feature can only use an SR repair tunnel to the remote LFA node.

The remote LFA algorithm for link protection is described in RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*. Unlike the regular LFA calculation, which is calculated per prefix, the LFA algorithm for link protection is a per-link LFA SPF calculation. It provides protection for all destination prefixes that share the protected link by using the neighbor on the other side of the protected link as a proxy for all destinations. The following figure shows an example of a remote LFA topology.

Figure 14: Example topology remote LFA algorithm



When the LFA SPF in node C computes the per-prefix LFA next-hop, prefixes that use link C to B as the primary next-hop have no LFA next-hop because of the ring topology. If node C uses node link C to D as a back-up next-hop, node D loops a packet back to node C. The remote LFA then runs the "PQ Algorithm" as described in RFC 7490.

1. Computes the extended P space of node C for link C to B. The extended P space is the set of nodes reachable from node C without any path transiting the protected link (C to B). The computation yields nodes D, E, and F.

The extended P space of node C is determined by running SPF on behalf of each of the neighbors of C; the same computation is used for the regular LFA.



Note:

According to the P space concept initially introduced in RFC 7490, node F would be excluded from the P space because, from the node C perspective, a few node C has a couple of ECMP paths would already exist in node C, including a path going through link C to B. However, because the remote LFA next-hop is activated when link C-B fails, this rule can be relaxed to include node F, which then yields the extended P space.

You can limit the search for candidate P nodes to reduce the number of SPF calculations in topologies where many eligible P nodes may exist. Use the following CLI commands to configure the maximum IGP cost from node C for a P node to be an eligible candidate:

- **configure>router>isis>loopfree-alternate remote-lfa max-pq-cost *value***
- **configure>router>ospf>loopfree-alternate remote-lfa max-pq-cost *value***

2. Compute the Q space of node B for link C-B. The Q space is the set of nodes from which the destination proxy (node B) can be reached without a path transiting the protected link (link C-B).

The Q space calculation is a reverse SPF on node B. A reverse SPF is run on behalf of each neighbor of C to protect all destinations that resolve over the link to the neighbor. This yields nodes F and A in the example shown in [Figure 14: Example topology remote LFA algorithm](#).

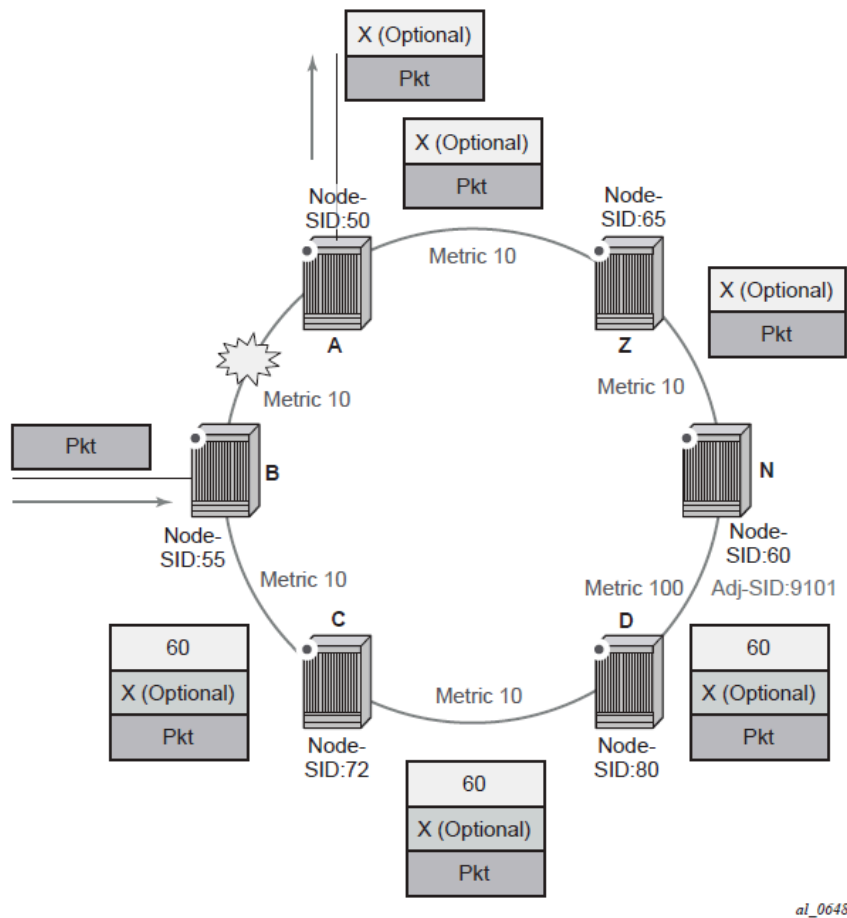
You can limit the search for candidate Q nodes to reduce the number of SPF calculations in topologies where many eligible Q nodes may exist. Use the CLI commands described in step 1 to configure the maximum IGP cost from node C for a Q node to be an eligible candidate.

3. Select the best alternate node, which is the intersection of extended P and Q spaces. In the [Figure 14: Example topology remote LFA algorithm](#) example, the best alternate node (PQ node) is node F. From node F onwards, traffic follows the IGP shortest path.

If many PQ nodes exist, the lowest IGP cost from node C is used to narrow the selection; if more than one PQ node remains, the node with the lowest router ID is selected.

The following figure shows label stack encoding for a packet that is forwarded over the remote LFA next-hop.

Figure 15: Remote LFA next-hop in segment routing



The label corresponding to the node SID of the PQ node is pushed on top of the original label of the SID of the resolved destination prefix. If node C has resolved multiple node SIDs corresponding to different prefixes of the selected PQ node, it pushes the lowest node SID label on the packet when forwarded over the remote LFA backup next-hop.

If the PQ node is also the advertising router for the resolved prefix, the label stack is compressed in the following cases depending on the IGP:

- In IS-IS, the label stack is always reduced to a single label, which is the label of the resolved prefix owned by the PQ node.
- In OSPF, the label stack is reduced to the single label of the resolved prefix when the PQ node advertises a single node SID in this OSPF instance. If the PQ node advertises a node SID for multiple of its loopback interfaces within the same OSPF instance, the label stack is reduced to a single label only in the case where the SID of the resolved prefix is the lowest SID value.

The following rules and limitations apply to the remote LFA implementation:

- LFA policy is supported for IP next-hops only. It is not supported with tunnel next-hops when IGP shortcuts are used for LFA backup. Remote LFA is also a tunnel next-hop and a user-configured LFA policy is not applied in the selection of a remote LFA backup next-hop when multiple candidates are available.

- As a result, if an LFA policy is applied and does not find an LFA IP next-hop for a set of prefixes, the remote LFA SPF searches for a remote LFA next-hop for the same prefixes. The selected remote LFA next-hops, if found, may not satisfy the LFA policy constraints.
- If the **loopfree-alternate-exclude** command (IS-IS or OSPF context of the interface) is used to exclude a network IP interface from being used as an LFA next-hop, the interface is also excluded from being used as the outgoing interface for a remote LFA tunnel next-hop.
- As with the regular LFA algorithm, the remote LFA algorithm computes a backup next-hop to the ABR advertising an inter-area prefix and not to the destination prefix.

4.1.9.4 Data path support

A packet received with a label matching either a node SID or an adjacency SID is forwarded according to the ILM type and operation, as described in the following table.

Table 42: Data path support

Label type	Operation
Top label is a local node SID	The label is popped and the packet is further processed. If the SID label of the popped node is at the bottom of the stack label, the IP packet is looked up and forwarded in the appropriate FIB.
Top or next label is a remote node SID	The label is swapped to the calculated label value for the next-hop and forwarded according to the primary or backup NHLFE. With ECMP, a number of primary next-hops (NHLFEs) are programmed for the same destination prefix and for each IGP instance. ECMP and LFA next-hops are mutually exclusive.
Top or next label is an adjacency SID	The label is popped and the packet is forwarded out of the interface to the next-hop associated with this adjacency SID label. In effect, the data path operation is modeled like a swap to an implicit-null label instead of a pop.
Next label is BGP 3107 label	The packet is further processed according to the ILM operation. The BGP label may be popped and the packet looked up in the appropriate FIB.
Next label is a service label	The packet is looked up and forwarded in the Layer 2 or VPRN FIB.

A router forwarding an IP or service packet over an SR tunnel pushes a maximum of three transport labels with a remote LFA next-hop.

4.1.9.4.1 Hash label

When the **hash-label** option is enabled in a service context, the hash label is always inserted at the bottom of the stack.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, hash label is supported only with specific services. See the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Services Guide for more information about services supported with hash label.

4.1.9.5 Control protocol changes

This section describes the [IS-IS control protocol changes](#) and [OSPF control protocol changes](#).

4.1.9.5.1 IS-IS control protocol changes

The following TLVs/sub-TLVs are defined in *draft-ietf-isis-segment-routing-extensions* and are supported in the implementation of SR in IS-IS:

- prefix SID sub-TLV
- adjacency SID sub-TLV
- SID/Label Binding TLV
- SR-Capabilities sub-TLV
- SR-Algorithm sub-TLV

This section describes the behaviors and limitations of using SR TLVs and sub-TLVs with IS-IS.

The 7210 SAS supports advertising the IS router capability TLV (RFC 4971) only for topology MT=0. As a result, the SR-Capabilities sub-TLV can be advertised only in MT=0, which restricts the segment routing feature to MT=0.

Similarly, if prefix SID sub-TLVs for the same prefix are received in different MT numbers of the same IS-IS instance, only the one in MT=0 is resolved. When the prefix SID index is also duplicated, an error is logged and a trap is generated, as described in [Error and resource exhaustion handling](#).

The I and V flags are both set to 1 when originating the SR-Capabilities sub-TLV to indicate support for processing SR MPLS encapsulated IPv4 and IPv6 packets on its network interfaces. These flags are not checked when the sub-TLV is received. Only the SRGB range is processed.

The algorithm field is set to 0, meaning it uses the SPF algorithm based on link metric, when the SR-Algorithm sub-TLV is originated but the field is not checked when the sub-TLV is received.

Only an IPv4 prefix and adjacency SID sub-TLVs can be originated within MT=0. An IPv6 prefix and adjacency SID sub-TLVs can, however, be received and ignored. Use the **show** command to display (dump) the octets of the received but unsupported sub-TLVs.

The 7210 SAS originates a single prefix SID sub-TLV per IS-IS IP reachability TLV and processes the first prefix SID sub-TLV only if multiple prefix SID sub-TLVs are received within the same IS-IS IP reachability TLV.

The 7210 SAS encodes the 32-bit index in the prefix SID sub-TLV. The 24-bit label is not supported.

The 7210 SAS originates a prefix SID sub-TLV with the following flag encoding and processing rules:

- The R-flag is set if the prefix SID sub-TLV, along with its corresponding IP reachability TLV, is propagated between levels.
- The N-flag is always set because the system supports a prefix SID of type node SID only.
- The P-flag (no-PHP flag) is always set, meaning that the label for the prefix SID is pushed by the penultimate hop popping (PHP) router when forwarding to this router. The 7210 SAS PHP router processes a received prefix SID with the P-flag set to zero and uses implicit-null for the outgoing label towards the router that advertised it, as long as the P-flag is also set to 1.
- The E-flag (Explicit-Null flag) is always set to zero. The PHP router, however, processes a received prefix SID with the E-flag set to 1 and, when the P-flag is also set to 1, it pushes explicit-null for the outgoing label toward the router that advertised it.
- The V-flag is always set to 0 to indicate an index value for the SID.
- The L-flag is always set to 0 to indicate that the SID index value is not locally significant.
- The algorithm field is always set to zero to indicate that the SPF algorithm is based on the link metric and is not checked on a received prefix SID sub-TLV.
- The system resolves a prefix SID sub-TLV received without the N-flag set but with the prefix length equal to 32. A trap, however, is raised by IS-IS.
- The system does not resolve a prefix SID sub-TLV received with the N flag set and a prefix length different than 32. A trap is raised by IS-IS.
- The system resolves a prefix SID received within an IP reachability TLV based on the following route preference:
 - SID received via level 1 in a prefix SID sub-TLV part of IP reachability TLV
 - SID received via level 2 in a prefix SID sub-TLV part of IP reachability TLV
- A prefix received in an IP reachability TLV is propagated, along with the prefix SID sub-TLV, by default from level 1 to level 2 by a level 1/2 router. A router in level 2 sets up an SR tunnel to the level 1 router via the level 1/2 router, which acts as an LSR.
- A prefix received in an IP reachability TLV is not propagated, along with the prefix SID sub-TLV, by default from level 2 to level 1 by a level 1/2 router. If the user adds a policy to propagate the received prefix, a router in level 1 sets up an SR tunnel to the level 2 router via the level 1/2 router, which acts as an LSR.
- If a prefix is summarized by an ABR, the prefix SID sub-TLV is not propagated with the summarized route between levels. To propagate the node SID for a /32 prefix, route summarization must be disabled.
- The 7210 SAS propagates the prefix SID sub-TLV when exporting the prefix to another IS-IS instance; however, it does not propagate it if the prefix is exported from a different protocol. When the corresponding prefix is redistributed from another protocol, such as OSPF, the prefix SID is removed.

The 7210 SAS originates an adjacency SID sub-TLV with the flags encoded as follows:

- The F-flag is set to 0 to indicate an IPv4 family and 1 to indicate an IPv6 family for the adjacency encapsulation.
- The B-Flag is set to 0 and is not processed on receipt.
- The V-flag is always set to 1.
- The L-flag is always set to 1.
- The S-flag is set to 0 because assigning an adjacency SID to parallel links between neighbors is not supported. A received adjacency SID with the S-flag set is not processed.

- The weight octet is not supported and is set to all zeros.

The system does not originate the SID/Label Binding TLV, but can process it if received. The following rules and limitations should be considered:

- Only the mapping server prefix-SID sub-TLV within the TLV is processed, and the ILMs are installed if the prefixes in the provided range are resolved.
- The range and FEC prefix fields are processed. Each FEC prefix is resolved in the same manner as the prefix SID sub-TLV. In other words, an IP reachability TLV must be received for the exact matching prefix.
- If the same prefix is advertised with both a prefix SID sub-TLV and a mapping server prefix-SID sub-TLV, the system uses the following route preference for resolution:
 - SID received via level 1 in a prefix SID sub-TLV part of the IP reachability TLV
 - SID received via level 2 in a prefix SID sub-TLV part of the IP reachability TLV
 - SID received via level 1 in a mapping server prefix-SID sub-TLV
 - SID received via level 2 in a mapping server prefix-SID sub-TLV
- No route leaking of the entire TLV is performed between levels. However, a level 1/2 router will propagate the prefix-SID sub-TLV from the SID/Label Binding TLV (received from a mapping server) into the IP reachability TLV if the latter is propagated between levels.
- The mapping server that advertises the SID/Label Binding TLV does not need to be in the shortest path for the FEC prefix.
- If the same FEC prefix is advertised in multiple binding TLVs by different routers, the SID in the binding TLV of the first router that is reachable is used. If that router becomes unreachable, the next reachable router is used.
- No check is performed of whether the content of the binding TLVs from different mapping servers is consistent.
- Other sub-TLV, for example, the SID/Label Sub-TLV, ERO metric, and unnumbered interface ID ERO, are ignored. However, the user can run the IGP **show** command to get a list of the octets of the received but unsupported sub-TLVs.

4.1.9.5.2 OSPF control protocol changes

The following TLVs/sub-TLVs are defined in *draft-ietf-ospf-segment-routing-extensions-04* and are required for the implementation of segment routing in OSPF:

- prefix SID sub-TLV part of the OSPFv2 Extended Prefix TLV
- prefix SID sub-TLV part of the OSPFv2 Extended Prefix Range TLV
- adjacency SID sub-TLV part of the OSPFv2 Extended Link TLV
- SID/Label Range Capability TLV
- SR-Algorithm Capability TLV

This section describes the behaviors and limitations of the OSPF support of segment routing TLVs and sub-TLVs.

The 7210 SAS originates a single prefix SID sub-TLV for each OSPFv2 Extended Prefix TLV and processes the first one only if multiple prefix SID sub-TLVs are received within the same OSPFv2 Extended Prefix TLV.

The 7210 SAS encodes the 32-bit index in the prefix SID sub-TLV. The 24-bit label or variable IPv6 SID is not supported.

The 7210 SAS originates a prefix SID sub-TLV with the following flag encoding:

- The NP-flag is always set, meaning that the label for the prefix SID is pushed by the PHP router when forwarding to this router. 7210 SAS PHP routers process a received prefix SID with the NP-flag set to zero and use implicit-null for the outgoing label toward the router that advertised it.
- The M-flag is always unset because the 7210 SAS does not support originating a mapping server prefix-SID sub-TLV.
- The E-flag is always set to 0. The 7210 SAS PHP routers properly process a received prefix SID with the E-flag set to 1, and when the NP-flag is also set to 1, they push explicit-null for the outgoing label towards the router that advertised it.
- The V-flag is always set to 0 to indicate an index value for the SID.
- The L-flag is always set to 0 to indicate that the SID index value is not locally significant.
- The algorithm field is always set to 0 to indicate the SPF algorithm is based on the link metric and is not checked on a received prefix SID sub-TLV.

The system resolves a prefix SID received within an extended prefix TLV based on the following route preference:

- SID received via an intra-area route in a prefix SID sub-TLV part of Extended Prefix TLV
- SID received via an inter-area route in a prefix SID sub-TLV part of Extended Prefix TLV

The 7210 SAS originates an adjacency SID sub-TLV with the following encoding of the flags:

- The F-flag is not set to indicate the adjacency SID refers to an adjacency with outgoing IPv4 encapsulation.
- The B-flag is set to 0 and is not processed on receipt.
- The V-flag is always set.
- The L-flag is always set.
- The S-flag is not supported.
- The weight octet is not supported and is set to all zeros.

The 7210 SAS does not originate the OSPFv2 Extended Prefix Range TLV but can process it if received. The following rules and limitations should be considered:

- Only the prefix SID sub-TLV within the TLV is processed, and the ILMs are installed if the prefixes are resolved.
- The range and address prefix fields are processed. Each prefix is resolved separately.
- If the same prefix is advertised with both a prefix SID sub-TLV in an IP reachability TLV and a mapping server Prefix-SID sub-TLV, the resolution follows the following route preference:
 - the SID received via an intra-area route in a prefix SID sub-TLV part of Extended Prefix TLV
 - the SID received via an inter-area route in a prefix SID sub-TLV part of Extended Prefix TLV
 - the SID received via an intra-area route in a prefix SID sub-TLV part of an OSPFv2 Extended Range Prefix TLV
 - the SID received via an inter-area route in a prefix SID sub-TLV part of an OSPFv2 Extended Range Prefix TLV

- No route leaking of any part of the TLV is allowed between areas. In addition, an ABR does not propagate the prefix-SID sub-TLV from the Extended Prefix Range TLV (received from a mapping server) into an Extended Prefix TLV if the latter is propagated between areas.
- The mapping server that advertised the OSPFv2 extended prefix range TLV does not need to be in the shortest path for the FEC prefix.
- If the same FEC prefix is advertised in multiple OSPFv2 extended prefix range TLVs by different routers, the SID in the TLV of the first router that is reachable is used. If that router becomes unreachable, the next reachable one is used.
- No check is performed to determine whether the contents of the OSPFv2 Extended Prefix Range TLVs received from different mapping servers are consistent.
- Any other sub-TLV (for example, the ERO metric and unnumbered interface ID ERO) is ignored, but the user can get a list of the octets of the received but unsupported sub-TLVs using the existing IGP **show** command.

The 7210 SAS supports the propagation on ABR of an external prefix LSA into other areas with the route type set to 3 as per *draft-ietf-ospf-segment-routing-extensions-04*.

The 7210 SAS supports the propagation on ABR of external prefix LSAs with route type 7 from the NSSA area into other areas with the route type set to 5, as described in *draft-ietf-ospf-segment-routing-extensions-04*. The system does not support the propagation of the prefix SID sub-TLV between OSPF instances.

If an OSPF import policy is configured, the outcome of the policy applies to prefixes resolved in RTM and the corresponding tunnels in TTM. A prefix that is removed by the policy is removed as both a route in the RTM and as an SR tunnel in the TTM.

4.1.9.6 BGP label route resolution using segment routing tunnels

Configure the following CLI commands to enable the resolution of RFC 3107 BGP label route prefixes using SR tunnels to BGP next-hops in the TTM:

```
configure>router>bgp>next-hop-resolution
    labeled-route-transport-tunnel
        [no] family family
            resolution {any | disabled | filter}
            resolution-filter
                [no] sr-isis
                [no] sr-ospf
        exit
    exit
exit
exit
```

If the **resolution** option is explicitly set to **disabled**, the default binding to LDP tunnel is used. If **resolution** option is set to **any**, a supported tunnel type from the BGP label route context is selected following the TTM preference.

The following tunnel types are supported in a BGP label route context and are listed in order of preference:

- RSVP
- LDP
- segment routing

When **sr-isis** or **sr-ospf** is configured using the **resolution-filter** option, a tunnel to the BGP next-hop is selected in the TTM from the lowest numbered IS-IS or OSPF instance.

See the [BGP](#) chapter for information about BGP label route resolution using SR tunnels.

4.1.9.7 Service packet forwarding with segment routing

The following SDP subtypes of the MPLS type allow service binding to an SR tunnel programmed in the TTM by OSPF or IS-IS:

- **config>service>sdp>sr-isis**
- **config>service>sdp>sr-ospf**

SDPs of type **sr-isis** or **sr-ospf** can be configured with the **far-end** CLI command. When the **sr-isis** or **sr-ospf** command is enabled, a tunnel to the far-end address is selected in the TTM from the lowest preference IS-IS or OSPF instance. If multiple instances have the same lowest preference from the lowest numbered IS-IS or OSPF instance, the SR-ISIS or SR-OSPF tunnel is selected at the time of the binding, using the tunnel selection rules. If a preferred tunnel is subsequently added to the TTM, the SDP will not automatically switch to the new tunnel until the next time the SDP is being re-resolved.

The **tunnel-far-end** option is not supported. In addition, the **mixed-lsp-mode** option does not support the **sr-isis** and **sr-ospf** tunnel types.

The signaling protocol for the service labels of an SDP using an SR tunnel can be configured to static (**off**), T-LDP (**tlsp**), or BGP (**bgp**).

SR tunnels can be configured in a VPRN service with the **auto-bind-tunnel** command.

VPN-IPv4 and VPN-IPv6 (6VPE) are supported in a VPRN or BGP EVPN service using segment routing transport tunnels with the **auto-bind-tunnel** command.

See [BGP](#) and the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Services Guide* for more information about the VPRN **auto-bind-tunnel** CLI command.

The following service contexts are supported with SR tunnels:

- VLL, LDP, VPLS, and RVPLS
- BGP-AD VPLS when the **use-provisioned-sdp** option is enabled in PW template binding
- intra-AS BGP VPRN for VPN-IPv4 and VPN-IPv6 prefixes with both auto-bind and explicit SDP

The following service contexts are not supported:

- inter-AS VPRN
- dynamic MS-PW, PW-switching
- BGP-AD VPLS with auto-generation of SDP using an SR tunnel when binding to a PW template

4.1.9.8 Mirror services



Note:

SR tunnels for mirror services are not supported on 7210 SAS platforms.

The user can configure a spoke-SDP bound to an SR tunnel to forward mirrored packets from a mirror source to a remote mirror destination. In the configuration of the mirror destination service at the destination node, the **remote-source** command must use a spoke-SDP with a VC-ID that matches the

VC-ID configured in the mirror destination service at the mirror source node. The **far-end** option is not supported with an SR tunnel.

Use the following syntax to configure a mirror source node.

```
config mirror mirror-dest service-id
no spoke-sdp <sdp-id:vc-id>
spoke-sdp <sdp-id:vc-id> [create]
    egress
        vc-label <egress-vc-label>
```



Note:

- *sdp-id* matches an SDP that uses an SR tunnel.
- For **vc-label**, both static and T-LDP egress VC labels are supported.

Use the following syntax to configure a mirror destination node.

```
configure mirror mirror-dest service-id remote-source
    spoke-sdp <SDP-ID>:<VC-ID> create <-- VC-ID matching that of spoke-sdp configured in
    mirror destination context at mirror source node.
        ingress
            vc-label <ingress-vc-label> <--- optional: both static and t-ldp ingress vc
            label are supported.
        exit
        no shutdown
    exit
exit
```



Note:

- The **far-end** command in the **config>mirror>mirror-dest>remote-source** context is not supported with SR tunnels at a mirror destination node; the user must reference a spoke-SDP using a segment routing SDP coming from a mirror source node.
- For **vc-label**, both static and T-LDP ingress VC labels are supported.


4.1.10 IGP-LDP synchronization

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C support IGP-LDP synchronization on IS-IS routes. For information, see "IGP-LDP and Static Route-LDP Synchronization on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C" in the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*.

4.2 IS-IS configuration process overview

The following figure shows the process to provision basic IS-IS parameters.

Figure 16: IS-IS configuration and implementation flow



4.3 Configuration notes

This section describes IS-IS configuration caveats.

4.3.1 General

- IS-IS must be enabled on each participating routers.
- There are no default network entity titles.
- There are no default interfaces.
- By default, routers are assigned a Level 1/Level 2 level capability.
- In network mode, the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C allow the configuration of a single instance at any specific time. The instance ID can be any number other than 0. This enables the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C devices to be used in a network where multi-instance IS-IS is deployed and the node needs to use an instance ID other than the default instance ID of 0.

4.4 Configuring IS-IS with CLI

This section provides information to configure intermediate-system-to-intermediate-system (IS-IS) using the command line interface.

4.5 IS-IS configuration overview

4.5.1 Router levels

The router level capability can be configured globally and on a per-interface basis. The interface-level parameters specify the interface's routing level. The neighbor capability and parameters define the adjacencies that are established.

IS-IS is not enabled by default. When IS-IS is enabled, the global default level capability is Level 1/2 which enables the router to operate as either a Level 1 and/or a Level 2 router with the associated databases. The router runs separate shortest path first (SPF) calculations for the Level 1 area routing and for the Level 2 multi-area routing to create the IS-IS routing table.

The level value can be modified on both or either of the global and interface levels to be only Level 1-capable, only Level 2-capable or Level 1 *and* Level 2-capable.

If the default value is not modified on any routers in the area, then the routers try to form both Level 1 and Level 2 adjacencies on all IS-IS interfaces. If the default values are modified to Level 1 or Level 2, then the number of adjacencies formed are limited to that level only.

4.5.2 Area address attributes

The **area-id** command specifies the area address portion of the NET which is used to define the IS-IS area to which the router will belong. At least one **area-id** command should be configured on each router participating in IS-IS. A maximum of three **area-id** commands can be configured per router.

The area address identifies a point of connection to the network, such as a router interface, and is called a *network service access point (NSAP)*. The routers in an area manage routing tables about destinations within the area. The Network Entity Title (NET) value is used to identify the IS-IS area to which the router belongs.

NSAP addresses are divided into three parts. Only the Area ID portion is configurable.

- **Area ID**

A variable length field between 1 and 13 bytes. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.

- **System ID**

A six-byte system identification. This value is not configurable. The system ID is derived from the system or router ID.

- **Selector ID**

A one-byte selector identification that must contain zeros when configuring a NET. This value is not configurable. The selector ID is always 00.

The following is a sample of ISO addresses in IS-IS address format.

MAC address 00:a5:c7:6b:c4:90	49.0011.00a5.c76b.c490.00
IP address: 218.112.14.5	49.0011.2181.1201.4005.00

4.5.3 Interface level capability

The level capability value configured on the interface level is compared to the level capability value configured on the global level to determine the type of adjacencies that can be established. The default level capability for routers and interfaces is Level 1/2.

The following table lists configuration combinations and the potential adjacencies that can be formed.

Table 43: Potential adjacency capabilities

Global level	Interface level	Potential adjacency
L 1/2	L 1/2	Level 1 and/or Level 2
L 1/2	L 1	Level 1 only
L 1/2	L 2	Level 2 only
L 2	L 1/2	Level 2 only
L 2	L 2	Level 2 only
L 2	L 1	none
L 1	L 1/2	Level 1 only
L 1	L 2	none
L 1	L 1	Level 1 only

4.5.4 Route leaking

The Nokia implementation of IS-IS route leaking is performed in compliance with RFC 2966, *Domain-wide Prefix Distribution with Two-Level IS-IS*. As previously stated, IS-IS is a routing domain (an autonomous system running IS-IS) which can be divided into Level 1 areas with a Level 2-connected subset (backbone) of the topology that interconnects all of the Level 1 areas. Within each Level 1 area, the routers exchange link state information. Level 2 routers also exchange Level 2 link state information to compute routes between areas.

Routers in a Level 1 area typically only exchange information within the Level 1 area. For IP destinations not found in the prefixes in the Level 1 database, the Level 1 router forwards PDUs to the nearest router that is in both Level 1/Level 2 with the attached bit set in its Level 1 link-state PDU.

There are many reasons to implement domain-wide prefix distribution. The goal of domain-wide prefix distribution is to increase the granularity of the routing information within the domain. The routing mechanisms specified in RFC 1195 are appropriate in many situations and account for excellent scalability properties. However, in certain circumstances, the amount of scalability can be adjusted which can distribute more specific information than described by RFC 1195.

Distributing more prefix information can improve the quality of the resulting routes. A well known property of default routing is that loss of information can occur. This loss of information affects the computation of a route based upon less information which can result in sub-optimal routes.

4.6 Basic IS-IS configuration

About this task

For IS-IS to operate on routers, IS-IS must be explicitly enabled, and at least one area address and interface must be configured. If IS-IS is enabled but no area address or interface is defined, the protocol is enabled but no routes are exchanged. When at least one area address and interface are configured, then adjacencies can be formed and routes exchanged.

To configure IS-IS, perform the following tasks:

Procedure

Step 1. Enable IS-IS (specifying the instance ID of multi-instance IS-IS is to be enabled).

Step 2. If necessary, modify the level capability on the global level (default is level-1/2).

Step 3. Define area address(es).

Step 4. Configure IS-IS interfaces.

Example: IS-IS default values

```
*A:Dut-D>config>router>isis# info detail
-----
no system-id
no router-id
level-capability level-1/2
no graceful-restart
no auth-keychain
no authentication-key
no authentication-type
authentication-check
csnp-authentication
no ignore-lsp-errors
no ignore-narrow-metric
lsp-lifetime 1200
lsp-mtu-size 1492
lsp-refresh-interval 600
no export-limit
no export
no import
hello-authentication
psnp-authentication
no traffic-engineering
no reference-bandwidth
no default-route-tag
no disable-ldp-sync
no advertise-passive-only
no advertise-router-capability
no hello-padding
no ldp-over-rsvp
no advertise-tunnel-link
no ignore-attached-bit
no suppress-attached-bit
no iid-tlv-enable
no poi-tlv-enable
no prefix-limit
no loopfree-alternate
no loopfree-alternate-exclude
no rib-priority high
```

```
ipv4-routing
no ipv6-routing
no multi-topology
no unicast-import-disable both
no strict-adjacency-check
igp-shortcut
shutdown
tunnel-next-hop
    family ipv4
        resolution disabled
        resolution-filter
        no rsvp
        no sr-te
    exit
    family ipv6
        resolution disabled
        resolution-filter
        no rsvp
        no sr-te
    exit
    family srv4
        resolution disabled
        resolution-filter
        no rsvp
        no sr-te
    exit
    family srv6
        resolution disabled
        resolution-filter
        no rsvp
        no sr-te
    exit
exit
exit
timers
    lsp-wait 5000 lsp-initial-wait 10 lsp-second-wait 1000
    sfp-wait 10000 sfp-initial-wait 1000 sfp-second-wait 1000
exit
level 1
    advertise-router-capability
    no hello-padding
    no lsp-mtu-size
    no auth-keychain
    no authentication-key
    no authentication-type
    csnp-authentication
    external-preference 160
    hello-authentication
    no loopfree-alternate-exclude
    preference 15
    psnp-authentication
    no wide-metrics-only
    default-metric 10
    default-ipv6-unicast-metric 10
exit
level 2
    advertise-router-capability
    no hello-padding
    no lsp-mtu-size
    no auth-keychain
    no authentication-key
    no authentication-type
    csnp-authentication
    external-preference 165
```

```
hello-authentication
no loopfree-alternate-exclude
preference 18
psnp-authentication
no wide-metrics-only
default-metric 10
default-ipv6-unicast-metric 10
exit
segment-routing
shutdown
adj-sid-hold 15
no export-tunnel-table
no prefix-sid-range
no tunnel-table-pref
no tunnel-mtu
mapping-server
shutdown
exit
exit
no shutdown
```

4.7 Common configuration tasks

To implement IS-IS in your network, you must enable IS-IS on each participating routers.

To assign different level capabilities to the routers and organize your network into areas, modify the level capability defaults on end systems from Level 1/2 to Level 1. Routers communicating to other areas can retain the Level 1/2 default.

On each router, at least one area ID also called the area address should be configured as well as at least one IS-IS interface.

- Enable IS-IS.
- Configure global IS-IS parameters.
 - Configure area address(es).
- Configure IS-IS interface-specific parameters.

4.8 Configuring IS-IS components

The following section describes the syntax used to configure the IS-IS components.

4.8.1 Enabling IS-IS

IS-IS must be enabled in order for the protocol to be active.



Note:

Careful planning is essential to implement commands that can affect the behavior of global and interface levels.

To configure IS-IS on a router, use the **config>router router-name>isis** [*isis-instance*] command.

IS-IS also supports the concept of multi-instance IS-IS which allows separate instances of the IS-IS protocol to run independently of the 7210 SAS router. Separate instances are created by adding a different instance ID as the optional parameter in the **config>router>isis** command.



Note:

Not all 7210 SAS platforms support use of multi-instances simultaneously. For more information, see the preceding configuration notes.

4.8.2 Modifying router-level parameters

When IS-IS is enabled, the default **level-capability** is Level 1/2. This means that the router operates with both Level 1 and Level 2 routing capabilities. To change the default value in order for the router to operate as a Level 1 router or a Level 2 router, you must explicitly modify the **level** value.

If the level is modified, the protocol shuts down and restarts. Doing this can affect adjacencies and routes.

The **level-capability** value can be configured on the global level and also on the interface level. The **level-capability** value determines which level values can be assigned on the router level or on an interface-basis.

In order for the router to operate as a Level 1 only router or as a Level 2 only router, you must explicitly specify the **level-number** value.

- Select **level-1** to route only within an area.
- Select **level-2** to route to destinations outside an area, toward other eligible Level 2 routers.

Use the following command syntax to configure the router level.

```
config>router# isis
  level-capability {level-1|level-2|level-1/2}
  level {1|2}
```

Example: Command usage to configure router level

```
config>router# isis
config>router>isis# level-capability 1/2
config>router>isis# level 2
```

Example: Configuration output

```
A:ALA-A>config>router>isis# info
#-----
echo "ISIS"
#-----

level-capability level-1/2
level 2

-----
A:ALA-A>config>router>isis#
```

4.8.3 Configuring ISO area addresses

Use the following syntax to configure an area ID also called an address. A maximum of 3 area-id can be configured.

```
config>router# isis
      area-id area-address
```

Example: Command usage to configure area ID

```
config>router>isis#
config>router>isis# area-id 49.0180.0001
config>router>isis# area-id 49.0180.0002
config>router>isis# area-id 49.0180.0003
```

Example: Area ID configuration output

```
A:ALA-A>config>router>isis# info
-----
      area-id 49.0180.0001
      area-id 49.0180.0002
      area-id 49.0180.0003
-----
A:ALA-A>config>router>isis#
```

4.8.4 Configuring global IS-IS parameters

Commands and parameters configured on the global level are inherited to the interface levels. Parameters specified in the interface and interface-level configurations take precedence over global configurations.

Example: Command usage to configure the global-level IS-IS

```
config>router# isis
config>router>isis#
config>router>isis# level-capability level-2
config>router>isis# authentication-check
config>router>isis# authentication-type password
config>router>isis# authentication-key test
config>router>isis# overload timeout 90
config>router>isis# traffic-engineering
```

Example: Modified global-level configuration output

```
A:ALA-A>config>router>isis# info
-----
      level-capability level-2
      area-id 49.0180.0001
      area-id 49.0180.0002
      area-id 49.0180.0003
      authentication-key "H5KBAWrAAQU" hash
      authentication-type password
      overload timeout 90
      traffic-engineering
-----
A:ALA-A>config>router>isis#
```

4.8.5 Migration to IS-IS multi-topology

To migrate to IS-IS multi-topology for IPv6, perform the following tasks:

Use the following syntax to enable the sending/receiving of IPv6 unicast reachability information in IS-IS MT TLVs on all the routers that support MT.

```
config>router# isis multi-topology ipv6-unicast
```

Example

```
A:SAS-12>config>router>isis# info detail
-----
...
    ipv4-routing
    ipv6-routing native
    multi-topology
        ipv6-unicast
    exit
...
-----
A:SAS-12>config>router>isis#
```

Use the following syntax to ensure that all MT routers have the IPv6 reachability information required by MT TLVs.

```
show>router# isis topology ipv6-unicast
```

Example

```
A:SAS-12>config>router>isis# show router isis topology ipv6-unicast
=====
Topology Table
=====
Node                               Interface                             Nexthop
-----
No Matching Entries
=====
A:SAS-12>config>router>isis#
```

```
show>router# isis database detail
```

Example

```
A:SAS-12>>config>router>isis# show router isis database detail
=====
Rtr Base ISIS Instance 0 Database (detail)
=====
Displaying Level 1 database
-----
LSP ID       : ALA-49.00-00                      Level       : L1
Sequence     : 0x22b                               Checksum    : 0x60e4  Lifetime    : 1082
Version      : 1                                   Pkt Type    : 18      Pkt Ver     : 1
Attributes: L1L2                                Max Area    : 3
SysID Len    : 6                                   Used Len    : 404     Alloc Len   : 1492

TLVs :
```

```

Area Addresses :
  Area Address : (13) 47.4001.8000.00a7.0000.ffdd.0007
Supp Protocols :
  Protocols : IPv4 IPv6
IS-Hostname :
  Hostname : ALA-49
TE Router ID :
  Router ID : 10.10.10.104
Internal Reach :
  IP Prefix : 10.10.10.104/32 (Dir. :Up) Metric : 0 (I)
  IP Prefix : 10.10.4.0/24 (Dir. :Up) Metric : 10 (I)
  IP Prefix : 10.10.5.0/24 (Dir. :Up) Metric : 10 (I)
  IP Prefix : 10.10.7.0/24 (Dir. :Up) Metric : 10 (I)
  IP Prefix : 10.10.0.0/24 (Dir. :Up) Metric : 10 (I)
  IP Prefix : 10.0.0.0/24 (Dir. :Up) Metric : 10 (I)
MT IPv6 Reach. :
  MT ID : 2
  IPv6 Prefix : 3ffe::101:100/120
    Flags : Up Internal Metric : 10
  IPv6 Prefix : 10::/64
    Flags : Up Internal Metric : 10
I/f Addresses :
  IP Address : 10.10.10.104
  IP Address : 10.10.4.3
  IP Address : 10.10.5.3
  IP Address : 10.10.7.3
  IP Address : 10.10.0.16
  IP Address : 10.0.0.104
I/f Addresses IPv6 :
  IPv6 Address : 3FFE::101:101
  IPv6 Address : 10::104
TE IP Reach. :
  IP Prefix : 10.10.10.104/32 (Dir. :Up) Metric : 0
  IP Prefix : 10.10.4.0/24 (Dir. :Up) Metric : 10
  IP Prefix : 10.10.5.0/24 (Dir. :Up) Metric : 10
  IP Prefix : 10.10.7.0/24 (Dir. :Up) Metric : 10
  IP Prefix : 10.10.0.0/24 (Dir. :Up) Metric : 10
  IP Prefix : 10.0.0.0/24 (Dir. :Up) Metric : 10
Authentication :
  Auth Type : Password(1) (116 bytes)

Level (1) LSP Count : 1

Displaying Level 2 database
-----
LSP ID : ALA-49.00-00
Sequence : 0x22c
Version : 1
Attributes: L1L2
SysID Len : 6
Checksum : 0xb888
Pkt Type : 20
Max Area : 3
Used Len : 304
Level : L2
Lifetime : 1082
Pkt Ver : 1
Alloc Len : 1492

TLVs :
Area Addresses :
  Area Address : (13) 47.4001.8000.00a7.0000.ffdd.0007
Supp Protocols :
  Protocols : IPv4 IPv6
IS-Hostname :
  Hostname : ALA-49
TE Router ID :
  Router ID : 10.10.10.104
Internal Reach :
  IP Prefix : 10.10.10.104/32 (Dir. :Up) Metric : 0 (I)
  IP Prefix : 10.10.4.0/24 (Dir. :Up) Metric : 10 (I)
  IP Prefix : 10.10.5.0/24 (Dir. :Up) Metric : 10 (I)

```



```

IP Prefix      : 10.10.7.0/24      (Dir. :Up) Metric : 10 (I)
IP Prefix      : 10.10.0.0/24      (Dir. :Up) Metric : 10 (I)
IP Prefix      : 10.0.0.0/24       (Dir. :Up) Metric : 10 (I)
MT IPv6 Reach. :
MT ID          : 2
IPv6 Prefix    : 3ffe::101:100/120
                Flags : Up Internal Metric : 10
IPv6 Prefix    : 10::/64
                Flags : Up Internal Metric : 10
I/f Addresses  :
IP Address     : 10.10.10.104
IP Address     : 10.10.4.3
IP Address     : 10.10.5.3
IP Address     : 10.10.7.3
IP Address     : 10.10.0.16
IP Address     : 10.0.0.104
I/f Addresses IPv6 :
IPv6 Address   : 3FFE::101:101
IPv6 Address   : 10::104
TE IP Reach.   :
IP Prefix      : 10.10.10.104/32    (Dir. :Up) Metric : 0
IP Prefix      : 10.10.4.0/24       (Dir. :Up) Metric : 10
IP Prefix      : 10.10.5.0/24       (Dir. :Up) Metric : 10
IP Prefix      : 10.10.7.0/24       (Dir. :Up) Metric : 10
IP Prefix      : 10.10.0.0/24       (Dir. :Up) Metric : 10
IP Prefix      : 10.0.0.0/24       (Dir. :Up) Metric : 10
Authentication :
Auth Type      : MD5(54) (16 bytes)

Level (2) LSP Count : 1
-----
Flags : D = Prefix Leaked Down
       : N = Node Flag
       : R = Re-advertisement Flag
       : S = Sub-TLVs Present
       : X = External Prefix Flag
=====
A:SAS-12>>config>router>isis#

```

Use the following syntax to configure MT TLVs for IPv6 SPF.

```
config>router# isis ipv6-routing mt
```

Example

```

A:SAS-12>>config>router>isis# info detail
-----
...
    ipv4-routing
    ipv6-routing mt
    multi-topology
        ipv6-unicast
    exit
...
-----
A:SAS-12>>config>router>isis#

```

Use the following syntax to verify IPv6 routes.

```
show>router# isis routes ipv6-unicast
```

Example

```
A:SAS-12>>config>router>isis# show router isis routes ipv6-unicast
=====
Rtr Base ISIS Instance 0 Route Table
=====
Prefix[Flags]                Metric    Lvl/Typ    Ver.  SysID/Hostname
NextHop                      MT        AdminTag/SID[F]
-----
No Matching Entries
=====
A:SAS-12>>config>router>isis#
```

```
show>router# route-table ipv6
```

Example

```
A:SAS-12>>show>router# route-table ipv6
=====
IPv6 Route Table (Router: Base)
=====
Dest Prefix                Type    Proto    Age        Pref
Next Hop[Interface Name]  Metric
-----
10::/64                    Local   Local    05h35m28s  0
to-104                    0
-----
No. of Routes: 1
=====
A:SAS-12>
```

4.8.6 Configuring interface parameters

There are no interfaces associated with IS-IS by default. An interface belongs to all areas configured on a router. Interfaces cannot belong to separate areas. There are no default interfaces applied to the router's IS-IS instance. You must configure at least one IS-IS interface in order for IS-IS to work.

To enable IS-IS on an interface, first configure an IP interface in the **config>router>interface** context. Then, apply the interface in the **config>router>isis>interface** context.

You can configure both the Level 1 parameters and the Level 2 parameters on an interface. The **level-capability** value determines which level values are used.



Note:

For point-to-point interfaces, only the values configured under Level 1 are used regardless of the operational level of the interface.

Example: Modified interface parameters output

```
config>router# isis
config>router>isis# level 1
config>router>isis>level# wide-metrics-only
config>router>isis>level# exit
config>router>isis# level 2
config>router>isis>level# wide-metrics-only
config>router>isis>level# exit
config>router>isis# interface ALA-1-2
```

```
config>router>isis>if# level-capability level-2
config>router>isis>if# mesh-group 85
config>router>isis>if# exit
config>router>isis# interface ALA-1-3
config>router>isis>if# level-capability level-1
config>router>isis>if# interface-type point-to-point
config>router>isis>if# mesh-group 101
config>router>isis>if# exit
config>router>isis# interface ALA-1-5
config>router>isis>if# level-capability level-1
config>router>isis>if# interface-type point-to-point
config>router>isis>if# mesh-group 85
config>router>isis>if# exit
config>router>isis# interface to-103
config>router>isis>if# level-capability level-1/2
>router>isis>if# mesh-group 101
config>router>isis>if# exit
config>router>isis#
```

Example: Global and interface-level configurations output

```
A:ALA-A>config>router>isis# info
-----
level-capability level-2
area-id 49.0180.0001
area-id 49.0180.0002
area-id 49.0180.0003
authentication-key "H5KBAWrAAQU" hash
authentication-type password
traffic-engineering
level 1
    wide-metrics-only
exit
level 2
    wide-metrics-only
exit
interface "system"
exit
interface "ALA-1-2"
    level-capability level-2
    mesh-group 85
exit
interface "ALA-1-3"
    level-capability level-1
    interface-type point-to-point
    mesh-group 101
exit
interface "ALA-1-5"
    level-capability level-1
    interface-type point-to-point
    mesh-group 85
exit
interface "to-103"
    mesh-group 101
exit
-----
A:ALA-A>config>router>isis#
```

4.8.6.1 Example: configuring a Level 1 area

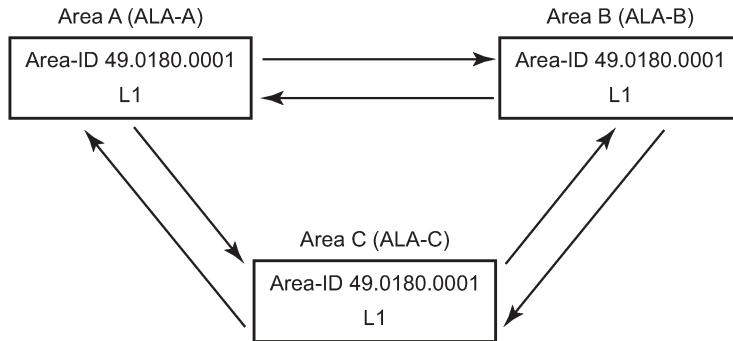


Note:

Interfaces are configured in the **config>router>interface** context.

The following figure shows the configuration of a Level 1 area.

Figure 17: Configuring a Level 1 area



OSRG031

Example: Command usage to configure a Level 1 area

```

A:ALA-A>config>router# isis
A:ALA-A>config>router>isis# area-id 47.0001
A:ALA-A>config>router>isis# level-capability level-1
A:ALA-A>config>router>isis# interface system
A:ALA-A>config>router>isis>if# exit
A:ALA-A>config>router>isis# interface A-B
A:ALA-A>config>router>isis>if# exit
A:ALA-A>config>router>isis# interface A-C
A:ALA-A>config>router>isis>if# exit
A:ALA-A>config>router>isis#

A:ALA-B>config>router# isis
A:ALA-B>config>router>isis# area-id 47.0001
A:ALA-B>config>router>isis# level-capability level-1
A:ALA-B>config>router>isis# interface system
A:ALA-B>config>router>isis>if# exit
A:ALA-B>config>router>isis# interface B-A
A:ALA-B>config>router>isis>if# exit
A:ALA-B>config>router>isis# interface B-C
A:ALA-B>config>router>isis>if# exit
A:ALA-B>config>router>isis#

A:ALA-C>config>router# isis
A:ALA-C>config>router>isis# area-id 47.0001
A:ALA-C>config>router>isis# level-capability level-1
A:ALA-C>config>router>isis# interface system
A:ALA-C>config>router>isis>if# exit
A:ALA-C>config>router>isis# interface "C-A"
A:ALA-C>config>router>isis>if# exit
A:ALA-C>config>router>isis# interface "C-B"
A:ALA-C>config>router>isis>if# exit

A:ALA-A>config>router>isis# info
-----
  
```

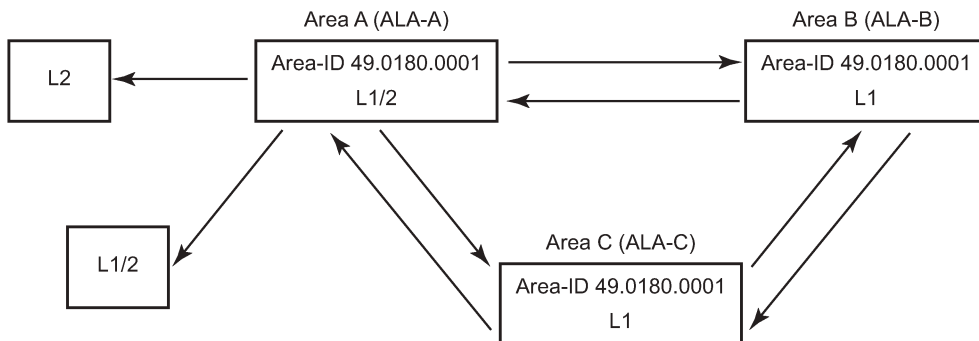
```
level-capability level-1
area-id 49.0180.0001
interface "system"
exit
interface "A-B"
exit
interface "A-C"
exit
-----
A:ALA-A>config>router>isis#
A:ALA-B>config>router>isis# info
-----
level-capability level-1
area-id 49.0180.0001
interface "system"
exit
interface "B-A"
exit
interface "B-C"
exit
-----
A:ALA-B>config>router>isis#
A:ALA-C>config>router>isis# info
#-----
echo "ISIS"
-----
level-capability level-1
area-id 49.0180.0001
interface "system"
exit
interface "C-A"
exit
interface "C-B"
exit
-----
A:ALA-C>config>router>isis#
```

4.8.6.2 Example: modifying a router's level capability

In [Example: configuring a Level 1 area](#), ALA-A, ALA-B, and ALA-C are configured as Level 1 systems. Level 1 systems communicate with other Level 1 systems in the same area. In this example, ALA-A is modified to set the level capability to Level 1/2. Now, the Level 1 systems in the area with NET 47.0001 forward PDUs to ALA-A for destinations that are not in the local area.

The following figure shows the configuration of Level 1/2 area.

Figure 18: Configuring a Level 1/2 area



OSRG036

Example: Command usage to configure a Level 1/2 system

```

A:ALA-A>config>router# isis
A:ALA-A>config>router>isis# level-capability level-1/2
    
```

4.9 IS-IS configuration management tasks

This section describes the IS-IS configuration management tasks.

4.9.1 Disabling IS-IS

The **shutdown** command disables the IS-IS protocol instance on the router. The configuration settings are not changed, reset, or removed.

Use the following syntax to disable IS-IS on a router.

```

config>router# isis
shutdown
    
```

4.9.2 Removing IS-IS

The **no isis** command deletes the IS-IS protocol instance. The IS-IS configuration reverts to the default settings.

Use the following syntax to remove the IS-IS configuration.

```

config>router#
no isis
    
```

4.9.3 Modifying global IS-IS parameters

You can modify, disable, or remove global IS-IS parameters without shutting down entities. Changes take effect immediately. Modifying the level capability on the global level causes the IS-IS protocol to restart.

Example: Command usage to modify various parameters

```
config>router>isis# overload timeout 500
config>router>isis# level-capability level-1/2
config>router>isis# no authentication-check
config>router>isis# authentication-key raiderslost
```

Example: Global modifications output

```
A:ALA-A>config>router>isis# info
-----
    area-id 49.0180.0001
    area-id 49.0180.0002
    area-id 49.0180.0003
    authentication-key "//oZrvtvFPn06S42lRIJsE" hash
    authentication-type password
    no authentication-check
    overload timeout 500 on-boot
    level 1
        wide-metrics-only
    exit
    level 2
        wide-metrics-only
    exit
    interface "system"
    exit
    interface "ALA-1-2"
        level-capability level-2
        mesh-group 85
    exit
    interface "ALA-1-3"
        level-capability level-1
        interface-type point-to-point
        mesh-group 101
    exit
    interface "ALA-1-5"
        level-capability level-1
        interface-type point-to-point
        mesh-group 85
    exit
    interface "to-103"
        mesh-group 101
    exit
    interface "A-B"
    exit
    interface "A-C"
    exit
-----
A:ALA-A>config>router>isis#
```

4.9.4 Modifying IS-IS interface parameters

You can modify, disable, or remove interface-level IS-IS parameters without shutting down entities. Changes take effect immediately. Modifying the level capability on the interface causes the IS-IS protocol on the interface to restart.

To remove an interface, issue the **no interface** *ip-int-name* command. To disable an interface, issue the **shutdown** command in the interface context.

Example: Command usage interface IS-IS modification

```
config>router# isis
config>router>isis# interface ALA-1-3
config>router>isis>if# mesh-group 85
config>router>isis>if# passive
config>router>isis>if# lsp-pacing-interval 5000
config>router>isis>if# exit
config>router>isis# interface to-103
config>router>isis>if# hello-authentication-type message-digest
config>router>isis>if# hello-authentication-key 49ersrule
config>router>isis>if# exit
```

Example: Modified interface parameters output

```
A:ALA-A>config>router>isis# info
-----
area-id 49.0180.0001
area-id 49.0180.0002
area-id 49.0180.0003
authentication-key "//oZrvtvFPn06S42lRIJsE" hash
authentication-type password
no authentication-check
overload timeout 500 on-boot
level 1
    wide-metrics-only
exit
level 2
    wide-metrics-only
exit
interface "system"
exit
interface "ALA-1-2"
    level-capability level-2
    mesh-group 85
exit
interface "ALA-1-3"
    level-capability level-1
    interface-type point-to-point
    lsp-pacing-interval 5000
    mesh-group 85
    passive
exit
interface "ALA-1-5"
    level-capability level-1
    interface-type point-to-point
    mesh-group 85
exit
interface "to-103"
    hello-authentication-key "DvR3l264KQ6vXMTvbAZ1mE" hash
    hello-authentication-type message-digest
```



```

        mesh-group 101
        exit
        interface "A-B"
        exit
    -----
A:ALA-A>config>router>isis#

```

4.9.5 Configuring leaking

IS-IS allows a two-level hierarchy to route PDUs. Level 1 areas can be interconnected by a contiguous Level 2 backbone.

The Level 1 link-state database contains information only about that area. The Level 2 link-state database contains information about the Level 2 system and each of the Level 1 systems in the area. A Level 1/2 router contains information about both Level 1 and Level 2 databases. A Level 1/2 router advertises information about its Level 1 area toward the other Level 1/2 or Level 2 (only) routers.

Packets with destinations outside the Level 1 area are forwarded toward the closest Level 1/2 router which, in turn, forwards the packets to the destination area.

Sometimes, the shortest path to an outside destination is not through the closest Level 1/2 router, or, the only Level 1/2 system to forward packets out of an area is not operational. Route leaking provides a mechanism to leak Level 2 information to Level 1 systems to provide routing information regarding inter-area routes. Then, a Level 1 router has more options to forward packets.

Configure a route policy to leak routers from Level 2 into Level 1 areas in the **config>router>policy-options>policy-statement** context, as shown in the following example.

Example

The following shows the command usage to configure prefix list and policy statement parameters in the **config>router** context.

```

config>router>policy-options# prefix-list loops
..>policy-options>prefix-list# prefix 10.1.1.0/24 longer
..>policy-options>prefix-list# exit
..>policy-options# policy-statement leak
..>policy-options>policy-statement# entry 10
..>policy-options>policy-statement>entry# from
..>policy-options>policy-statement>entry>from# prefix-list loops
..>policy-options>policy-statement>entry>from# level 2
..>policy-options>policy-statement>entry>from# exit
..>policy-options>policy-statement>entry# to
..>policy-options>policy-statement>entry>to# level 1
..>policy-options>policy-statement>entry>to# exit
..>policy-options>policy-statement>entry# action accept
..>policy-options>policy-statement>entry>action# exit
..>policy-options>policy-statement>entry# exit
..>policy-options>policy-statement# exit
..>policy-options# commit
..>policy-options#

```

```

A:ALA-A>config>router>policy-options# info
-----
        prefix-list "loops"
            prefix 10.1.1.0/24 longer
        exit
        policy-statement "leak"
            entry 10

```

```

        from
        prefix-list "loop"
        level 2
    exit
    to
        level 1
    exit
    action accept
    exit
exit
exit
-----
A:ALA-A>config>router>policy-options#

```

Next, use the following commands to apply the policy to leak routes from Level 2 info Level 1 systems on ALA-A.

```

config>router#isis
config>router>isis# export leak

```

```

A:ALA-A>config>router>isis# info
-----
    area-id 49.0180.0001
    area-id 49.0180.0002
    area-id 49.0180.0003
    authentication-key "//oZrvtvFPn06S42lRIJsE" hash
    authentication-type password
    no authentication-check
    export "leak"
    ...
-----
A:ALA-A>config>router>isis#

```

After the policy is applied, create a policy to redistribute external IS-IS routes from Level 1 systems into the Level 2 backbone (see [Redistributing external IS-IS routers](#)).

In the **config>router** context, the following commands can be used to configure the following policy statement parameters.

```

config>router>policy-options# begin
..>policy-options# policy-statement "isis-ext"
..>policy-options>policy-statement# entry 10
..>policy-options>policy-statement>entry$ from
..>policy-options>policy-statement>entry>from$ external
..>policy-options>policy-statement>entry>from# exit
..>policy-options>policy-statement>entry# to
..>policy-options>policy-statement>entry>to$ level 2
..>policy-options>policy-statement>entry>to# exit
..>policy-options>policy-statement>entry# action accept
..>policy-options>policy-statement>entry>action# exit
..>policy-options>policy-statement>entry# exit
..>policy-options>policy-statement# exit
..>policy-options# commit

A:ALA-A>config>router>policy-options# info
-----
    prefix-list "loops"
    prefix 10.1.1.0/24 longer
exit

```

```
policy-statement "leak"
  entry 10
    from
      prefix-list "loop"
      level 2
    exit
    to
      level 1
    exit
    action accept
    exit
  exit
exit
policy-statement "isis-ext"
  entry 10
    from
      external
    exit
    to
      level 2
    exit
    action accept
    exit
  exit
exit
-----
A:ALA-A>config>router>policy-options#
```

4.9.6 Redistributing external IS-IS routers

IS-IS does not redistribute Level 1 external routes into Level 2 by default. You must explicitly apply the policy to redistribute external IS-IS routes. Policies are created in the **config>router>policy-options** context. See the [Route policies](#) section of this manual for more information.

Example: Policy statement configuration output

```
config>router>policy-options# info
-----
prefix-list "loops"
  prefix 10.1.1.0/24 longer
exit
policy-statement "leak"
  entry 10
    from
      prefix-list "loop"
      level 2
    exit
    to
      level 1
    exit
    action accept
    exit
  exit
exit
policy-statement "isis-ext"
  entry 10
    from
      external
    exit
    to
      level 2
```

```
        exit
        action accept
        exit
    exit
exit
-----
config>router>policy-options#
```

4.10 IS-IS command reference

4.10.1 Command hierarchies

4.10.1.1 Configuration commands

- [Global commands](#)
- [Interface command](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

4.10.1.2 Global commands

```
config
- router
- [no] isis [isis-instance]
- [no] advertise-passive-only
- advertise-router-capability {area | as}
- no advertise-router-capability
- advertised-max-area-addr advertised-max-area-addr
- no advertised-max-area-addr
- all-l1isis ieee-address
- all-l2isis ieee-address
- [no] area-id area-address
- [no] authentication-check
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- authentication-type {password | message-digest}
- no authentication-type
- [no] csnp-authentication
- [no] default-route-tag tag
- [no] disable-ldp-sync
- export policy-name [... policy-name... up to 5 max)]
- no export
- export-limit number [log percentage]
- no export-limit
- [no] graceful-restart
- - [no] helper-disable
- [no] hello-authentication
- [no] iid-tlv-enable
- segment-routing
- no segment-routing
```

```
- prefix-sid-range {global | start-label label-value max-index index-value}  
- no prefix-sid-range  
- tunnel-mtu bytes  
- no tunnel-mtu  
- tunnel-table-pref preference  
- no tunnel-table-pref  
- [no] shutdown
```

4.10.1.3 Interface command

```
config  
- router  
  - [no] isis [isis-instance]  
    - [no] interface ip-int-name  
      - [no] bfd-enable ipv4  
      - csnp-interval seconds  
      - no csnp-interval  
      - hello-authentication-key [authentication-key | hash-key] [hash | hash2]  
      - no hello-authentication-key  
      - hello-authentication-type {password | message-digest}  
      - no hello-authentication-type  
      - interface-type {broadcast | point-to-point}  
      - no interface-type  
      - [no] loopfree-alternate-exclude  
      - level {1 | 2}  
        - hello-authentication-key [authentication-key | hash-key] [hash | hash2]  
        - no hello-authentication-key  
        - hello-authentication-type [password | message-digest]  
        - no hello-authentication-type  
        - hello-interval seconds  
        - no hello-interval  
        - hello-multiplier multiplier  
        - no hello-multiplier  
        - ipv6-unicast-metric ipv6 metric  
        - no ipv6-unicast-metric  
        - metric ipv4-metric  
        - no metric  
        - [no] passive  
        - priority number  
        - no priority  
      - level-capability {level-1 | level-2 | level-1/2}  
      - no lfa-policy-map  
      - lfa-policy-map route-nh-template template-name  
      - loopfree-alternate-exclude  
      - no loopfree-alternate-exclude  
      - lsp-pacing-interval milli-seconds  
      - no lsp-pacing-interval  
      - mesh-group [value | blocked]  
      - no mesh-group  
      - ipv4-node-sid index value  
      - ipv4-node-sid label value  
      - no ipv4-node-sid  
      - [no] passive  
      - retransmit-interval seconds  
      - no retransmit-interval  
      - [no] shutdown  
      - tag tag  
      - no tag  
    - [no] ipv4-routing  
    - [no] ipv6-routing {native | mt}  
    - loopfree-alternate [remote-lfa]
```

```

- loopfree-alternate remote-lfa [max-pq-cost value]
- no loopfree-alternate
- loopfree-alternate-exclude
- no loopfree-alternate-exclude
- level level
  - authentication-key [authentication-key | hash-key] [hash | hash2]
  - no authentication-key
  - authentication-type {password | message-digest}
  - no authentication-type
  - [no] csnp-authentication
  - [no] default-ipv4-unicast-metric ipv4 multicast metric
  - default-ipv4-unicast-metric ipv4 metric
  - no default-ipv4-unicast-metric
  - default-ipv6-unicast-metric ipv6 metric
  - no default-ipv6-unicast-metric
  - external-preference external-preference
  - no external-preference
  - [no] hello-authentication
  - [no] loopfree-alternate-exclude
  - preference preference
  - no preference
  - [no] psnp-authentication
  - [no] wide-metrics-only
- level-capability {level-1 | level-2 | level-1/2}
- lsp-lifetime seconds
- no lsp-lifetime
- loopfree-alternate
- no loopfree-alternate
- lsp-mtu-size size
- no lsp-mtu-size
- [no] lsp-wait lsp-wait [lsp-initial-wait [lsp-second-wait]]
- multi-topology
- no multi-topology
  - ipv6-unicast
  - no ipv6-unicast
- overload
- no overload
- overload-on-boot [timeout seconds]
- no overload-on-boot
- [no] psnp-authentication
- reference-bandwidth bandwidth-in-kbps
- reference-bandwidth [tbps Tera-bps] [gbps Giga-bps] [mbps Mega-bps] [kbps Kilo-
bps]
- no reference-bandwidth
- [no] shutdown
- [no] spf-wait spf-wait [spf-initial-wait [spf-second-wait]]
- [no] strict-adjacency-check
- [no] suppress-default
- summary-address {ip-prefix/mask | ip-prefix [netmask]} level
- no summary-address {ip-prefix/mask | ip-prefix [netmask]}
- [no] traffic-engineering

```

4.10.1.4 Show commands

```

show
- router
  - isis all
  - isis [isis-instance]
    - adjacency [ip-address | ip-int-name | nbr-system-id] [detail]
    - capabilities [system-id | lsp-id ] [level level]
    - database [system-id | lsp-id ] [detail] [level level]

```

```
- hostname
- interface [ip-int-name | ip-address] [detail]
- lfa-coverage
- link-group-member-status name [level level]
- link-group-status name [level level]
- prefix-sids [ipv4-unicast] [ip-prefix[/prefix-length]] [sid sid] [adv-
router system-id | hostname]
- routes [ipv4-unicast | ipv6-unicast | mt mt-id-number] [ip-prefix/prefix-length]
[alternative]
- spf-log [detail]
- statistics
- status
- summary-address [ip-prefix[/prefix-length]]
- topology [[ipv4-unicast | ipv6-unicast | mt mt-id-number][detail]]
```

4.10.1.5 Clear commands

```
clear
- router
  - isis [isis-instance]
    - adjacency [system-id]
    - database [system-id]
    - export
    - spf-log
    - statistics
```

4.10.1.6 Debug commands

```
debug
- router
  - isis [isis-instance]
    - [no] adjacency [ip-int-name | ip-address | nbr-system-id]
    - [no] cspf
    - [no] graceful-restart
    - interface [ip-int-name | ip-address]
    - no interface
    - leak [ip-address]
    - no leak
    - [no] lsdb [level-number] [system-id | lsp-id]
    - [no] misc
    - packet [packet-type] [ip-int-name | ip-address] [detail]
    - rtm [ip-address]
    - no rtm
```

4.10.2 Command descriptions

4.10.2.1 IS-IS configuration commands

4.10.2.1.1 Generic commands

```
isis
```

Syntax

```
isis [isis-instance]  
no isis [isis-instance]
```

Context

```
config>router
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure the Intermediate-System-to-Intermediate-System (IS-IS) protocol instance.

The IS-IS protocol instance is enabled with the **no shutdown** command in the **config>router>isis** context. Alternatively, the IS-IS protocol instance is disabled with the **shutdown** command in the **config>router>isis** context.

The **no** form of this command deletes the IS-IS protocol instance. Deleting the protocol instance removes all configuration parameters for this IS-IS instance.



Note:
The platforms as described in this document allow for the configuration of a single IS-IS instance at any time. The instance ID can be any number other than 0. This enables these platforms to be used in a network where multi-instance IS-IS is deployed, and the node needs to use an instance ID other than the default instance ID of 0.

Parameters

<i>isis-instance</i>	
Specifies the IS-IS instance.	
Values	0 to 31
Default	0

shutdown

Syntax

[no] shutdown

Context

config>router>isis

config>router>isis>interface

config>router>isis>if>level

config>router>isis>segment-routing

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity and entities contained within is disabled. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables an entity.

Default

no shutdown

Special Cases

IS-IS Global

In the **config>router>isis** context, the **shutdown** command disables the IS-IS protocol instance. By default, the protocol is enabled, **no shutdown**.

IS-IS Interface

In the **config>router>isis>interface** context, the command disables the IS-IS interface. By default, the IS-IS interface is enabled, **no shutdown**.

IS-IS Interface and Level

In the **config>router>isis>interface>level** context, the command disables the IS-IS interface for the level. By default, the IS-IS interface at the level is enabled, **no shutdown**.

advertise-passive-only

Syntax

[no] advertise-passive-only

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables and disables IS-IS to advertise only prefixes that belong to passive interfaces.

The **no** form of this command disables IS-IS to advertise only prefixes that belong to passive interfaces.

advertise-router-capability

Syntax

advertise-router-capability {area | as}

no advertise-router-capability

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables advertisement of the capabilities of a router to its neighbors for informational and troubleshooting purposes. A TLV, as defined in RFC 4971, advertises the TE Node Capability Descriptor capability.

The **area** and **as** keywords control the scope of the capability advertisements.

The **no** form of this command disables this advertisement capability.

Default

no advertise-router-capability

Parameters

area

Keyword specifying advertisement only within the area of origin.

as

Keyword specifying advertisement throughout the entire autonomous system.

advertised-max-area-addr

Syntax

advertised-max-area-addr *advertised-max-area-addr*
no advertised-max-area-addr

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the value for the maximum area addresses advertised in IS-IS messages.



Note:

- This command does not affect the number of area addresses allowed to be configured on the node.
- The value configured must be set to the same value on the adjacent IS-IS neighbor for the adjacency to be established successfully.

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

Default

advertised-max-area-addr 3

Parameters

advertised-max-area-addr

Specifies the maximum advertised area address.

Values 3 to 64

authentication-check

Syntax

[no] authentication-check

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets an authentication check to reject PDUs that do not match the type or key requirements.

The default behavior when authentication is configured is to reject all IS-IS protocol PDUs that have a mismatch in either the authentication type or authentication key.

When **no authentication-check** is configured, authentication PDUs are generated and IS-IS PDUs are authenticated on receipt. However, mismatches cause an event to be generated and will not be rejected.

The **no** form of this command allows authentication mismatches to be accepted and generate a log event.

Default

authentication-check

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

config>router>isis

config>router>isis>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets the authentication key used to verify PDUs sent by neighboring routers on the interface.

Neighboring routers use passwords to authenticate PDUs sent from an interface. For authentication to work, both the authentication key and the authentication type on a segment must match. The [authentication-type](#) statement must also be included.

To configure authentication on the global level, configure this command in the **config>router>isis** context. When this parameter is configured on the global level, all PDUs are authenticated including the hello PDU.

To override the global setting for a specific level, configure the **authentication-key** command in the **config>router>isis>level** context. When configured within the specific level, hello PDUs are not authenticated.

The **no** form of this command removes the authentication key.

Default

no authentication-key

Parameters

authentication-key

Specifies the authentication key. The key can be any combination of ASCII characters up to 255 characters (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 342 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" "). This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

authentication-type

Syntax

authentication-type {password | message-digest}

no authentication

Context

config>router>isis

config>router>isis>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables either simple password or message digest authentication or must go in either the global IS-IS or IS-IS level context.

Both the authentication key and the authentication type on a segment must match. The **authentication-key** statement must also be included.

Configure the authentication type on the global level in the **config>router>isis** context.

Configure or override the global setting by configuring the authentication type in the **config>router>isis>level** context.

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

Default

no authentication-type

Parameters

password

Specifies that simple password (plain text) authentication is required.

message-digest

Specifies that MD5 authentication in accordance with RFC2104 is required.

bfd-enable

Syntax

[no] bfd-enable ipv4

Context

config>router>isis>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the use of bidirectional forwarding (BFD) to control IPv4 adjacencies. By enabling BFD on an IPv4 protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set by the BFD command under the IP interface.

For more information about the protocols and platforms that support BFD, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*.

The **no** form of this command removes BFD from the associated adjacency.

Default

no bfd-enable ipv4

default-route-tag

Syntax

default-route-tag tag

no default-route-tag

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the route tag for the default route.

Parameters

<i>tag</i>	Specifies a default tag.
Values	Accepts decimal or hex formats: ISIS: [0x0..0xFFFFFFFF]H
Values	1 to 4294967295

csnp-authentication

Syntax

[no] csnp-authentication

Context

config>router>isis
config>router>isis>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables authentication of individual ISIS packets of the complete sequence number PDUs (CSNP) type.
The **no** form of this command suppresses authentication of CSNP packets.

csnp-interval

Syntax

csnp-interval *seconds*
no csnp-interval

Context

config>router>isis>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the time interval, in seconds, to send complete sequence number (CSN) PDUs from the interface. IS-IS must send CSN PDUs periodically.

By default, CSN PDUs are sent every 10 seconds for LAN interfaces and every 5 seconds for point-to-point interfaces

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

Default

csnp-interval 10

csnp-interval 5

Parameters

seconds

Specifies the time interval, in seconds, between successive CSN PDUs sent from this interface expressed as a decimal integer.

Values 1 to 65535

default-ipv4-unicast-metric

Syntax

default-ipv4-unicast-metric *metric*

no default-ipv4-unicast-metric

Context

config>router>isis>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the default metric used for IPv4 routes for both level 1 and level 2 on the interface, only when IS-IS multi-topology is configured for use.

To calculate the lowest cost to reach a specific destination, each configured level on each interface must have a cost. The costs for each level on an interface may be different. The value specified with this command is used only if the metric is not specified using the CLI command **ipv4-unicast-metric** under the specific level.

If the metric is not configured, the default of 10 is used unless reference bandwidth is configured.

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

Default

default-ipv4-unicast-metric 10

Parameters

metric

Specifies the metric assigned for this level on this interface.

Values 1 to 16777215

default-ipv6-unicast-metric

Syntax

default-ipv6-unicast-metric *ipv6 metric*

no default-ipv6-unicast-metric

Context

config>router>isis>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the default metric used for IPv6 routes for both level 1 and level 2 on the interface, only when IS-IS multi-topology is configured for use.

To calculate the lowest cost to reach a specific destination, each configured level on each interface must have a cost. The costs for each level on an interface may be different. The value specified with this command is used only if the metric is not specified using the command **ipv6-unicast-metric** under the specific level.

If the metric is not configured, the default of 10 is used unless reference bandwidth is configured.

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

Default

default-ipv6-unicast-metric 10

Parameters

ipv6 metric

The metric assigned for this level on this interface.

Values 1 to 16777215

default-metric

Syntax

default-metric *ipv4 metric*
no default-metric

Context

config>router>isis>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the configurable default metric used for all IS-IS interfaces on this level. This value is not used if a metric is configured for an interface.

Default

default-metric 10

Parameters

ipv4 metric

Specifies the default metric for IPv4 unicast.

Values 1 to 16777215

disable-ldp-sync

Syntax

[no] disable-ldp-sync

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command disables the IGP-LDP synchronization feature on all interfaces participating in the IS-IS routing protocol. When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces which have the IGP-LDP synchronization enabled if the currently advertised cost is different. It then disables IGP-LDP synchronization for all interfaces. This command does not

delete the interface configuration. The **no** form of this command has to be entered to re-enable IGP-LDP synchronization for this routing protocol.

For information about LDP synchronization, see “IGP-LDP and static route-LDP synchronization on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C” and the **ldp-sync** and **ldp-sync-timer** commands in the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*.

The **no** form of this command restores the default settings and re-enables IGP-LDP synchronization on all interfaces participating in the IS-IS routing protocol and for which the **ldp-sync-timer** is configured.

Default

no disable-ldp-sync

export

Syntax

[no] **export** *policy-name* [*policy-name*...up to 5 max]

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures export routing policies that determine the routes exported from the routing table to IS-IS.

If no export policy is defined, non IS-IS routes are not exported from the routing table manager to IS-IS.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered overrides the previous command. A maximum of five policy names can be specified.

If an **aggregate** command is also configured in the **config>router** context, the aggregation is applied before the export policy is applied.

Routing policies are created in the **config>router>policy-options** context.

The **no** form of this command removes the specified *policy-name* or all policies from the configuration if no *policy-name* is specified.

Default

no export

Parameters

policy-name

Specifies the export policy name, up to 32 characters. Up to five *policy-name* arguments can be specified.

export-limit

Syntax

export-limit *number* [**log percentage**]

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum number of routes (prefixes) that can be exported into IS-IS from the route table.

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

Default

no export-limit

Parameters

number

Specifies the maximum number of routes (prefixes) that can be exported into ISIS from the route table.

Values 1 to 4294967295

log percentage

Specifies the percentage of the export-limit when a warning log message and SNMP notification will be sent.

Values 1 to 100

external-preference

Syntax

external-preference *external-preference*

no external-preference

Context

config>router>isis>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the external route preference for the IS-IS level.

The **external-preference** command configures the preference level of either IS-IS level 1 or IS-IS level 2 external routes. By default, the preferences are as listed in [Table 44: Default route preferences](#).

A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference decides the route to use.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is dependent on the default preference table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of the route to use is determined by the configuration of **ecmp** in the **config>router** context.

Default

Default preferences are listed in the following table.

Table 44: Default route preferences

Route type	Preference	Configurable
Direct attached	0	No
Static-route	5	Yes
OSPF internal routes	10	No
IS-IS Level 1 internal	15	Yes ¹¹
IS-IS Level 2 internal	18	Yes*
OSPF external	150	Yes
IS-IS Level 1 external	160	Yes
IS-IS Level 2 external	165	Yes
BGP	170	Yes
BGP	170	Yes

Parameters

external-preference

Specifies the preference for external routes at this level as expressed.

¹¹ Internal preferences are changed using the **preference** command in the **config>router>isis>level** context.

Values 1 to 255

graceful-restart

Syntax

[no] graceful-restart

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables graceful-restart helper support for IS-IS. The router will act as a helper to neighbors who are graceful-restart-capable and are restarting.

When the control plane of a graceful-restart-capable router fails, the neighboring routers (graceful-restart helpers) temporarily preserve adjacency information so packets continue to be forwarded through the failed graceful-restart router using the last known routes. If the control plane of the graceful-restart router comes back up within the timer limits, the routing protocols reconverge to minimize service interruption.

The **no** form of this command disables graceful restart and removes all graceful restart configurations in the IS-IS instance.

Default

disabled

helper-disable

Syntax

[no] helper-disable

Context

config>router>isis>graceful-restart

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command disables the helper support for graceful restart.

When **graceful-restart** is enabled, the router can act as a helper router (the router is helping a neighbor to restart) or a restarting router or both. The router supports only helper mode. This facilitates the graceful

restart of neighbors but the router does not act as a restarting router (meaning that the router will not help the neighbors to restart).

The **no** form of this command enables helper support and is the default when **graceful-restart** is enabled.

Default

helper-disable

loopfree-alternate

Syntax

loopfree-alternate [**remote-lfa**]

loopfree-alternate remote-lfa [**max-pq-cost** *value*]

no loopfree-alternate

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the Loop-Free Alternate (LFA) computation by SPF for the IS-IS routing protocol instance.

The IGP SPF is instructed to precompute both a primary next-hop and an LFA next-hop for every learned prefix. When found, the LFA next-hop is populated into the routing table along with the primary next-hop for the prefix.

The IGP LFA SPF uses the **remote-lfa** option to enable the remote LFA next-hop calculation. When this option is enabled in an IGP instance, SPF performs the remote LFA additional computation following the regular LFA next-hop calculation when the latter results in no protection for one or more prefixes that are resolved to a specific interface.

Remote LFA extends the protection coverage of LFA-FRR to any topology by automatically computing and establishing or tearing down shortcut tunnels (repair tunnels) to a remote LFA node (PQ node). This puts the packets back into the shortest path without looping them to the node that forwarded them over the repair tunnel. A repair tunnel can be an RSVP LSP, an LDP-in-LDP tunnel, or a segment routing tunnel. The use of segment routing repair tunnels is restricted to the remote LFA node.

Unlike the regular LFA algorithm, which is per-prefix, the remote LFA algorithm is a per-link LFA SPF calculation. It provides protection to all destination prefixes that share the protected link by using the neighbor on the other side of the protected link as a proxy for those prefixes.

Default

no loopfree-alternate

Parameters

remote-lfa

Keyword to enable the remote LFA next-hop calculation by the IGP LFA SPF.

max-pq-lfa value

Specifies the maximum IGP cost from the router that is performing the remote LFA calculation to the candidate P or Q node.

Values 0 to 4294967295

loopfree-alternate-exclude

Syntax

[no] loopfree-alternate

Context

configure>router>isis>level

configure>router>isis>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command instructs IGP to exclude a specific interface or all interfaces participating in a specific IS-IS level or OSPF area from the SPF LFA computation. This reduces the LFA SPF calculation where it is not needed.

When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the preceding OSPF command can only be executed under the area in which the specified interface is primary and when enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

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

Default

no loopfree-alternate-exclude

hello-authentication

Syntax

[no] hello-authentication

Context

config>router>isis


```
config>router>isis>level
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables authentication of individual IS-IS hello packets.

The **no** form of this command suppresses authentication of hello packets.

iid-tlv-enable

Syntax

```
[no] iid-tlv-enable
```

Context

```
config>router>isis>graceful-restart
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies whether the Instance Identifier (IID) TLV has been enabled or disabled for this IS-IS instance.

hello-authentication-key

Syntax

```
hello-authentication-key [authentication-key | hash-key] [hash | hash2]
```

```
no hello-authentication-key
```

Context

```
config>router>isis>interface
```

```
config>router>isis>if>level
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the authentication key (password) for hello PDUs. Neighboring routers use the password to verify the authenticity of hello PDUs sent from this interface. Both the hello authentication key and hello authentication type on a segment must match. The **hello-authentication-type** must be specified.

To configure the hello authentication key in the interface context, use the **hello-authentication-key** command in the **config>router>isis>interface** context.

To configure or override the hello authentication key for a specific level, use the **hello-authentication-key** command in the **config>router>isis>interface>level** context.

If both IS-IS and hello authentication are configured, hello messages are validated using hello authentication. If only IS-IS authentication is configured, it will be used to authenticate all IS-IS protocol PDUs (including hello).

When the hello authentication key is configured in the **config>router>isis>interface** context, it applies to all levels configured for the interface.

The **no** form of this command removes the authentication-key from the configuration.

Default

no hello-authentication-key

Parameters

authentication-key

Specifies the hello authentication key (password). The key can be any combination of ASCII characters up to 254 characters (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 342 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

hello-authentication-type

Syntax

hello-authentication-type {password | message-digest}

no hello-authentication-type

Context

config>router>isis>interface

config>router>isis>if>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables hello authentication at either the interface or level context. Both the hello authentication key and the hello authentication type on a segment must match. The hello **authentication-key** statement must also be included.

To configure the hello authentication type at the interface context, use the **hello-authentication-type** command in the **config>router>isis>interface** context.

To configure or override the hello authentication setting for a specific level, configure the **hello-authentication-type** command in the **config>router>isis>interface>level** context.

The **no** form of this command disables hello authentication.

Default

no hello-authentication-type

Parameters

password

Specifies simple password (plain text) authentication is required.

message-digest

Specifies MD5 authentication (in accordance with RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*) is required.

hello-interval

Syntax

hello-interval *seconds*

no hello-interval

Context

config>router>isis>if>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the interval between IS-IS Hello PDUs issued on the interface at this level. The **hello-interval** command, along with the **hello-multiplier** command, is used to calculate a hold time, which is communicated to a neighbor in a Hello PDU.



Note:

The neighbor hold time is (hello multiplier × hello interval) on non-designated intermediate system broadcast interfaces and point-to-point interfaces and is (hello multiplier × hello interval / 3) on

designated intermediate system broadcast interfaces. Hello values can be adjusted for faster convergence, but the hold time should always be > 3 to reduce routing instability.

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

Default

hello-interval 3 — Hello interval default for the designated intersystem.

hello-interval 9 — Hello interval default for non-designated intersystems.

Parameters

seconds

Specifies the Hello interval in seconds expressed as a decimal integer.

Values 1 to 20000

hello-multiplier

Syntax

hello-multiplier *multiplier*

no hello-multiplier

Context

config>router>isis>if>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a hello multiplier. The **hello-multiplier** command, along with the **hello-interval** command, is used to calculate a hold time, which is communicated to a neighbor in a Hello PDU.

The hold time is the time during which the neighbor expects to receive the next Hello PDU. If the neighbor receives a Hello within this time, the hold time is reset. If the neighbor does not receive a Hello within the hold time, it brings the adjacency down.



Note:

The neighbor hold time is (hello multiplier × hello interval) on non-designated intermediate system broadcast interfaces and point-to-point interfaces and is (hello multiplier × hello interval / 3) on designated intermediate system broadcast interfaces. Hello values can be adjusted for faster convergence, but the hold time should always be greater than three to reduce routing instability.

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

Default

hello-multiplier 3

Parameters

multiplier

Specifies the multiplier for the hello interval expressed as a decimal integer.

Values 2 to 100

ipv6-unicast-metric

Syntax

ipv6-unicast-metric *ipv6 metric*

ipv6-unicast-metric

Context

config>router>isis>if>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the default metric used for IPv6 routes for both level 1 and level 2 on the interface, only when IS-IS multi-topology is configured for use.

To calculate the lowest cost to reach a specific destination, each configured level on each interface must have a cost. The costs for each level on an interface may be different.

If the metric is not configured, the default of 10 is used unless reference bandwidth is configured.

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

Default

ipv6-unicast-metric 10

Parameters

ipv6 metric

Specifies the metric assigned for this level on this interface.

Values 1 to 16777215

interface

Syntax

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

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an IS-IS interface.

When an area is defined, the interfaces belong to that area. Interfaces cannot belong to separate areas.

When the interface is a POS channel, the OSINCP is enabled when the interface is created and removed when the interface is deleted.

The **shutdown** command in the **config>router>isis>interface** context administratively disables IS-IS on the interface without affecting the IS-IS configuration.

The **no** form of this command removes IS-IS from the interface.

Default

no interface

Parameters

ip-int-name

Specifies the IP interface name created in the **config>router>interface** context. The IP interface name must already exist.

tag

Syntax

tag *tag*

no tag

Context

config>router>isis>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a route tag to the specified IP address of an interface.

Parameters

tag

Specifies the route tag number.

Values 1 to 4294967295

interface-type

Syntax

```
interface-type {broadcast | point-to-point}  
no interface-type
```

Context

```
config>router>isis>interface
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the default metric used for IPv6 routes for both level 1 and level 2 on the interface, only when IS-IS multi-topology is configured for use.

Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the designated IS-IS overhead if the link is used as a point-to-point.

If the interface type is not known at the time the interface is added to IS-IS, and subsequently the IP interface is bound (or moved) to a different interface type, this command must be entered manually.

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

Default

```
interface-type point-to-point  
interface-type broadcast
```

Special Cases

SONET

Interfaces on SONET channels default to the point-to-point type.

Ethernet or Unknown

Physical interfaces that are Ethernet or unknown default to the broadcast type.

Parameters

broadcast

Specifies to maintain this link as a broadcast network.

point-to-point

Specifies to maintain this link as a point-to-point link.

ipv4-routing

Syntax

[no] ipv4-routing

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies whether this IS-IS instance supports IPv4.

The **no** form of this command disables IPv4 on the IS-IS instance.

Default

ipv4-routing

ipv6-routing

Syntax

[no] ipv6-routing {native | mt}

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables IPv6 routing.

The **no** form of this command disables support for IS-IS IPv6 TLVs for IPv6 routing.

Default

no ipv6-routing

Parameters

native

Specifies to enable IS-IS IPv6 TLVs for IPv6 routing and enables support for native IPv6 TLV.

mt

Specifies to enable IS-IS multi-topology TLVs for IPv6 routing. When this parameter is specified, the support for native IPv6 TLVs is disabled.

level

Syntax

level *level-number*

Context

config>router>isis

config>router>isis>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure IS-IS Level 1 or Level 2 area attributes.

A router can be configured as a Level 1, Level 2, or Level 1-2 system. A Level 1 adjacency can be established if there is at least one area address shared by this router and a neighbor. A Level 2 adjacency cannot be established over this interface.

Level 1-2 adjacency is created if the neighbor is also configured as a Level 1-2 router and has at least one area address in common. A Level 2 adjacency is established if there are no common area IDs.

A Level 2 adjacency is established if another router is configured as Level 2 or a Level 1-2 router with interfaces configured as Level 1-2 or Level 2. Level 1 adjacencies are not established over this interface.

To reset global and interface level parameters to the default, the following commands must be entered independently:

- **level>no hello-authentication-key**
- **level>no hello-authentication-type**
- **level>no hello-interval**
- **level>no hello-multiplier**
- **level>no metric**
- **level>no passive**
- **level>no priority**

Default

level 1 or level 2

Special cases

Global IS-IS Level

The **config>router>isis** context configures default global parameters for both Level 1 and Level 2 interfaces.

IS-IS Interface Level

The **config>router>isis>interface** context configures IS-IS operational characteristics of the interface at Level 1 and Level 2. A logical interface can be configured on one Level 1 and one Level 2. In this case, each level can be configured independently and parameters must be removed independently.

By default, an interface operates in both Level 1 and Level 2 modes.

Parameters

level-number

Specifies the IS-IS level number.

Values 1, 2

level-capability

Syntax

level-capability {level-1 | level-2 | level-1/2}

no level-capability

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the routing level for an instance of the IS-IS routing process.
An IS-IS router and an IS-IS interface can operate at Level 1, Level 2 or both Level 1 and 2.
The following table displays configuration combinations and the potential adjacencies that can be formed.

Table 45: Potential adjacency capabilities

Global level	Interface level	Potential adjacency
L 1/2	L 1/2	Level 1 and/or Level 2
L 1/2	L 1	Level 1 only
L 1/2	L 2	Level 2 only

Global level	Interface level	Potential adjacency
L 2	L 1/2	Level 2 only
L 2	L 2	Level 2 only
L 2	L 1	none
L 1	L 1/2	Level 1 only
L 1	L 2	none
L 1	L 1	Level 1 only

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

Default

level-capability level-1/2

Special cases

IS-IS Router

In the **config>router>isis** context, changing the **level-capability** performs a restart on the IS-IS protocol instance.

IS-IS Interface

In the **config>router>isis>interface** context, changing the **level-capability** performs a restart of IS-IS on the interface.

Parameters

level-1

Specifies that the router/interface can operate at Level 1 only.

level-2

Specifies that the router/interface can operate at Level 2 only.

level-1/2

Specifies that the router/interface can operate at both Level 1 and Level 2.

lsp-pacing-interval

Syntax

lsp-pacing-interval *milliseconds*

no lsp-pacing-interval

Context

config>router>isis>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the interval between LSP PDUs sent from this interface.

To avoid bombarding adjacent neighbors with excessive data, pace the Link State Protocol Data Units (LSPs). If a value of zero is configured, no LSPs are sent from the interface.

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

Default

lsp-pacing-interval 100

Parameters

milliseconds

Specifies the interval in milliseconds that IS-IS LSPs can be sent from the interface, expressed as a decimal integer.

Values 0 to 65535

lsp-lifetime

Syntax

lsp-lifetime *seconds*

no lsp-lifetime

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets the time interval, in seconds, for LSPs originated by the router to be considered valid by other router in the domain.

Each LSP received is maintained in an LSP database until the **lsp-lifetime** expires, unless the originating router refreshes the LSP. By default, each router refreshes its LSPs every 20 minutes (1200 seconds) so other routers will not age out the LSP.

The LSP refresh timer is derived from the following formula:

$\text{lsp-lifetime}/2$

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

Default

lsp-lifetime 1200

Parameters

seconds

Specifies the interval, for LSPs originated by the route to be considered valid by other routers in the domain.

Values 350 to 6553

lsp-mtu-size

Syntax

lsp-mtu-size *size*

no lsp-mtu-size

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the LSP MTU size. If the *size* value is changed from the default using CLI or SNMP, IS-IS must be restarted for the change to take effect. This can be done by performing a **shutdown** command and a **no shutdown** command in the **config>router>isis** context.



Note:

Using the **exec** command to execute a configuration file to change the LSP MTU size from the default value will automatically bounce IS-IS for the change to take effect.

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

Default

lsp-mtu-size 1492

Parameters

size

Specifies the LSP MTU size.

Values 490 to 9190

lsp-wait

Syntax

lsp-wait *lsp-wait* [*lsp-initial-wait* [*lsp-second-wait*]]

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command is used to customize IS-IS LSP generation throttling. Timers that determine when to generate the first, second, and subsequent LSPs can be controlled with this command. Subsequent LSPs are generated at increasing intervals of the second **lsp-wait** timer until a maximum value is reached.

Parameters

lsp-max-wait

Specifies the maximum interval in seconds between two consecutive occurrences of an LSP being generated.

Values 1 to 120

Default 5

lsp-initial-wait

Specifies the initial LSP generation delay in seconds.

Values 0 to 100

Default 0

lsp-second-wait

Specifies the hold time in seconds between the first and second LSP generation.

Values 1 to 100

Default 1

mesh-group

Syntax

mesh-group {*value* | **blocked**}

no mesh-group

Context

config>router>isis>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns an interface to a mesh group. Mesh groups limit the amount of flooding that occurs when a new or changed LSP is advertised throughout an area.

All routers in a mesh group should be fully meshed. When LSPs need to be flooded, only a single copy is received instead of a copy per neighbor.

To create a mesh group, configure the same mesh group value for each interface that is part of the mesh group. All routers must have the same mesh group value configured for all interfaces that are part of the mesh group.

To prevent an interface from flooding LSPs, the optional **blocked** parameter can be specified. Configure mesh groups carefully. It is easy to create isolated islands that do not receive updates as (other) links fail.

The **no** form of this command removes the interface from the mesh group.

Default

no mesh-group

Parameters

value

Specifies the unique decimal integer value that distinguishes this mesh group from other mesh groups on any router that is part of this mesh group.

Values 1 to 2000000000

blocked

Keyword to prevent an interface from flooding LSPs.

multi-topology

Syntax

[no] multi-topology

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables IS-IS multi-topology support.

The **no** form of this command disables IS-IS multi-topology support.

Default

no multi-topology

ipv4-node-sid

Syntax

ipv4-node-sid *index value*

ipv4-node-sid *label value*

no ipv4-node-sid

Context

config>router>isis>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns a node SID index or label value to the prefix representing the primary address of an IPv4 network interface of type loopback. Only a single node SID can be assigned to an interface. The secondary address of an IPv4 interface cannot be assigned a node SID index and does not inherit the SID of the primary IPv4 address.

This command fails if the network interface is not of type loopback or if the interface is defined in an IES or a VPRN context. Also, assigning an identical SID index or label value to the same interface in two different IGP instances is not allowed within the same node.

The value of the label or index SID is extracted from the range configured for this IGP instance. When the global mode of operation is used, a new segment routing module checks that the same index or label value is not assigned to more than one loopback interface address. When the per-instance mode of operation is used, this check is not required because the index and label ranges of the various IGP instances are not allowed to overlap.

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

Default

no ipv4-node-sid

Parameters

index *value*

Specifies the IPv4 SID node index value.

Values 0 to 4294967295

label *value*

Specifies the IPv4 SID node label value.

Values 0 to 4294967295

ipv6-unicast

Syntax

[no] ipv6-unicast

Context

config>router>isis>multi-topology

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables multi-topology TLVs.

This **no** form of this command disables multi-topology TLVs.

Default

no ipv6-unicast

metric

Syntax

metric *ipv4-metric*

no metric

Context

config>router>isis>if>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the metric used for the level on the interface.

To calculate the lowest cost to reach a specific destination, each configured level on each interface must have a cost. The costs for each level on an interface may be different.

If the metric is not configured, the default of 10 is used, unless reference bandwidth is configured.

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

Default

metric 10

Parameters

ipv4-metric

Specifies the metric assigned for this level on this interface.

Values 1 to 16777215

all-l1isis

Syntax

[no] **all-l1isis** *ieee-address*

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the MAC address to use for all L1 IS-IS routers. The MAC address should be a multicast address. The user should configure **shutdown** and **no shutdown** in the IS-IS instance to make the change operational.

The MAC address, 01-80-C2-00-02-11, is used in the IS-IS base instance ID (ID==0). This cannot be modified by the user.

Default

no all-l1isis

Parameters

ieee-address

Specifies the destination MAC address for all L1 IS-IS neighbors on the link for this IS-IS instance.

all-l2isis

Syntax

[no] **all-l2isis** *ieee-address*

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the MAC address to use for all L2 IS-IS routers. The MAC address should be a multicast address. The user should configure **shutdown** and **no shutdown** in the IS-IS instance to make the change operational.

The MAC address, 01-80-C2-00-01-00, is used in the IS-IS base instance ID (ID==0). This cannot be modified by the user.

Default

no all-l2isis

Parameters

ieee-address

Specifies the destination MAC address for all L2 IS-IS neighbors on the link for this ISIS instance.

area-id

Syntax

[no] **area-id** *area-address*

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command was previously named the **net** *network-entity-title* command. The **area-id** command enables the context to configure the area ID portion of Network Service Access Point (NSAP) addresses, which identify a point of connection to the network, such as a router interface. Addresses in the IS-IS protocol are based on the ISO NSAP addresses and Network Entity Titles (NETs), not IP addresses.

A maximum of 3 area addresses can be configured.

NSAP addresses are divided into the three following parts; only the area ID portion is configurable:

- **Area ID**

A variable length field between 1 and 13 bytes. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.

- **System ID**

A six-byte system identification. This value is not configurable. The system ID is derived from the system or router ID.

- **Selector ID**

A one-byte selector identification that must contain zeros when configuring a NET. This value is not configurable. The selector ID is always 00.

The NET is constructed like an NSAP but the selector byte contains a 00 value. NET addresses are exchanged in Hello and LSP PDUs. All NET addresses configured on the node are advertised to its neighbors.

For Level 1 interfaces, neighbors can have different area IDs, but they must have at least one area ID (AFI + area) in common. Because they share a common area ID, they become neighbors and area merging between the potentially different areas can occur.

For Level 2 (only) interfaces, neighbors can have different area IDs. However, if they have no area IDs in common, they become only Level 2 neighbors and Level 2 LSPs are exchanged.

For Level 1 and Level 2 interfaces, neighbors can have different area IDs. If they have at least one area ID (AFI + area) in common, they become neighbors. In addition to exchanging Level 2 LSPs, area merging between potentially different areas can occur.

If multiple **area-id** commands are entered, the system ID of all subsequent entries must match the first **area** address.

The **no** form of this command removes the area address.

Parameters

area-address

Specifies the 1 to 13-byte address. Of the total 20 bytes comprising the NET, only the first 13 bytes can be manually configured. As few as one byte can be entered up to a maximum of 13 bytes. If less than 13 bytes are entered, the rest is padded with zeros.

lfa-policy-map

Syntax

lfa-policy-map route-nh-template *template-name*

no lfa-policy-map

Context

config>router>isis>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command applies a route next-hop policy template to an OSPF or IS-IS interface.

When a route next-hop policy template is applied to an interface in IS-IS, it is applied in both Level 1 and Level 2.

However, the command in an OSPF interface context can only be executed under the area in which the specified interface is primary, and then applied in that area and in all other areas where the interface is

secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

If the user excluded the interface from LFA using the command **loopfree-alternate-exclude**, the LFA policy, if applied to the interface, has no effect.

Finally, if the user applied a route next-hop policy template to a loopback interface or to the system interface, the command will not be rejected, but it will result in no action being taken.

The **no** form deletes the mapping of a route next-hop policy template to an OSPF or IS-IS interface.

Parameters

template-name

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

loopfree-alternate-exclude

Syntax

[no] loopfree-alternate-exclude

Context

config>router>isis>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command instructs IGP to exclude a specific interface or all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.

When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both Level 1 and Level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the preceding OSPF command can only be executed under the area in which the specified interface is primary and when enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

The **no** form of this command reinstates the default value for this command.

Default

no loopfree-alternate-exclude

overload

Syntax

overload [timeout seconds]

no overload

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command administratively sets the IS-IS router to operate in the overload state for a specific time period, in seconds, or indefinitely.

During normal operation, the router may be forced to enter an overload state because of a lack of resources. When in the overload state, the router is only used if the destination is reachable by the router and will not be used for other transit traffic.

If a time period is specified, the overload state persists for the configured length of time. If no time is specified, the overload state operation is maintained indefinitely.

The **overload** command can be useful in circumstances where the router is overloaded or used before executing a **shutdown** command to divert traffic around the router.

The **no** form of this command causes the router to exit the overload state.

Default

no overload

Parameters

seconds

Specifies the time, in seconds, that this router must operate in the overload state.

Values 60 to 1800

Default infinity (overload state maintained indefinitely)

overload-on-boot

Syntax

overload-on-boot [timeoutseconds]

no overload-on-boot

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occurs.

- The timeout timer expires.
- A manual override of the current overload state is entered with the **config>router>isis>no overload** command.

The **no overload** command does not affect the **overload-on-boot** function.

If no timeout is specified, IS-IS will go into overload indefinitely after a reboot. After the reboot, the IS-IS status will display a permanent overload state:

- L1 LSDB Overload: Manual on boot (Indefinitely in overload)
- L2 LSDB Overload: Manual on boot (Indefinitely in overload)

This state can be cleared with the **config>router>isis>no overload** command.

If a timeout value is specified, IS-IS will go into the overload state for the configured timeout after a reboot. After the reboot, the IS-IS status will display the remaining time the system stays in overload:

- L1 LSDB Overload: Manual on boot (Overload Time Left: 17)
- L2 LSDB Overload: Manual on boot (Overload Time Left: 17)

The overload state can be cleared before the timeout expires with the **no overload** command. Use the **show router isis status** commands to display the administrative and operational state as well as all timers.

The **no** form of this command removes the overload-on-boot functionality from the configuration.

Default

no overload-on-boot

Parameters

timeout seconds

Specifies the number of seconds that the router remains in the overload state after rebooting.

Values 60 to 1800

passive

Syntax

[no] passive

Context

config>router>isis>interface

config>router>isis>if>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds the passive attribute to the IS-IS interface, which causes the interface to be advertised as an IS-IS interface without running the IS-IS protocol. Normally, only interface addresses that are configured for IS-IS are advertised as IS-IS interfaces at the level that they are configured.

When the passive mode is enabled, the interface or the interface at the specified level ignores ingress IS-IS protocol PDUs and will not transmit IS-IS protocol PDUs.

The **no** form of this command removes the passive attribute.

Default

passive

no passive

Special Cases

Service Interfaces

Service interfaces (defined using the service-prefix command in **config>router**) are passive by default.

All other Interfaces

All other interfaces are not passive by default.

preference

Syntax

preference *preference*

no preference

Context

config>router>isis>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the preference level of either IS-IS Level 1 or IS-IS Level 2 internal routes. The default preferences are listed in the default values section.

A route can be learned by the router by different protocols, in which case the costs are not comparable. When this occurs, the preference is used to decide the route that will be used by the router.

Protocols should not be configured with the same preference. If this occurs, the default preferences defined in the following table are used as the tiebreaker. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with

an identical preference using the same protocol and the costs (metrics) are equal, the route to use is determined by the configuration of the **ecmp** command in the **config>router** context.

Default

The following table lists default preferences for route types.

Table 46: Default preferences

Route type	Preference	Configurable
Direct attached	0	No
Static-route	5	Yes
OSPF internal routes	10	No
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes ¹²
IS-IS level 2 external	165	Yes ¹³
BGP	170	Yes

Parameters

preference

Specifies the preference for external routes at this level expressed as a decimal integer.

Values 1 to 255

priority

Syntax

priority *number*
no priority

Context

config>router>isis>if>level

¹² External preferences are changed using the **external-preference** command in the config>router>isis>level *level-number* context.
¹³ Internal preferences are changed using the **preference** command in the **config>router>isis>level** context.

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the priority of the IS-IS router interface for designated router election on a multi-access network.

The priority is included in Hello PDUs transmitted by the interface on a multi-access network. The router with the highest priority is the preferred designated router. The designated router is responsible for sending LSPs with regard to this network and the routers that are attached to it.

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

Default

priority 64

Parameters

number

Specifies the priority for this interface at this level.

Values 0 to 127

psnp-authentication

Syntax

[no] psnp-authentication

Context

config>router>isis

config>router>isis>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables authentication of individual IS-IS packets of partial sequence number PDU (PSNP) type.

The **no** form of this command suppresses authentication of PSNP packets.

reference-bandwidth

Syntax

reference-bandwidth *bandwidth-in-kbps*

reference-bandwidth [**tbps** *Tera-bps*] [**gbps** *Giga-bps*] [**mbps** *Mega-bps*] [**kbps** *Kilo-bps*]
no reference-bandwidth

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the reference bandwidth that provides the basis of bandwidth relative costing. To calculate the lowest cost to reach a specific destination, each configured level on each interface must have a cost. If the reference bandwidth is defined, the cost is calculated using the following formula:

cost = reference-bandwidth # bandwidth

If the reference bandwidth is configured as 10 Gigabits (10,000,000,000), a 100 Mb/s interface has a default metric of 100. For metrics in excess of 63 to be configured, wide metrics must be deployed. See [wide-metrics-only](#) for more information.

If the reference bandwidth is not configured, all interfaces have a default metric of 10.

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

Default

no reference-bandwidth

Parameters

bandwidth-in-kbps

Specifies the reference bandwidth in kilobits per second, expressed as a decimal integer.

Values 1 to 4000000000

Tera-bps

Specifies the reference bandwidth in terabits per second, expressed as a decimal integer.

Values 1 to 4

Giga-bps

Specifies the reference bandwidth in gigabits per second, expressed as a decimal integer.

Values 1 to 999

Mega-bps

Specifies the reference bandwidth in megabits per second, expressed as a decimal integer.

Values 1 to 999

Kilo-bps

Specifies the reference bandwidth in kilobits per second, expressed as a decimal integer.

Values 1 to 999

retransmit-interval

Syntax

retransmit-interval *seconds*

no retransmit-interval

Context

config>router>isis>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the minimum time between LSP PDU retransmissions on a point-to-point interface.

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

Default

retransmit-interval 100

Parameters

seconds

Specifies the interval, in seconds, that IS-IS LSPs can be sent on the interface.

Values 1 to 65535

segment-routing

Syntax

segment-routing

no segment-routing

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure segment routing parameters within an IGP instance.

Segment routing adds to IS-IS routing protocols the ability to perform shortest path routing and source routing using the concept of abstract segment. A segment can represent a local prefix of a node, a specific adjacency of the node (interface/next-hop), a service context, or a specific explicit path over the network. For each segment, the IGP advertises a segment identifier (SID).

When segment routing is used with the MPLS data plane, the SID is used as a standard MPLS label. A router forwarding a packet using segment routing pushes one or more MPLS labels.

Segment routing using MPLS labels is used in both shortest path routing applications and in traffic engineering applications. The commands in the **segment-routing** context configure the shortest path forwarding application.

After segment routing is configured in the IS-IS instance, the router will perform the following operations.

1. Advertise the segment routing capability sub-TLV to routers in all areas and levels of this IGP instance. However, only neighbors with which it established an adjacency will interpret the SID/label range information and use it for calculating the label to swap to or push for a given resolved prefix SID.
2. Advertise the assigned index for each configured node SID in the new prefix SID sub-TLV with the N-flag (node-SID flag) set. The segment routing module then programs the incoming label map (ILM) with a pop operation for each local node SID in the datapath.
3. Assign and advertise automatically an adjacency SID label for each formed adjacency over a network IP interface in the new adjacency SID sub-TLV. The segment routing module programs the incoming label map (ILM) with a pop operation, with a swap to an implicit null label operation, for each advertised adjacency SID.
4. Resolve received prefixes, and if a prefix SID sub-TLV exists, the Segment Routing module programs the ILM with a swap operation and also an LTN with a push operation both pointing to the primary/LFA NHLFE. An SR tunnel is also added to the TTM.

When the user enables segment routing in an IGP instance, the main SPF and LFA SPF are computed and the primary next-hop and LFA backup next-hop for a received prefix are added to the RTM without the label information advertised in the prefix SID sub-TLV.

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

prefix-sid-range

Syntax

prefix-sid-range {*global* | **start-label** *label-value* **max-index** *index-value*}

no prefix-sid-range

Context

config>router>isis>segment-routing

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the prefix SID index range and offset label value for an IGP instance.

The user must configure the prefix SID index range and the offset label value that this IGP instance uses. Because each prefix SID represents a network global IP address, the SID index for a prefix must be unique in the network. Therefore, all routers in the network configure and advertise the same prefix SID index range for an IGP instance. However, the label value used by each router to represent this prefix, that is, the label programmed in the ILM, can be local to that router by the use of an offset label, referred to as a start label, as in the following:

$$\text{Local Label (Prefix SID)} = \text{start-label} + \{\text{SID index}\}$$

The label operation in the network becomes similar to LDP when operating in the independent label distribution mode (RFC 5036), with the difference that the label value used to forward a packet to each downstream router is computed by the upstream router based on the advertised prefix SID index using the above formula.

There are two mutually exclusive modes of operation for the prefix SID range on the router.

In the global mode of operation, the global value is configured and this IGP instance assumes that the start label value is the lowest label value in the SRGB, and the prefix SID index range size is equal to the range size of the SRGB. When one IGP instance selects the global option for the prefix SID range, all IGP instances on the system are restricted to do the same. The user must shut down the segment routing context and delete the **prefix-sid-range** command in all IGP instances to change the SRGB. After the SRGB is changed, the user must re-enter the **prefix-sid-range** command. The SRGB range change fails if an already allocated SID index or label goes out of range.

In the per-instance mode of operation, the user partitions the SRGB into non-overlapping sub-ranges among the IGP instances. The user therefore configures a subset of the SRGB by specifying the start label value and the prefix SID index range size. All resulting net label values (start-label + index) must be within the SRGB or the configuration will be failed.

Furthermore, the code checks for overlaps of the resulting net label value range across IGP instances and strictly enforces that these ranges do not overlap. The user must shut down the segment routing context of an IGP instance to change the SID index or label range of that IGP instance using the **prefix-sid-range** command.

In addition, any range change will fail if an already allocated SID index or label goes out of range. The user can, however, change the SRGB on the fly as long as it does not reduce the current per-IGP instance SID index or label range defined in the **prefix-sid-range** command. Otherwise, the user must shut down the segment routing context of the IGP instance and delete and reconfigure the **prefix-sid-range** command.

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

Default

no prefix-sid-range

Parameters

start-label *label-value*

Specifies the label offset for the SR label range of this IGP instance.

Values 0 to 524287

max-index *index-value*

Specifies the maximum value of the prefix SID index range for this IGP instance.

Values 1 to 524287

tunnel-mtu

Syntax

tunnel-mtu *bytes*

no tunnel-mtu

Context

config>router>isis>segment-routing

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the MTU of all SR tunnels within each IGP instance.

The MTU of an SR tunnel populated into the TTM is determined in the same way as for an IGP tunnel; for example, LDP LSP, based on the outgoing interface MTU minus the label stack size. Remote LFA can add at least two more labels to the tunnel for a total of three labels. There is no default value. If the user does not configure an SR tunnel MTU, the MTU will be determined by IGP.

The MTU of the SR tunnel in bytes is determined as follows:

$$\text{SR_Tunnel_MTU} = \text{MIN} \{ \text{Cfg_SR_MTU}, \text{IGP_Tunnel_MTU} - (1 + \text{frr-overhead}) * 4 \}$$

Where:

- Cfg_SR_MTU is the MTU configured by the user for all SR tunnels within a specific IGP instance using this CLI command. If no value is configured by the user, the SR tunnel MTU will be determined by the *IGP_Tunnel_MTU* calculated value.
- IGP_Tunnel_MTU is the minimum of the IS-IS or OSPF interface MTU among all the ECMP paths or among the primary and LFA backup paths of this SR tunnel.
- frr-overhead is set to 1 if **segment-routing** and **remote-lfa** options are enabled in the IGP instance. Otherwise, it is set to 0.

The SR tunnel MTU is dynamically updated when any of the preceding parameters that are used in its calculation change. This includes when the set of the tunnel next-hops changes, or the user changes the configured SR MTU or interface MTU value.

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

Default

no tunnel-mtu

Parameters

bytes

Specifies the size of the MTU in bytes.

Values 512 to 9198

tunnel-table-pref

Syntax

tunnel-table-pref *preference*

no tunnel-table-pref

Context

config>router>isis>segment-routing

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the TTM preference of the shortest path SR tunnels created by the IGP instance. The TTM preference is used in the case of VPRN auto-bind or BGP transport tunnels when the new tunnel binding commands are configured to the **any** value, which parses the TTM for tunnels in the protocol preference order. The user can either use the global TTM preference or list the tunnel types they want to use. When they list the tunnel types, the TTM preference is used to select one type over the other. In both cases, a fallback to the next preferred tunnel type is performed if the selected one fails. A reversion to a more preferred tunnel type is performed as soon as one is available.

The segment routing module adds an SR tunnel entry to the TTM for each resolved remote node SID prefix and programs the datapath that has the corresponding LTN with the push operation pointing to the primary and LFA backup NHLFEs.

The default preference for shortest path SR tunnels in the TTM is set lower than LDP tunnels but higher than BGP tunnels to allow controlled migration of customers without disrupting their current deployment when they enable segment routing. The following is the setting of the default preference of the various tunnel types. This includes the preference of SR tunnels based on the shortest path (referred to as SR-ISIS).

The global default TTM preference for the tunnel types is as follows:

- ROUTE_PREF_RSVP 7
- ROUTE_PREF_SR_TE 8
- ROUTE_PREF_LDP 9
- ROUTE_PREF_OSPF_TTM 10
- ROUTE_PREF_ISIS_TTM 11
- ROUTE_PREF_BGP_TTM 12
- ROUTE_PREF_GRE 255

The default value for SR-ISIS is the same regardless of whether one or more IS-IS instances programmed a tunnel for the same prefix. The selection of an SR tunnel in this case will be based on the lowest IGP instance ID.

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

Default

no tunnel-table-pref

Parameters

preference

Specifies an integer value that represents the preference of IS-IS SR tunnels in the TTM.

Values 1 to 255

Default 11

spf-wait

Syntax

[no] **spf-wait** *spf-wait* [*spf-initial-wait* [*spf-second-wait*]]

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum interval between two consecutive SPF calculations in seconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, the next SPF will run after 2000 milliseconds, and then the next SPF will run after 4000 milliseconds, and so on, until it reaches the *spf-wait* value.

The SPF interval stays at the *spf-wait* value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval drops back to *spf-initial-wait*.

Default

no spf-wait

Parameters

spf-wait

Specifies the maximum interval in seconds between two consecutive SPF calculations.

Values 1 to 120

Default 10

spf-initial-wait

Specifies the initial SPF calculation delay in milliseconds after a topology change.

Values 10 to 100000

Default 1000

spf-second-wait

Specifies the hold time in milliseconds between the first and second SPF calculation.

Values 1 to 100000

Default 1000

strict-adjacency-check

Syntax

[no] **strict-adjacency-check**

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables strict checking of address families (IPv4 and IPv6) for IS-IS adjacencies. When enabled, adjacencies will not come up unless both routers have exactly the same address families configured. If there is an existing adjacency with unmatched address families, it will be torn down. This command is used to prevent black-holing traffic when IPv4 and IPv6 topologies are different. When disabled (no strict-adjacency-check) a BFD session failure for either IPv4 or Ipv6 will cause the routes for the other address family to be removed as well.

The **no** form of this command disables the strict checking of address families. When strict checking of address families is disabled, both routers only need to have one common address family to establish the adjacency.

Default

no strict-adjacency-check

summary-address

Syntax

```
summary-address {ip-prefix/mask | ip-prefix [netmask]} level [tag tag]
no summary-address {ip-prefix/mask | ip-prefix [netmask]}
```

Context

```
config>router>isis
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates summary addresses.

Parameters

ip-prefix/mask

Specifies information for the specified IP prefix and mask length.

Values	
ipv4-address:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 — 32
ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x: [0 — FFFF]H d: [0 — 255]D
ipv6-prefix-length:	[0 — 128]

netmask

Specifies the subnet mask in dotted-decimal notation.

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

level

Specifies IS-IS level area attributes.

Values	level-1, level-2, level-1/2
--------	-----------------------------

tag tag

Assigns an OSPF, RIP or ISIS tag to routes matching the entry.

Values	Accepts decimal or hex formats: [1..4294967295]
--------	---

OSPF and ISIS: [0x0..0xFFFFFFFF]H

suppress-default

Syntax

[no] suppress-default

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables IS-IS to suppress the installation of default routes.

traffic-engineering

Syntax

[no] traffic-engineering

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures traffic engineering and determines if IGP shortcuts are required.

Default

no traffic-engineering

wide-metrics-only

Syntax

[no] wide-metrics-only

Context

config>router>isis>level

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the exclusive use of wide metrics in the LSPs for the level number. Narrow metrics can have values between 1 and 63. IS-IS can generate two TLVs, one for the adjacency and one for the IP prefix. To support traffic engineering, wider metrics are required. When wide metrics are used, a second pair of TLVs are added, again, one for the adjacency and one for the IP prefix.

By default, both sets of TLVs are generated. When **wide-metrics-only** is configured, IS-IS only generates the pair of TLVs with wide metrics for that level.

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

4.10.2.2 Show commands

isis

Syntax

isis all
isis [isis-instance]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information for a specified IS-IS instance.

Parameters

instance-id
Specifies the instance ID for an IS-IS instance.

Values	0 to 31
Default	0

adjacency

Syntax

adjacency [ip-int-name | ip-address | nbr-system-id] [detail]

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information about IS-IS neighbors. If no parameters are specified, all adjacencies are displayed. If **detail** is specified, operational and statistical information is displayed.

Parameters

ip-int-name

Specifies the IP interface name. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, and spaces), the entire string must be enclosed within double quotes.

ip-address

Specifies the interface IP address.

Values

ipv4-address:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 to 32
ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
	x: [0 ..FFFF]H
	d: [0 ..255]D

nbr-system-id

Displays only the adjacency with the specified system ID.

Values 6-octet system identifier (xxxx.xxxx.xxxx)

detail

Displays detailed information about the adjacency.

Output

The following output is an example of ISIS adjacency information, and [Table 47: Output fields: ISIS adjacency](#) describes the output fields.

Sample output

```
*A:Dut-A# show router isis adjacency
=====
Router Base ISIS Instance 1 Adjacency
=====
System ID           Usage State Hold Interface           MT Enab
```

```

-----
Dut-B          L1    Up    2    ip-3FFE::A0A:101        Yes
Dut-B          L2    Up    2    ip-3FFE::A0A:101        Yes
Dut-F          L1L2  Up    5    ies-1-3FFE::A0A:1501    Yes
-----
Adjacencies : 3
=====
*A:Dut-A#

*A:ALA-A# show router isis adjacency 180.0.7.12
=====
Router Base ISIS Instance 1 Adjacency
=====
System ID          Usage State Hold Interface
-----
asbr_east          L2    Up    25    if2/5
-----
Adjacencies : 1
=====
*A:ALA-A#

*A:ALA-A# show router isis adjacency if2/5
=====
Router Base ISIS Instance 1 Adjacency
=====
System ID          Usage State Hold Interface
-----
asbr_east          L2    Up    20    if2/5
-----
Adjacencies : 1
=====
*A:ALA-A#

*A:Dut-A# show router isis adjacency detail
=====
Router Base ISIS Instance 1 Adjacency
=====
SystemID      : Dut-B          SNPA      : 20:81:01:01:00:01
Interface     : ip-3FFE::A0A:101 Up Time   : 0d 00:56:10
State        : Up             Priority    : 64
Nbr Sys Typ  : L1             L. Circ Typ : L1
Hold Time    : 2              Max Hold   : 2
Adj Level    : L1             MT Enabled  : Yes

IPv4 Neighbor : 10.10.1.2
Restart Support : Disabled
Restart Status : Not currently being helped
Restart Supressed : Disabled
Number of Restarts: 0
Last Restart at : Never

SystemID      : Dut-B          SNPA      : 20:81:01:01:00:01
Interface     : ip-3FFE::A0A:101 Up Time   : 0d 00:56:10
State        : Up             Priority    : 64
Nbr Sys Typ  : L2             L. Circ Typ : L2
Hold Time    : 2              Max Hold   : 2
Adj Level    : L2             MT Enabled  : Yes
Topology     : Unicast

IPv4 Neighbor : 10.10.1.2
Restart Support : Disabled

```

```

Restart Status      : Not currently being helped
Restart Supressed   : Disabled
Number of Restarts: 0
Last Restart at     : Never

SystemID      : Dut-F                      SNPA      : 00:00:00:00:00:00
Interface     : ies-1-3FFE::A0A:1501      Up Time   : 0d 01:18:34
State         : Up                        Priority   : 0
Nbr Sys Typ   : L1L2                     L. Circ Typ : L1L2
Hold Time     : 5                        Max Hold  : 6
Adj Level     : L1L2                     MT Enabled : Yes
Topology      : Unicast

IPv4 Neighbor    : 10.10.21.6
Restart Support  : Disabled
Restart Status   : Not currently being helped
Restart Supressed : Disabled
Number of Restarts: 0
Last Restart at  : Never
=====
*A:Dut-A#

A:Dut-A# show router isis status
=====
Router Base ISIS Instance 1 Status
=====
System Id          : 0100.2000.1001
Admin State        : Up
Ipv4 Routing       : Enabled
Last Enabled       : 08/28/2006 10:22:17
Level Capability    : L2
Authentication Check : True
Authentication Type : None
CSNP-Authentication : Enabled
HELLO-Authentication : Enabled
PSNP-Authentication : Enabled
Traffic Engineering : Enabled
Graceful Restart    : Disabled
GR Helper Mode      : Disabled
LSP Lifetime       : 1200
LSP Wait           : 1 sec (Max)  1 sec (Initial)  1 sec (Second)
Adjacency Check     : loose
L1 Auth Type        : none
L2 Auth Type        : none
L1 CSNP-Authenticati*: Enabled
L1 HELLO-Authenticat*: Enabled
L1 PSNP-Authenticati*: Enabled
L1 Preference       : 15
L2 Preference       : 18
L1 Ext. Preference  : 160
L2 Ext. Preference  : 165
L1 Wide Metrics     : Disabled
L2 Wide Metrics     : Enabled
L1 LSDB Overload    : Disabled
L2 LSDB Overload    : Disabled
L1 LSPs             : 0
L2 LSPs             : 15
Last SPF            : 08/28/2006 10:22:25
SPF Wait            : 1 sec (Max)  10 ms (Initial)  10 ms (Second)
Export Policies     : None
Area Addresses      : 49.0001
=====
* indicates that the corresponding row element may have been truncated.

```


A:Dut - A#

Table 47: Output fields: ISIS adjacency

Label	Description
Interface	Displays the interface name associated with the neighbor
System-id	Displays the neighbor system ID
Level	Displays the level: L1 only, L2 only, L1 and L2
State	Displays the state: Up, down, new, one-way, initializing, or rejected
Hold	Displays the hold time remaining for the adjacency
SNPA	Displays the subnetwork point of attachment (MAC address of the next hop)
Circuit type	Displays the level on the interface: L1, L2, or both
Expires In	Displays the number of seconds until adjacency expires
Priority	Displays the priority to become designated router
Up/down transitions	Displays the number of times the neighbor state has changed
Event	Displays the event causing the last transition
Last transition	Displays the amount of time since last transition change
Speaks	Displays supported protocols (only IP)
IP address	Displays the IP address of the neighbor
MT enab	Yes — The neighbor is advertising at least 1 non MTID#0
Topology	Derived from the MT TLV in the IIH <ul style="list-style-type: none"> MT#0, MT#2 => "Topology : Unicast" Native IPv4 Not supported MTID's => Topology line suppressed

capabilities

Syntax

capabilities [*system-id* | *lsp-id*] [**level** *level*]

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays IS-IS capability information.

Parameters

system-id

Displays only the IS-IS capabilities related to the specified system ID. If no parameters are specified, all database entries are displayed.

lsp-id

Displays only IS-IS capabilities related to the specified LSP ID. If no system ID or LSP ID is specified, all database entries are displayed.

level

Displays the interface level capabilities (1, 2, or 1 and 2).

Output

The following output is an example of IS-IS capability information, and [Table 48: Output fields: IS-IS capabilities](#) describes the output fields.

Sample output

```
*A:Dut-C# show router isis capabilities
=====
Rtr Base ISIS Instance 0 Capabilities
=====
Displaying Level 1 capabilities
-----
LSP ID : Dut-A.00-00
Router Cap : 10.20.1.1, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-A.02-00
LSP ID : Dut-A.03-00
LSP ID : Dut-B.00-00
Router Cap : 10.20.1.2, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-C.00-00
Router Cap : 10.20.1.3, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-C.02-00
LSP ID : Dut-D.00-00
Router Cap : 10.20.1.4, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
```

```
LSP ID : Dut-D.01-00
LSP ID : Dut-E.00-00
Router Cap : 10.20.1.5, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-E.01-00
LSP ID : Dut-E.02-00
LSP ID : Dut-F.00-00
Router Cap : 10.20.1.6, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-F.01-00
LSP ID : Dut-F.03-00
Level (1) Capability Count : 14
Displaying Level 2 capabilities
```

```
-----
LSP ID : Dut-A.00-00
Router Cap : 10.20.1.1, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-A.02-00
LSP ID : Dut-A.03-00
LSP ID : Dut-B.00-00
Router Cap : 10.20.1.2, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-C.00-00
Router Cap : 10.20.1.3, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-C.02-00
LSP ID : Dut-D.00-00
Router Cap : 10.20.1.4, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-D.01-00
LSP ID : Dut-E.00-00
Router Cap : 10.20.1.5, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-E.01-00
LSP ID : Dut-E.02-00
LSP ID : Dut-F.00-00
Router Cap : 10.20.1.6, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-F.01-00
LSP ID : Dut-F.03-00
```

```
Level (2) Capability Count : 14
=====
*A:ALA-A#
```

Table 48: Output fields: IS-IS capabilities

Label	Description
LSP ID	Displays the LSP ID of the specified system ID or hostname
Router Cap	Displays the router IP address and capability
TE Node Cap	Displays the TE node capability
SR Cap	Displays the segment routing capability
SRGB Base	Displays the Segment Routing Global Block (SRGB) base index value and range
SR Alg	Displays the type of SR algorithm used for the specified LSP ID
Level (n) Capability Count	Displays the capability count for the specified level

database

Syntax

database [*system-id* | *lsp-id*] [**detail**] [**level** *level*]

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the entries in the IS-IS link state database. If no parameters are specified, all entries are displayed.

Parameters

system-id

Displays only the LSPs related to the specified *system-id*. If no *system-id* or *lsp-id* is specified, all database entries are listed.

lsp-id

Displays only the specified LSP (hostname). If no *system-id* or *lsp-id* is specified, all database entries are listed.

- detail**
Keyword to specify that all output is displayed in the detailed format.
- level**
Displays only the specified IS-IS protocol level attributes.

Output

The following output is an example of IS-IS database information, and [Table 49: Output fields: router IS-IS database](#) describes the output fields.

Sample output

```
*A:ALA-A# show router isis database
=====
Router Base ISIS Instance 1 Database
=====
LSP ID                               Sequence Checksum Lifetime Attributes
-----
Displaying Level 1 database
-----
abr_dfw.00-00                        0x50      0x164f    603      L1L2
Level (1) LSP Count : 1
Displaying Level 2 database
-----
asbr_east.00-00                      0x53      0xe3f5    753      L1L2
abr_dfw.00-00                        0x57      0x94ff    978      L1L2
abr_dfw.03-00                        0x50      0x14f1    614      L1L2
Level (2) LSP Count : 3
=====

*A:ALA-A#

*A:Dut-B# show router isis database Dut-A.00-00 detail
=====
Router Base ISIS Instance 1 Database
=====
Displaying Level 1 database
-----
Level (1) LSP Count : 0

Displaying Level 2 database
-----
LSP ID      : Dut-A.00-00                Level      : L2
Sequence    : 0x6                        Checksum   : 0xb7c4  Lifetime   : 1153
Version     : 1                          Pkt Type   : 20     Pkt Ver    : 1
Attributes  : L1L2                       Max Area   : 3
SysID Len   : 6                          Used Len   : 311    Alloc Len  : 311

TLVs :
Area Addresses:
Area Address : (2) 30.31
Supp Protocols:
Protocols    : IPv4
IS-Hostname   : Dut-A
Router ID    :
Router ID    : 10.20.1.1
I/F Addresses :
I/F Address  : 10.20.1.1
I/F Address  : 10.10.1.1
I/F Address  : 10.10.2.1
TE IS Nbrs   :
```

```
Nbr      : Dut-B.01
Default Metric : 1000
Sub TLV Len   : 98
IF Addr      : 10.10.1.1
MaxLink BW: 100000 kbps
Resvble BW: 100000 kbps
Unresvd BW:
    BW[0] : 10000 kbps
    BW[1] : 40000 kbps
    BW[2] : 40000 kbps
    BW[3] : 40000 kbps
    BW[4] : 50000 kbps
    BW[5] : 50000 kbps
    BW[6] : 50000 kbps
    BW[7] : 10000 kbps
Admin Grp : 0x0
TE Metric : 1000
SUBTLV BW CONSTS      : 8
    BW Model : 1
    BC[0]: 10000 kbps
    BC[1]: 0 kbps
    BC[2]: 40000 kbps
    BC[3]: 0 kbps
    BC[4]: 0 kbps
    BC[5]: 50000 kbps
    BC[6]: 0 kbps
    BC[7]: 0 kbps
TE IP Reach :
    Default Metric : 0
    Control Info:   , prefLen 32
    Prefix : 10.20.1.1
    Default Metric : 1000
    Control Info:   , prefLen 24
    Prefix : 10.10.1.0
    Default Metric : 1000
    Control Info:   , prefLen 24
    Prefix : 10.10.2.0

Level (2) LSP Count : 1
=====
*A:Dut-B#
```

Table 49: Output fields: router IS-IS database

Label	Description
LSP ID	<p>Displays the LSP ID</p> <p>LSP IDs are auto-assigned by the originating IS-IS node. The LSP ID is comprised of three sections. The first 6 bytes is the system ID for that node, followed by a single byte value for the pseudonode generated by that router, and a fragment byte which starts at zero.</p> <p>For example, if a router system ID is 1800.0000.0029, the first LSP ID is 1800.0000.0029.00-00. If there are too many routes, LSP ID 1800.0000.0029.00-01 is created to contain the excess routes. If the router is the Designated Intermediate System (DIS) on a broadcast network, a pseudo-node LSP is created. Usually the internal circuit ID is used to determine the ID assigned to the</p>

Label	Description
	<p>pseudonode. For instance, for circuit 4, an LSP pseudonode with ID 1800.0000.0029.04-00 is created.</p> <p>The router learns hostnames and uses the hostname in place of the system ID. An example of LDP IDs are:</p> <p>acc_arl.00-00</p> <p>acc_arl.00-01</p> <p>acc_arl.04-00</p>
Sequence	Displays the sequence number of the LSP that allows other systems to determine whether they have received the latest information from the source
Checksum	Displays the checksum of the entire LSP packet
Lifetime	Displays the amount of time, in seconds, that the LSP remains valid.
Attributes	<p>OV — The overload bit is set</p> <p>L1 — Specifies a Level 1 IS type</p> <p>L2 — Specifies a Level 2 IS type</p> <p>ATT — The attach bit is set. When this bit is set, the router can also act as a Level 2 router and can reach other areas</p>
LSP Count	Displays the sum of all the configured Level 1 and Level 2 LSPs.
LSP ID	Displays a unique identifier for each LSP composed of SysID, Pseudonode ID, and LSP name
Lifetime	Displays the remaining time until the LSP expires
Version	Displays the version/protocol ID extension. This value is always set to 1.
Pkt Type	Displays the PDU type number
Pkt Ver	Displays the version/protocol ID extension. This value is always set to 1.
Max Area	Displays the maximum number of area addresses supported
Sys ID Len	Displays the length of the system ID field (0 or 6 for 6 digits)
Use Len	Displays the actual length of the PDU
Alloc Len	Displays the amount of memory space allocated for the LSP
Area Address	Displays the area addresses to which the router is connected
Supp Protocols	Displays the data protocols that are supported

Label	Description
IS-Hostname	Displays the name of the router originating the LSP
Virtual Flag	0 — Level 1 intermediate systems report this octet as 0 to all neighbors 1 — Indicates that the path to a neighbor is a Level 2 virtual path used to repair an area partition
Neighbor	Displays the routers running interfaces to which the router is connected
Internal Reach	Displays a 32-bit metric. A bit is added for the ups and downs resulting from Level 2 to Level 1 route-leaking.
IP Prefix	Displays the IP addresses that the router knows about by externally-originated interfaces
Metrics	Displays a routing metric used in the IS-IS link-state calculation

hostname

Syntax

hostname

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the hostname database. There are no options or parameters.

Output

The following output is an example of IS-IS hostname database information, and [Table 50: Output fields: router IS-IS hostname](#) describes the output fields.

Sample output

```
A:ALA-A# show router isis hostname
=====
Router Base ISIS Instance 1 Hostnames
=====
System Id          Hostname
-----
1800.0000.0002     core_west
1800.0000.0005     core_east
1800.0000.0008     asbr_west
1800.0000.0009     asbr_east
```


1800.0000.0010	abr_sjc
1800.0000.0011	abr_lax
1800.0000.0012	abr_nyc
1800.0000.0013	abr_dfw
1800.0000.0015	dist_oak
1800.0000.0018	dist_nj
1800.0000.0020	acc_nj
1800.0000.0021	acc_ri
1800.0000.0027	dist_arl
1800.0000.0028	dist_msq
1800.0000.0029	acc_arl
1800.0000.0030	acc_msq
=====	
A:ALA-A#	

Table 50: Output fields: router IS-IS hostname

Label	Description
System-id	Displays the system identifier mapped to hostname
Hostname	Displays the hostname for the specific <i>system-id</i>
Type	Displays the type of entry (static or dynamic)

interface

Syntax

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

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command shows IS-IS interface information.
If no parameters are specified, all entries are displayed.

Parameters

ip-int-name

Specifies the IP interface name. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 character composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, and spaces), the entire string must be enclosed within double quotes.

ip-address

Specifies the interface IP address.

Values ipv4-address: a.b.c.d (host bits must be 0) ipv4-prefix-length: 0 — 32
 ipv6-address: x:x:x:x:x:x:x (eight 16-bit pieces)
 x:x:x:x:x:d.d.d.d
 x: [0 — FFFF]H
 d: [0 — 255]D

nbr-system-id
Displays only the adjacency with the specified system ID.

Values 6-octet system identifier (xxxx.xxxx.xxxx)

detail
Displays detailed information about the adjacency.

Output
The following output is an example of IS-IS interface information, and [Table 51: Output fields: IS-IS interface](#) describes the output fields.

Sample output

```
A:ALA-A# show router isis interface
=====
ISIS Interfaces
=====
Interface                Level CircID  Oper State  L1/L2 Metric
-----
system                   L1L2   1          Up          10/10
if2/1                    L2     8          Up          -/10
if2/2                    L1     5          Up          10/-
if2/3                    L1     6          Up          10/-
if2/4                    L1     7          Up          10/-
if2/5                    L2     2          Up          -/10
lag-1                    L2     3          Up          -/10
if2/8                    L2     4          Up          -/10
-----
Interfaces : 8
=====
A:ALA-A#

Router Base ISIS Instance 1 Interfaces

Router Base ISIS Instance 1 Interfaces

*A:JC-NodeA# show router isis interface detail
=====
Router Base ISIS Instance 1 Interfaces
=====
Interface      : ip-10.10.1.1          Level Capability: L1L2
Oper State     : Up                  Admin State      : Up
Auth Type      : None
Circuit Id     : 2                    Retransmit Int.  : 5
Type           : Broadcast            LSP Pacing Int.  : 100
Mesh Group     : Inactive              CSNP Int.        : 10
Bfd Enabled    : No
```

Level	: 1	Adjacencies	: 0
Desg. IS	: JC-NodeA		
Auth Type	: None	Metric	: 10
Hello Timer	: 9	Hello Mult.	: 3
Priority	: 64		
Passive	: No	Te-Metric	: 2
Level	: 2	Adjacencies	: 0
Desg. IS	: JC-NodeA		
Auth Type	: None	Metric	: 10
Hello Timer	: 9	Hello Mult.	: 3
Priority	: 64 : 10		
Passive	: No	Te-Metric	: 21
=====			
*A:JC-NodeA#			

Table 51: Output fields: IS-IS interface

Label	Description
Interface	Displays the interface name
Level	Displays the interface level (1, 2, or 1 and 2)
CirID	Displays the circuit identifier
Oper State	Up — The interface is operationally up
	Down — The interface is operationally down
L1/L2 Metric	Displays the interface metric for Level 1 and Level 2, if none are set to 0

lfa-coverage

Syntax

lfa-coverage

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command shows IS-IS LFA coverage information.

Output

The following output is an example of IS-IS LFA coverage information, and [Table 52: Output fields: LFA coverage](#) describes the output fields.

Sample output

```
A:ALA-A# show router isis lfa-coverage

Rtr Base ISIS Instance 0 LFA Coverage
=====
Topology Level Node IPv4
-----
IPv4 Unicast L1 4/4(100%) 826/826(100%)
IPv4 Unicast L2 2/2(100%) 826/826(100%)
IPv6 Unicast L1 3/3(100%) 0/0(0%)
IPv6 Unicast L2 0/0(0%) 0/0(0%)
=====
A:ALA-A#
```

Table 52: Output fields: LFA coverage

Label	Description
Topology	Displays the type of network
Level	Displays the IS-IS level in which LFA is enabled
Node	Displays the number of nodes in the level on which LFA is enabled
IPv4	Displays the number of IPv4 interfaces on the nodes on which LFA is enabled

link-group-member-status

Syntax

link-group-member-status *name* [**level** *level*]

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays IS-IS link group member status information.

Parameters

- name**
Specifies the link group member name, up to 32 characters.
- level**
Specifies the interface level.

Values 1, 2, or 1 and 2

Output

The following output is an example of IS-IS link group member status information, and [Table 53: Output fields: link group member status](#) describes the output fields.

Sample output

```
A:ALA-A# show router isis link-group-member-status
=====
Rtr Base ISIS Instance 0 Link-Group Member
=====
Link-group I/F name Level State
-----
toDutB ip-10.10.12.3 L1 Up
toDutB ip-10.10.3.3 L1 Up
toDutB ip-10.10.12.3 L2 Up
toDutB ip-10.10.3.3 L2 Up
-----
Legend: BER = bitErrorRate
=====
A:ALA-A#
```

Table 53: Output fields: link group member status

Label	Description
Link-group	Displays the link group
I/F name	Displays the interface name
Level	Displays the interface level (1, 2, or 1 and 2)
State	Up — The interface is operationally up Down — The interface is operationally down

link-group-status

Syntax

link-group-status *name* [**level** *level*]

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays IS-IS Link Group status information.

Parameters

- name**
Specifies the link group name, up to 32 characters.
- level**
Specifies the interface level.
- Values** 1, 2, or 1 and 2

Output

The following output is an example of IS-IS link group status information.

Sample output

```
A:ALA-A# show router isis link-group-status
=====
Rtr Base ISIS Instance 0 Link-Group Status
=====
Link-group Mbrs Oper Revert Active Level State
           Mbr  Mbr   Mbr
-----
toDutB      2    2    2    2    L1   normal
toDutB      2    2    2    2    L2   normal
toDutE      2    2    2    2    L1   normal
toDutE      2    2    2    2    L2   normal
=====
A:ALA-A#
```

prefix-sids

Syntax

prefix-sids [**ipv4-unicast**] [*ip-prefix[/prefix-length]*] [**sid** *sid*] [**adv-router** *system-id | hostname*]

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays IS-IS prefix SIDs.

Parameters

- ipv4-unicast**
Displays information about the IPv4 unicast prefix SIDs.
- ip-prefix* [*/prefix-length*]**
Displays the IP prefix and mask length.

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

sid

Displays information related to the specified segment routing ID.

Values	0 to 524287
--------	-------------

system-id | hostname

Displays only the prefix SIDs related to the specified system ID or host name, up to 38 characters.

Output

The following output is an example of IS-IS prefix SID information, and [Table 54: Output fields: prefix SIDs](#) describes the output fields.

Sample output

*A:Dut-C# show router isis prefix-sids

Rtr Base ISIS Instance 0 Prefix/SID Table				
Prefix	SID	Lvl/Typ	SRMS MT	AdvRtr Flags
10.0.0.1/32	1	2/Int.	N 0	Dut-B RNnP
10.0.0.1/32	1	2/Int.	N 0	Dut-C RNnP
10.0.0.1/32	1	1/Int.	N 0	Dut-D NnP
10.0.0.1/32	1	2/Int.	N 0	Dut-D NnP
10.0.0.1/32	1	2/Int.	N 0	Dut-E RNnP
10.20.1.2/32	1002	1/Int.	N 0	Dut-B NnP
10.20.1.2/32	1002	2/Int.	N 0	Dut-B NnP
10.20.1.2/32	1002	2/Int.	N 0	Dut-C RNnP
10.20.1.2/32	1002	2/Int.	N 0	Dut-D RNnP
10.20.1.2/32	1002	2/Int.	N 0	Dut-E RNnP
10.20.1.3/32	1003	2/Int.	N 0	Dut-B RNnP
10.20.1.3/32	1003	1/Int.	N 0	Dut-C NnP
10.20.1.3/32	1003	2/Int.	N 0	Dut-C NnP
10.20.1.3/32	1003	2/Int.	N 0	Dut-D RNnP
10.20.1.3/32	1003	2/Int.	N 0	Dut-E RNnP
10.20.1.4/32	1004	2/Int.	N 0	Dut-B RNnP

```

10.20.1.4/32          1004      2/Int.    N      Dut-C
                      0        RnNP
10.20.1.4/32          1004      1/Int.    N      Dut-D
                      0        NnP
10.20.1.4/32          1004      2/Int.    N      Dut-D
                      0        NnP
10.20.1.4/32          1004      2/Int.    N      Dut-E
                      0        RnNP
10.20.1.5/32          1005      2/Int.    N      Dut-B
                      0        RnNP
10.20.1.5/32          1005      2/Int.    N      Dut-C
                      0        RnNP
10.20.1.5/32          1005      2/Int.    N      Dut-D
                      0        RnNP
10.20.1.5/32          1005      1/Int.    N      Dut-E
                      0        NnP
10.20.1.5/32          1005      2/Int.    N      Dut-E
                      0        NnP
-----
No. of Prefix/SIDs: 25
Flags: R = Re-advertisement
       N = Node-SID
       nP = no penultimate hop POP
       E = Explicit-Null
       V = Prefix-SID carries a value
       L = value/index has local significance
=====
*A:Dut-C#

```

Table 54: Output fields: prefix SIDs

Label	Description
Prefix	Displays the IP prefix for the SID
SID	Displays the segment routing identifier (SID)
Lvl/Typ	Displays the level and type of SR
SRMS	Displays whether the prefix SID is advertised by the SR mapping service: Y (yes) or N (no)
MT	Displays the multicast tunnel number (0, 2, 3, or 4)
AdvRtr	Displays the advertised router name
Flags	Displays the flags related to the advertised router: R = re-advertisement N = node-SID nP = no penultimate hop POP E = explicit-null V = prefix-SID carries a value L = value/index has local significance

routes

Syntax

routes [ipv4-unicast | ipv6-unicast | mt *mt-id-number*] [*ip-prefix/prefix-length*] [alternative]

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the routes in the IS-IS route table.

Parameters

ipv4-unicast

Displays IPv4 unicast parameters.

ipv6-unicast

Displays IPv6 unicast parameters.

mt-idnumber

Displays unicast multi-topology information.

Values 0, 2

ip-prefix/prefix-length

Displays information for the specified IP address.

Values	
ipv4-address:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 to 32
ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D
ipv6-prefix-length:	[0 to 128]

alternative

Displays the level of protection per prefix.

Output

The following output is an example of IS-IS route information, and [Table 55: Output fields: IS-IS routes](#) describes the output fields.

Sample output

*A:Dut-A# show router isis routes

Router Base ISIS Instance 1 Route Table

Prefix	Metric	Lvl/Typ	Ver.	SysID/Hostname	MT	
NextHop						
10.10.1.0/24				10	1/Int.	5
0.0.0.0				0		Dut-A
10.10.3.0/24	20	1/Int.		137		Dut-B
10.10.1.2				0		
10.10.4.0/24				20	1/Int.	137
10.10.1.2				0		Dut-B
10.10.5.0/24				30	1/Int.	137
10.10.1.2				0		Dut-B
10.10.9.0/24				60	1/Int.	52
10.10.21.6				0		Dut-F
10.10.10.0/24				70	1/Int.	52
10.10.21.6				0		Dut-F
10.10.12.0/24				20	1/Int.	137
10.10.1.2				0		Dut-B
10.10.13.0/24				10	1/Int.	7
0.0.0.0				0		Dut-A
10.10.14.0/24				20	1/Int.	52
10.10.21.6				0		Dut-F
10.10.15.0/24				30	1/Int.	137
10.10.1.2				0		Dut-B
10.10.16.0/24				30	1/Int.	137
10.10.1.2				0		Dut-B
10.10.21.0/24				10	1/Int.	48
0.0.0.0				0		Dut-A
10.10.22.0/24				30	1/Int.	137
10.10.1.2				0		Dut-B
10.20.1.1/32				0	1/Int.	10
0.0.0.0				0		Dut-A
10.20.1.2/32				10	1/Int.	137
10.10.1.2				0		Dut-B
10.20.1.3/32				20	1/Int.	137
10.10.1.2				0		Dut-B
10.20.1.4/32				20	1/Int.	137
10.10.1.2				0		Dut-B
10.20.1.5/32				30	1/Int.	137
10.10.1.2				0		Dut-B
10.20.1.6/32				10	1/Int.	52
10.10.21.6				0		Dut-F
10.10.1.0/24				10	1/Int.	65
0.0.0.0				2		Dut-A
10.10.13.0/24				10	1/Int.	65
0.0.0.0				2		Dut-A
10.10.21.0/24				10	1/Int.	65
0.0.0.0				2		Dut-A
10.20.1.1/32				0	1/Int.	65
0.0.0.0				2		Dut-A

No. of Routes: 20

*A:Dut-A#

Table 55: Output fields: IS-IS routes

Label	Description
Prefix	Displays the route prefix and mask
Metric MT	Displays the route metric.
Lvl/Type	Displays the level (1 or 2) and the route type, Internal (Int) or External (Ext)
Version	Displays the SPF version that generated the route
Nexthop	Displays the system ID of next hop
Hostname	Displays the hostname for the specific <i>system-id</i>

spf-log

Syntax

spf-log [detail]

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the last 20 IS-IS SFP events.

Output

The following output is an example of IS-IS SPF log information, and [Table 56: Output fields: IS-IS SPF log](#) describes the output fields.

Sample output

```
A:ALA-48# show router isis spf-log
=====
Router Base ISIS Instance 1 SPF Log
=====
When                Duration      L1 Nodes   L2 Nodes   Event Count Type
-----
01/30/2007 11:01:54  <0.01s     1           1           3
-----
Log Entries : 1
=====
A:ALA-48#
```

Table 56: Output fields: IS-IS SPF log

Label	Description
When	Displays the timestamp when the SPF run started on the system
Duration	Displays the time (in hundredths of a second) required to complete the SPF run
L1 Nodes	Displays the number of Level 1 nodes involved in the SPF run
L2 Nodes	Displays the number of Level 2 nodes involved in the SPF run
Event Count	Displays the number of SPF events that triggered the SPF calculation
Log Entries	Displays the total number of log entries

statistics

Syntax

statistics

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information about IS-IS traffic statistics.

Output

The following output is an example of IS-IS traffic statistics information, and [Table 57: Output fields: IS-IS statistics](#) describes the output fields.

Sample output

```
A:dut-b>show>router>isis# statistics

=====
Router Base ISIS Instance 0 Statistics
=====
ISIS Instance      : 0                SPF Runs          : 2
Purge Initiated    : 0                LSP Regens.       : 36

CSPF Statistics
Requests           : 0                Request Drops     : 0
Paths Found        : 0                Paths Not Found   : 0
```

```

LFA Statistics
LFA Runs          : 1

-----
PDU Type   Received   Processed   Dropped    Sent      Retransmitted
-----
LSP         0          0          0          0          0
IIH         0          0          0          0          0
CSNP        0          0          0          0          0
PSNP        0          0          0          0          0
Unknown     0          0          0          0          0
=====
A:dut-b>show>router>isis#

```

Table 57: Output fields: IS-IS statistics

Label	Description
Purge Initiated	Displays the number of times purges have been initiated
SPF Runs	Displays the number of times shortest path first calculations have been made
LSP Regens	Displays the count of LSP regenerations
Requests	Displays the number of CSPF requests made to the protocol
Paths Found	Displays the number of responses to CSPF requests for which paths satisfying the constraints were found
PDU Type	Displays the PDU type
Received	Displays the count of link state PDUs received by this instance of the protocol
Processed	Displays the count of link state PDUs processed by this instance of the protocol
Dropped	Displays the count of link state PDUs dropped by this instance of the protocol
Sent	Displays the count of link state PDUs sent out by this instance of the protocol
Retransmitted	Displays the count of link state PDUs that were retransmitted by this instance of the protocol
LFA Runs	Displays the number of time the shortest path first algorithm has been run to calculate the LFA (backup path to a destination).

status

Syntax
status

Context
show>router>isis

Platforms
Supported on all 7210 SAS platforms as described in this document

Description
This command displays information about IS-IS status.

Output
The following output is an example of IS-IS status information, and [Table 58: Output fields: IS-IS status](#) describes the output fields.

Sample output

```
*A:ALU_SIM11>show>router>isis# status

=====
Router Base ISIS Instance 1 Status
=====
System Id           : 0010.0100.1002
Admin State         : Up
Ipv4 Routing         : Enabled
Last Enabled        : 07/06/2010 12:28:12
Level Capability     : L1L2
Authentication Check : True
Authentication Type  : None
CSNP-Authentication : Enabled
HELLO-Authentication : Enabled
PSNP-Authentication : Enabled
Traffic Engineering : Disabled
Graceful Restart     : Disabled
GR Helper Mode       : Disabled
LSP Lifetime         : 1200
LSP Wait             : 5 sec (Max)   0 sec (Initial)   1 sec (Second)
Adjacency Check      : loose
L1 Auth Type         : none
L2 Auth Type         : none
L1 CSNP-Authenticati*: Enabled
L1 HELLO-Authenticat*: Enabled
L1 PSNP-Authenticati*: Enabled
L1 Preference        : 15
L2 Preference        : 18
L1 Ext. Preference   : 160
L2 Ext. Preference   : 165
L1 Wide Metrics      : Disabled
L2 Wide Metrics      : Disabled
L1 LSDB Overload     : Disabled
```

```

L2 LSDB Overload      : Disabled
L1 LSPs               : 3
L2 LSPs               : 3
Last SPF              : 07/06/2010 12:28:17
SPF Wait              : 10 sec (Max)   1000 ms (Initial)   1000 ms (Second)
Export Policies       : None
Multicast Import      : None
Multi-topology        : Disabled
Advertise-Passive-On* : Disabled
Suppress Default      : Disabled
Default Route Tag     : None
Area Addresses        : 01
Ldp Sync Admin State  : Up
LDP-over-RSVP         : Disabled
Loopfree-Alternate    : Enabled
L1 LFA                : Included
L2 LFA                : Included

```

```

=====
* indicates that the corresponding row element may have been truncated.
*A:ALU_SIM11>show>router>isis#

```

Table 58: Output fields: IS-IS status

Label	Description
System-id	Displays the neighbor system ID
Admin State	Up — IS-IS is administratively up Down — IS-IS is administratively down
Ipv4 Routing	Enabled — IPv4 routing is enabled Disabled — IPv4 routing is disabled
Ipv6 Routing	Disabled — IPv6 routing is disabled Enabled, Native — IPv6 routing is enabled Enabled, Multi-topology — Multi-topology TLVs for IPv6 routing is enabled
Multi-topology	Disabled — Multi-topology TLVs for IPv6 routing is disabled Enabled — Multi-topology TLVs for IPv6 routing is enabled
Last Enabled	Displays the date/time when IS-IS was last enabled in the router
Level Capability	Displays the routing level for the IS-IS routing process
Authentication Check	True — All IS-IS mismatched protocol packets are rejected False — Authentication is performed on received IS-IS protocol packets but mismatched packets are not rejected
Authentication Type	Displays the method of authentication used to verify the authenticity of packets sent by neighboring routers on an IS-IS interface.

Label	Description
Traffic Engineering	Enabled — TE is enabled for the router Disabled — TE is disabled so that TE metrics are not generated and are ignored when received by this node
Graceful Restart	Enabled — Graceful restart is enabled for this instance of IS-IS on the router Disabled — Graceful restart capability is disabled for this instance of IS-IS on the router
Ldp Sync Admin State	Displays whether the IGP-LDP synchronization feature is enabled or disabled on all interfaces participating in the IS-IS routing protocol
Loopfree-Alternate	Displays the interface LFA status (included in LFA computation or excluded in LFA computations)
L1 LFA	Displays the LFA status for an IS-IS Level 1 (included in LFA computation or excluded in LFA computations)
L2 LFA	Displays the LFA status for an IS-IS Level 2 (included in LFA computation or excluded in LFA computations)

summary-address

Syntax

summary-address [*ip-prefix[/prefix-length]*]

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays IS-IS summary address information.

Parameters

ip-prefix/prefix-length

Displays information for the specified IP address.

Values

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

ipv4-prefix-length: 0 to 32

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

Output

The following output is an example of IS-IS summary address information, and [Table 59: Output fields: IS-IS summary address](#) describes the output fields.

Sample output

```
A:ALA-48# show router isis summary-address
=====
Router Base ISIS Instance 1 Summary Address
=====
Address Level
-----
10.0.0.0/8 L1
10.1.0.0/24 L1L2
10.1.2.3/32 L2
-----
Summary Addresses : 3
=====
A:ALA-48#
```

Table 59: Output fields: IS-IS summary address

Label	Description
Address	Displays the IP address
Level	Displays the IS-IS level from which the prefix should be summarized

topology

Syntax

topology [[ipv4-unicast | ipv6-unicast | mt *mt-id-number*][detail]]

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays IS-IS unicast topology information.

Parameters

ipv4-unicast

Displays IPv4 unicast parameters.

ipv6-multicast

Displays IPv6 unicast parameters.

mt *mt-id-number*

Displays unicast topology parameters.

Values 0, 2

detail

Displays detailed unicast topology information.

Output

The following output is an example of IS-IS unicast topology information, and [Table 60: Output fields: IS-IS topology](#) describes the output fields.

Sample output

```
*A:duth# show router isis 1 topology

=====
Router Base ISIS Instance 1 Topology Table
=====
Node                               Interface      Nexthop
-----
IS-IS IP paths (MT-ID 0),   Level 1
-----
dutb.00                           toB            dutb
dutg.00                           toG            dutg
duti.00                           toB            dutb
duti.01                           toB            dutb
duti.02                           toG            dutg
-----
IS-IS IP paths (MT-ID 0),   Level 2
-----
dutb.00                           toB            dutb
dutg.00                           toG            dutg
duti.00                           toB            dutb
duti.01                           toB            dutb
duti.02                           toG            dutg
=====
```

Table 60: Output fields: IS-IS topology

Label	Description
Node	Displays the IP address

Label	Description
Interface	Displays the interface name
Nexthop	Displays the next-hop IP address

4.10.2.3 Clear commands

isis

Syntax

isis [*isis-instance*]

Context

clear>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context clear and reset IS-IS protocol entities.

Parameters

isis-instance

Specifies the IS-IS instance.

Values 0 to 31

adjacency

Syntax

adjacency [*system-id*]

Context

clear>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears and resets the entries from the IS-IS adjacency database.

Parameters

system-id

Specifies the system ID. When the system ID is entered, only the specified entries are removed from the IS-IS adjacency database.

Values 6-octet system identifier (xxxx.xxxx.xxxx)

database

Syntax

database [*system-id*]

Context

clear>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command removes the entries from the IS-IS link-state database, which contains information about PDUs.

Parameters

system-id

Specifies the system ID. When the system ID is entered, only the specified entries are removed from the IS-IS link-state database.

Values 6-octet system identifier (xxxx.xxxx.xxxx)

export

Syntax

export

Context

clear>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command reevaluates route policies participating in the export mechanism, either as importers or exporters of routes.

spf-log

Syntax

spf-log

Context

clear>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears the SPF log.

statistics

Syntax

statistics

Context

clear>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears and resets IS-IS statistics.

4.10.2.4 Debug commands

isis

Syntax

isis [*isis-instance*]

Context

debug>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging of IS-IS protocol entities.

Parameters

isis-instance

Specifies the IS-IS instance.

Values 0 to 31

adjacency

Syntax

[no] adjacency [*ip-int-name* | *ip-address* | *nbr-system-id*]

Context

debug>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for IS-IS adjacency.
The **no** form of this command disables debugging.

Parameters

ip-int-name

Specifies the IP interface name. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, and spaces), the entire string must be enclosed within double quotes.

ip-address

Specifies the interface IP address.

Values

ipv4-address:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 to 32
ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D

cspf

Syntax

[no] cspf

Context

debug>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for IS-IS CSPF.

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

graceful-restart

Syntax

[no] graceful-restart

Context

debug>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for IS-IS graceful restart.

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

interface

Syntax

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

no interface

Context

debug>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for the IS-IS interface.
The **no** form of this command disables debugging.

Parameters

- ip-int-name*
Specifies the IP interface name. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, and spaces), the entire string must be enclosed within double quotes.
- ip-address*
Specifies the interface IP address.

Values	
ipv4-address:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 to 32
ipv6-address:	x::x::x::x::x::x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D

leak

Syntax

- leak ip-address*
- no leak**

Context

debug>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for IS-IS leaks.
The **no** form of this command disables debugging.

Parameters

ip-address

Specifies the interface IP address.

Values

ipv4-address:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 to 32
ipv6-address:	x::x::x::x::x::x (eight 16-bit pieces) x::x::x::x::x::x::d.d.d.d x: [0 to FFFF]H d: [0 to 255]D

lsdb

Syntax

[no] lsdb [*level-number*] [*system-id* | *lsp-id*]

Context

debug>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for Link State Database (LSDB).
The **no** form of this command disables debugging.

misc

Syntax

[no] misc

Context

debug>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for IS-IS misc.
The **no** form of this command disables debugging.

packet

Syntax

packet [*packet-type*] [*ip-int-name* | *ip-address*] [**detail**]

Context

debug>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for IS-IS packets.
The **no** form of this command disables debugging.

Parameters

ip-int-name
Specifies the IP interface name. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, and spaces), the entire string must be enclosed within double quotes.

ip-address
Specifies the interface IP address.

Values	
ipv4-address:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 to 32
ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D

rtm

Syntax

rtm [*ip-address*]
no rtm

Context

debug>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for the IS-IS route table manager (RTM).
The **no** form of this command disables debugging.

Parameters

ip-address

Specifies the interface IP address.

Values

ipv4-address:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 to 32
ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D

spf

Syntax

[no] spf [*level-number*] [*system-id*]

Context

debug>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for IS-IS SFP.

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

5 BGP

This chapter provides information about configuring BGP.

5.1 BGP overview

Border Gateway Protocol (BGP) is an inter-autonomous system routing protocol. An autonomous system (AS) is a network or group of routers logically organized and controlled by a common network administration. BGP enables routers to exchange network reachability information, including information about other ASs that traffic must traverse to reach other routers in other ASs. To implement BGP, the ASN must be specified in the **config>router** context. A 7210 SAS BGP configuration must contain at least one group and include information about at least one 7210 SAS neighbor (peer).

AS paths are the routes to each destination. Other attributes, such as the origin of the path, the multiple exit discriminator (MED), the local preference, and communities included with the route are called path attributes. When BGP interprets routing and topology information, loops can be detected and eliminated. Route preference for routes learned from the configured peers can be enabled among groups of routes to enforce administrative preferences and routing policy decisions.



Note:

- MP-BGP (family IPv4 and IPv6) for use in Layer 3 VPN services (also known as VPRN services) is supported on all 7210 SAS platforms as described in this document.
- BGP (family IPv4 and IPv6) is not available for use in the base routing instance. It is only available for use as a PE-CE routing protocol.
- The L2-VPN (BGP-AD) and EVPN BGP address families are supported on all 7210 SAS platforms as described in this document.

5.2 BGP communication

There are two types of BGP peers, internal BGP (iBGP) and external BGP (eBGP) ([Figure 19: BGP configuration](#)).

- iBGP is used to communicate with peers in the same autonomous system. Routes received from an iBGP peer in the same autonomous system are not advertised to other iBGP peers (unless the router is a route reflector) but can be advertised to an eBGP peer.
- eBGP is used to communicate with peers in different autonomous systems. Routes received from a router in a different AS can be advertised to both eBGP and iBGP peers.

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 known routers, reachable addresses, and associated path cost metrics to each router. BGP uses the information and path attributes to compile a network topology.

5.2.1 Message types

Four message types are used by BGP to negotiate parameters, exchange routing information and indicate errors. They are:

- **Open message**

After a transport protocol connection is established, the first message sent by each side is an Open message. If the Open message is acceptable, a Keepalive message confirming the Open is sent back. When the Open is confirmed, Update, Keepalive, and Notification messages can be exchanged.

Open messages consist of the BGP header and the following fields:

- **version**

The current BGP version number is 4.

- **local ASN**

The autonomous system number is configured in the **config>router** context.

- **hold time**

Configure the maximum time BGP will wait between successive messages (either keep alive or update) from its peer, before closing the connection. Configure the local hold time with in the **config>router>bgp** context.

- **BGP identifier**

IP address of the BGP system or the router ID. The router ID must be a valid host address.

- **Update message**

Update messages are used to transfer routing information between BGP peers. The information contained in the packet can be used to construct a graph describing the relationships of the various autonomous systems. By applying rules, routing information loops and some other anomalies can be detected and removed from the inter-AS routing.

The Update messages consist of a BGP header and the following optional fields:

- **unfeasible routes length**

The field length which lists the routes being withdrawn from service because they are considered unreachable.

- **withdrawn routes**

The associated IP address prefixes for the routes withdrawn from service.

- **total path attribute length**

The total length of the path field that provides the attributes for a possible route to a destination.

- **path attributes**

The path attributes presented in variable length TLV format.

- **network layer reachability information (NLRI)**

IP address prefixes of reachability information.

- **Keepalive message**

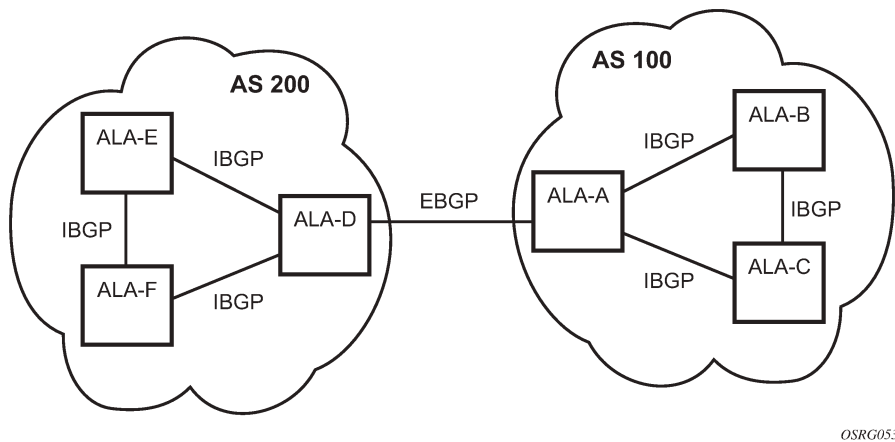
Keepalive messages, consisting of only a 19 octet message header, are exchanged between peers frequently so hold timers do not expire. The keepalive messages determine whether a link is unavailable.

- **Notification message**

A Notification message is sent when an error condition is detected. The peering session is terminated and the BGP connection (TCP connection) is closed immediately after sending it.

The following figure shows BGP configuration.

Figure 19: BGP configuration



5.3 Group configuration and peers

To enable BGP routing, participating routers must have BGP enabled and be assigned to an autonomous system and the neighbor (peer) relationships must be specified. A router typically belongs to only one AS. TCP connections must be established in order for neighbors to exchange routing information and updates. Neighbors exchange BGP open messages that includes information such as AS numbers, BGP versions, router IDs, and hold-time values. Keepalive messages determine whether a connection is established and operational. The hold-time value specifies the maximum time BGP will wait between successive messages (either keep alive or update) from its peer, before closing the connection.

In BGP, peers are arranged into groups. A group must contain at least one neighbor. A neighbor must belong to a group. Groups allow multiple peers to share similar configuration attributes.

Although neighbors do not have to belong to the same AS, they must be able to communicate with each other. If TCP connections are not established between two neighbors, the BGP peering will not be established and updates will not be exchanged.

Peer relationships are defined by configuring the IP address of the routers that are peers of the local BGP system. When neighbor and peer relationships are configured, the BGP peers exchange Update messages to advertise network reachability information.

5.4 Hierarchical levels

BGP parameters are initially applied on the global level. These parameters are inherited by the group and neighbor (peer) levels. Parameters can be modified and overridden on a level-specific basis. BGP command hierarchy consists of three levels:

- global level
- group level
- neighbor level

Many of the hierarchical BGP commands can be modified on different levels. The most specific value is used. That is, a BGP group-specific command takes precedence over a global BGP command. A neighbor-specific statement takes precedence over a global BGP and group-specific command; for example, if you modify a BGP neighbor-level command default, the new value takes precedence over group- and global-level settings.



Note:

Careful planning is essential to implement commands that can affect the behavior of global, group, and neighbor-levels. Because the BGP commands are hierarchical, analyze the values that can disable features on the global or group levels that must be enabled at the neighbor level. For example, if you enable the damping command on the global level but want it disabled only for a specific neighbor (not for all neighbors within the group), you cannot configure a double **no** command (**no no damping**) to enable the feature.

5.5 Route reflection

In a standard BGP configuration, all BGP speakers within an AS, must have full BGP mesh to ensure that all externally learned routes are redistributed through the entire AS. iBGP speakers do not re-advertise routes learned from one iBGP peer to another iBGP peer. If a network grows, scaling issues could emerge because of the full mesh configuration requirement. Instead of peering with all other iBGP routers in the network, each iBGP router only peers with a router configured as a route reflector.

Route reflection circumvents the full mesh requirement but maintains the full distribution of external routing information within an AS. Route reflection is effective in large networks because it is manageable, scalable, and easy to implement. Route reflection is implemented in autonomous systems with a large internal BGP mesh to reduce the number of iBGP sessions required within an AS.



Note:

The 7210 SAS can be configured only as route reflector clients. Only the client functionality of a route reflector described here is available for use with the 7210 SAS. The route reflector server-side functionality cannot be used on the 7210 SAS.


A large AS can be sub-divided into one or more clusters. Each cluster contains at least one route reflector which is responsible for redistributing route updates to all clients. Route reflector clients do not need to maintain a full peering mesh between each other. They only require a peering to the route reflectors in their cluster. The route reflectors must maintain a full peering mesh between all non-clients within the AS.

Each route reflector must be assigned a cluster ID and specify which neighbors are clients and which are non-clients to determine which neighbors should receive reflected routes and which should be treated as a

standard iBGP peer. Additional configuration is not required for the route reflector besides the typical BGP neighbor parameters.

The following figure shows a simple full-mesh configuration with several BGP routers. When SR-A receives a route from SR-1 (an external neighbor), it must advertise route information to all of its iBGP peers (SR-B, SR-C, SR-D, and so on). To prevent loops, iBGP learned routes are not re-advertised to other iBGP peers.

Figure 20: Fully meshed BGP configuration

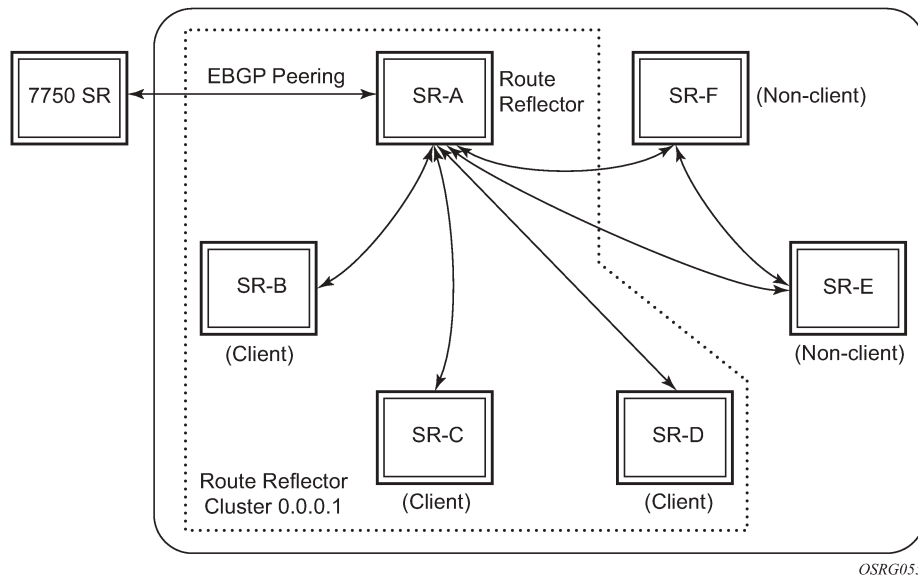


al_0138

When route reflectors are configured, the routers within a cluster do not need to be fully meshed. The preceding figure shows a fully meshed network and [Figure 21: BGP configuration with route reflectors](#) shows the same network but with route reflectors configured to minimize the iBGP mesh between SR-A, SR-B, SR-C, and SR-D. SR-A, configured as the route reflector, is responsible for redistributing route updates to clients SR-B, SR-C, and SR-D. iBGP peering between SR-B, SR-C and SR-D is not necessary because even iBGP learned routes are reflected to the route reflector's clients.

In the following figure, SR-E and SR-F are shown as non-clients of the route reflector. As a result, a full mesh of iBGP peerings must be maintained between, SR-A, SR-E and SR-F.

Figure 21: BGP configuration with route reflectors



A route reflector enables communication between the clients and non-client peers. Clients of a route reflector do not need to be fully meshed but non-client peers need to be fully meshed within an AS.

A grouping, called a cluster, is composed of a route reflector (or a redundant pair of route reflectors configured with the same cluster-id) and its client peers. Each route reflector is assigned a cluster ID and this defines the cluster that it and its clients belong to. Multiple route reflectors can be configured within a cluster for redundancy. A router assumes the role as a route reflector by configuring the **cluster cluster-id** command. No other command is required unless you want to disable reflection to specific clients.

When a route reflector receives an advertised route, it selects the best path. If the best path was received from an eBGP peer then it is typically advertised, with next hop unchanged, to all clients and non-client peers of the route reflector. If the best path was received from a non-client peer then it is advertised to all clients of the route reflector. If the best path was received from a client then it is advertised to all clients and non-client peers.

5.6 Fast external failover

Fast external failover on a group and neighbor basis is supported. For eBGP neighbors, this feature controls whether the router should drop an eBGP session immediately upon an interface-down event, or whether the BGP session should be kept up until the hold-time expires.

When fast external failover is disabled, the eBGP session stays up until the hold-time expires or the interface comes back up. If the BGP routes become unreachable as a result of the down IP interface, BGP withdraws the unavailable route immediately from other peers.

5.7 Sending of BGP communities

The capability to explicitly enable or disable the sending of the BGP community attribute to BGP neighbors, other than through the use of policy statements, is supported.

This feature allows an administrator to enable or disable the sending of BGP communities to an associated peer. This feature overrides communities that are already associated with a specific route or that may have been added via an export route policy. That is, even if the export policies leave BGP communities attached to a specific route, when the disable-communities feature is enabled, no BGP communities are advertised to the associated BGP peers.

5.8 ECMP and BGP route tunnels



Note:

ECMP is not supported for BGP route tunnels.

ECMP is not available for BGP route tunnels and also not on the transport LSP that is used to resolve BGP next-hop. If multiple LSP next-hops are available, only then the first next-hop is used and the rest ignored.

5.9 Next-hop resolution of BGP labeled routes to tunnels

The user enables the resolution of RFC 3107 BGP label route prefixes using tunnels to BGP next hops in the TTM with using following commands:

```
config>router>bgp>next-hop-res
      labeled-route-transport-tunnel
      [no] family family
            resolution {any | disabled | filter}
            resolution-filter
                  [no] ldp
                  [no] rsvp
                  [no] sr-isis
                  [no] sr-ospf
```

The **transport-tunnel** context allows the user to configure different types of BGP label routes: label-IPv4 and VPN routes (which includes both VPN-IPv4 and VPN-IPv6 routes). By default, all labeled routes resolve to LDP, even if the preceding CLI commands are not configured in the system.

If the **resolution** command is set to **disabled**, the default binding to LDP tunnels resumes. If **resolution** is set to **any**, the supported tunnel type selection is based on the TTM preference. The order of preference of TTM tunnels is the following:

- RSVP
- LDP
- segment routing OSPF
- segment routing IS-IS

If the **rsvp** option is enabled, BGP searches for the best metric RSVP LSP to the address of the BGP next-hop. The address can correspond to the system interface or to another loopback used by the BGP instance

on the remote node. MPLS provides the LSP metric in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel ID.

If the **ldp** option is enabled, BGP searches for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.

If the **sr-isis** or **sr-ospf** option is enabled, an SR tunnel to the BGP next-hop is selected in the TTM from the lowest preference IS-IS or OSPF instance. If many instances have the same lowest preference, the lowest numbered IS-IS or OSPF instance is chosen.

If one or more explicit tunnel types are specified using the **resolution-filter** option, only these tunnel types are selected again following the TTM preference. The **resolution** command must be set to **filter** to activate the list of tunnel types configured in the **resolution-filter** context.

5.9.1 VPN-IPv4 and VPN-IPv6 route resolution

The user enables the resolution of VPN-IPv4 and VPN-IPv6 prefixes using tunnels to MP-BGP peers using the following commands:

```
config>service>vprn
  auto-bind-tunnel
    resolution {any|disabled|filter}
    resolution-filter
      [no] ldp
      [no] rsvp
      [no] sr-isis
      [no] sr-ospf
```

The **auto-bind-tunnel** context configures the binding of VPRN routes to tunnels. The user must configure the **resolution** command to enable auto-bind resolution to tunnels in the TTM. If the **resolution** command is set to **disabled**, auto-binding to a tunnel is removed.

If the **resolution** command is set to **any**, any supported tunnel type in the **vprn** context is selected following the TTM preference. If one or more explicit tunnel types are specified using the **resolution-filter** command, only these tunnel types are selected again following the TTM preference. The following tunnel types are supported in a **vprn** context in order of preference: RSVP, LDP, and segment routing (SR).

If the **rsvp** command is enabled, BGP searches for the best metric RSVP LSP to the address of the BGP next-hop. This address can correspond to the system interface or to another loopback that the BGP instance uses on the remote node. MPLS provides the LSP metric in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel ID.

If the **ldp** command is enabled, BGP searches for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.

If the **sr-isis** or **sr-ospf** command is configured, an SR tunnel to the BGP next-hop is selected in the TTM from the lowest preference ISIS or OSPF instance. If many instances have the same lowest preference, the lowest numbered IS-IS or OSPF instance is chosen.

The BGP tunnel type is not explicitly configured in VPRN resolution and is therefore implicit. It is always preferred over any other tunnel type enabled in the **auto-bind-tunnel** context. However, the BGP tunnel type is configurable as a new tunnel type for BGP EVPN prefixes. The user must enable the BGP tunnel type to ensure that inter-area or inter-as prefixes are resolved.

The user must set the **resolution** command to **filter** to activate the list of tunnel types configured under **resolution-filter**.

When configured in a VPRN service (using the **configure>service>vprn>spoke-sdp** command), an explicit SDP to a BGP next-hop overrides the **auto-bind-tunnel** selection for that BGP next-hop only. There is no support for reverting automatically to the **auto-bind-tunnel** selection if the explicit SDP goes down. The user must delete the explicit spoke-SDP in the VPRN service to resume using the **auto-bind-tunnel** selection for the BGP next-hop.

5.10 Route selection criteria

For each prefix in the routing table, the routing protocol selects the best path. Then, the best path is compared to the next path in the list until all paths in the list are exhausted. The following parameters are used to determine the best path:

1. Routes are not considered if they are unreachable.
2. An RTM's preference is lowered as well as the hierarchy of routes from a different protocol. The lower the preference the higher the chance of the route being the active route.
3. Routes with higher local preference have preference.
4. Routes with the shorter AS path have preference.
5. Routes with the lower origin have preference. IGP = 0 EGP = 1 INCOMPLETE = 2
6. Routes with the lowest MED metric have preference. Routes with no MED value are exempted from this step unless `always-compare-med` is configured.
7. Routes learned by an eBGP peer rather than those learned from an iBGP peer are preferred.
8. Routes with the lowest IGP cost to the next-hop path attribute are preferred.
9. Routes with the lowest BGP-ID are preferred.
10. Routes with shortest cluster list are preferred.
11. Routes with lowest next-hop IP address are preferred.



Note:

- For BGP-VPN routes with the same prefix but a different Route Distinguisher (RD) that are imported in a VRF, if ECMP is not enabled in that VRF, the preceding selection criteria are used until parameter point 8. If all selection criteria are still the same after that point, the last updated route will be selected.
- For BGP-VPN routes with the same prefix but a different Route Distinguisher (RD) that reach parameter point 8 in the selection criteria, all routes are flagged as BEST and USED, although the actual number of used routes depends on the ECMP value configured in the VRF.



Note:

7210 SAS devices do not support BGP ECMP (multi-path). That is, an ECMP value of 1 is always used.

- For BGP-VPN routes with the same prefix and same Route Distinguisher (RD) that reach parameter point 8 in the selection criteria, such routes are flagged as BEST but parameter points 9 to 11 determine which routes are submitted to the VRF and marked as USED in accordance with the ECMP value configured in the VRF.

5.11 BGP path attributes

A BGP route for a specific NLRI is distinguished from other BGP routes for the same NLRI by its set of path attributes. Each path attribute is encoded as a TLV in the Path Attributes field of the Update message, and describes a property of the path. The type field of the TLV identifies the path attribute and the value field carries data specific to the attribute type.

The 7210 SAS supports the following path attributes:

- ORIGIN (well-known mandatory)
- AS_PATH (well-known mandatory)
- NEXT_HOP (well-known; required only in Update messages that have IPv4 prefixes in the NLRI field); see [Next-hop indirection](#) for information about the NEXT_HOP attribute
- MED (optional non-transitive)
- LOCAL_PREF (well-known; required only in Update messages sent to iBGP peers)
- ATOMIC_AGGR (well-known discretionary)
- AGGREGATOR (optional transitive)
- COMMUNITY (optional transitive)
- ORIGINATOR_ID (optional non-transitive)
- CLUSTER_LIST (optional non-transitive)
- MP_REACH_NLRI (optional non-transitive)
- MP_UNREACH_NLRI (optional non-transitive)
- EXT_COMMUNITY (optional transitive)
- AS4_PATH (optional transitive)
- AS4_AGGREGATOR (optional transitive)
- CONNECTOR (optional transitive)
- PMSI_TUNNEL (supported only on platforms that support NG-MVPN with BGP signaling; see the *7210 SAS Software Release Notes 25.x.Rx* for more information about NG-MVPN with BGP signaling).

5.11.1 NEXT_HOP attribute

The NEXT_HOP attribute indicates the IPv4 address of the BGP router that is the next hop to reach the IPv4 prefixes in the NLRI field. If the Update message is advertising routes other than IPv4 unicast routes, next hop of these routes is encoded in the MP_REACH_NLRI attribute and the NEXT_HOP attribute is not included in the Update message.

5.11.1.1 Next-hop indirection

The 7210 SAS supports next-hop indirection for most types of BGP routes. Next-hop indirection means that BGP next hops are logically separated from resolved next hops in the forwarding plane (IOMs). The separation allows the grouping of routes that share the same BGP next hops such that if the method of BGP next-hop resolution changes, only one forwarding plane update is required, instead of one update

for each route in the group. The convergence time after the next-hop resolution change is uniform, and not linear, with the number of prefixes. The next-hop indirection technology supports Prefix-Independent Convergence (PIC). The 7210 SAS uses next-hop indirection whenever possible; there is no option to disable the functionality.

On the 7210 SAS, the following families support next-hop indirection:

- label-IPv4
- VPN-IPv4
- label-IPv6
- VPN-IPv6
- L2-VPN
- PW route

5.12 BGP Routing Information Base

The entire set of BGP routes learned and advertised by a BGP router make up its BGP Routing Information Base (RIB). Conceptually, the BGP RIB contains three parts:

- RIB-IN
- LOC-RIB
- RIB-OUT

The RIB-IN (or Adj-RIBs-In, as defined in RFC 4271) contains the BGP routes received from peers that the router has stored in its memory.

The LOC-RIB contains modified versions of the BGP routes in the RIB-IN. The path attributes of a RIB-IN route can be modified using BGP import policies. All LOC-RIB routes for the NLRI are compared using the BGP decision process, which selects the best path for each NLRI. The local router uses the best paths in the LOC-RIB for forwarding, filtering, auto-discovery, and other tasks.

The RIB-OUT (or Adj-RIBs-Out, as defined in RFC 4271) contains the BGP routes selected for advertisement to peers. A BGP route is generally not advertised to a peer; that is, the router is not held in the RIB-OUT unless it is used locally, but there are exceptions. BGP export policies modify the path attributes of a LOC-RIB route to create the path attributes of the RIB-OUT route. A specific LOC-RIB route can be advertised with different path attribute values to different peers, and a 1:N relationship may exist between LOC-RIB and RIB-OUT routes.

5.12.1 LOC-RIB features

The 7210 SAS implements the following LOC-RIB processing features:

- BGP decision process
- BGP route installation in the route table
- BGP route installation in the tunnel table
- BGP fast reroute
- route flap damping (RFD)

See [BGP fast reroute](#) for more information about BGP fast reroute.

5.12.2 BGP fast reroute

BGP fast reroute uses indirection techniques in the forwarding plane and BGP backup path precomputation in the control plane to support the fast reroute of BGP traffic around unreachable or failed BGP next hops. BGP fast reroute is supported for label-IPv4 routes.

The following table describes the scenarios supported in the base router BGP context.

Table 61: BGP fast reroute scenarios (base router context)

Ingress packet	Primary route	Backup route	PIC
IPv4	IPv4 route with next-hop A resolved by an IPv4 route or any shortcut tunnel	IPv4 route with next-hop B resolved by an IPv4 route or any shortcut tunnel	No
IPv4	Label-IPv4 route with next-hop A resolved by any transport tunnel	Label-IPv4 route with next-hop B resolved by any transport tunnel	Yes
IPv4	Label-IPv4 route with next-hop A resolved by a local route	Label-IPv4 route with next-hop B resolved by a local route	Yes
IPv4	Label-IPv4 route with next-hop A resolved by a static route	Label-IPv4 route with next-hop B resolved by a static route	Yes

5.12.2.1 Calculating backup paths

BGP fast reroute is optional on the 7210 SAS. Use the **bgp backup-path** command to enable the feature.



Note:

On the 7210 SAS, the **backup-path** command is supported only for label-IPv4 routes.

In the base router context, the **backup-path** command is used to control fast reroute on a per-RIB basis (labeled IPv4 routes). When the command specifies a particular family, BGP attempts to find a backup path for every prefix learned by the associated BGP RIB.

The backup path is the best path after the primary path and any paths using the same BGP next hop as the primary path have been removed.

5.12.2.2 Failure detection and switchover to the backup path

When BGP fast reroute is enabled, BGP decides when a primary path is no longer usable and notifies the IOM. Based on BGP input, the IOM immediately reroutes affected traffic to the backup path.

When BGP fast reroute is enabled, the IOM reroutes traffic onto a backup path based on input from BGP. When BGP decides that a primary path is no longer usable, it notifies the IOM and affected traffic is immediately switched to the backup path.

The following events trigger failure notifications to the IOM and traffic rerouting to backup paths:

- peer IP address is unreachable and peer tracking is enabled
- BFD session associated with the BGP peer goes down

- BGP session is terminated with the peer (for example, send or receive NOTIFICATION)
- there is no longer any route (allowed by the next-hop resolution policy, if configured) that can resolve the BGP next-hop address
- BGP tunnel that resolves the next hop goes down because the BGP label-IPv4 route is withdrawn by the peer or becomes invalid due to an unresolved next hop

5.12.3 RIB-OUT features

This section describes features related to RIB-OUT processing.

5.12.3.1 BGP export policies

The **export** command is used to apply one or more policies (up to 15) to a neighbor or group, or to the entire BGP context. The **export** command that is most specific to a peer is applied. An **export** policy command applied at the **neighbor** level takes precedence over the same command applied at the **group** or global level. An **export** policy command applied at the **group** level takes precedence over the same command specified on the global level. The **export** policies applied at different levels are not cumulative. The policies listed in an **export** command are evaluated in the order in which they are specified in the configuration.



Note:

The **export** command can reference a policy before the policy has been created as a **policy-statement**.

The most common uses for BGP export policies are the following.

- BGP export policies can be used to locally originate a BGP route by exporting (or redistributing) a non-BGP route that is installed in the route table and actively used for forwarding. The non-BGP route is most frequently a direct, static, or aggregate route (exporting IGP routes into BGP is generally not recommended).
- BGP export policies can be used to block the advertisement of certain BGP routes towards specific BGP peers. The routes may be blocked on the basis of IP prefix, communities, and so on.
- BGP export policies can be used to modify the attributes of BGP routes advertised to specific BGP peers. The following path attribute modifications are possible using BGP export policies:
 - change the ORIGIN value
 - add a sequence of AS numbers to the start of the AS_PATH. When a route is advertised to an eBGP peer, the addition of the local-AS or global-AS numbers to the AS_PATH is always the final step (done after export policy).
 - replace the AS_PATH with a new AS_PATH. When a route is advertised to an eBGP peer, the addition of the local-AS or global-AS numbers to the AS_PATH is always the final step (done after export policy).
 - prepend an AS number multiple times to the start of the AS_PATH. When a route is advertised to an eBGP peer, the addition of the local-AS or global-AS numbers to the AS_PATH is always the final step (done after export policy). The add or replace action on the AS_PATH supersedes the prepend action if both are specified in the same policy entry.
 - change the NEXT_HOP to a specific IP address. When a route is advertised to an eBGP peer, the next hop cannot be changed from the local-address.

- change the NEXT_HOP to the local-address used with the peer (**next-hop-self**)
- add a value to the MED. If the MED attribute does not exist, it is added.
- subtract a value from the MED. If the MED attribute does not exist, it is added with a value of 0. If the result of the subtraction is a negative number, the MED metric is set to 0.
- set the MED to a specific value
- set the MED to the cost of the IP route (or tunnel) used to resolve the BGP next hop
- set LOCAL_PREF to a specific value when advertising to an iBGP peer
- add, remove, or replace standard communities
- add, remove, or replace extended communities
- add a static value to the AIGP metric when advertising the route to an AIGP-enabled peer with a modified BGP next hop. The static value is incremental to the automatic adjustment of the LOC-RIB AIGP metric to reflect the distance between the local router and the received BGP next hop.
- increment the AIGP metric by a fixed amount when advertising the route to an AIGP-enabled peer with a modified BGP next hop. The static value is a substitute for the dynamic value of the distance between the local router and the received BGP next hop.

5.12.3.2 Outbound Route Filtering

The ORF mechanism allows the ORF-sending router to signal to a peer, the ORF-receiving router, a set of route filtering rules (ORF entries) that the ORF-receiving router should apply to its route advertisements toward the ORF-sending router. The ORF entries are encoded in Route Refresh messages.

The use of ORF on a session must be negotiated; that is, both routers must advertise the ORF capability in their Open messages. The ORF capability describes the address families that support ORF, and for each address family, the ORF types that are supported and the ability to send and receive each type. 7210 SAS routers support ORF type 3, which is ORF based on extended communities, for only the following address families:

- VPN-IPv4
- VPN-IPv6
- MVPN-IPv4

On the 7210 SAS, the send and receive capability for ORF type 3 is configurable using the **send-orf** and **accept-orf** commands, but the setting applies to all supported address families.

The 7210 SAS support for ORF type 3 allows a PE router that imports VPN routes with a particular set of route target extended communities to indicate to a peer (for example, a route reflector) that it only wants to receive VPN routes that contain one or more of these extended communities. When the PE router wants to inform its peer about a new RT extended community, it sends a Route Refresh message to the peer containing an ORF type 3 entry instructing the peer to add a permit entry for the 8-byte extended community value. When the PE router wants to inform its peer about an RT extended community that is no longer needed, it sends a Route Refresh message to the peer containing an ORF type 3 entry instructing the peer to remove the permit entry for the 8-byte extended community value.

On the 7210 SAS, the type-3 ORF entries that are sent to a peer can be generated dynamically (if no route target extended communities are specified with the **send-orf** command) or specified statically. Dynamically generated ORF entries are based on the route targets that are imported by all locally-configured VPRNs.

A router that has installed ORF entries received from a peer can still apply BGP export policies to the session. If the evaluation of a BGP export policy results in a reject action for a VPN route that matches a permit ORF entry, the route is not advertised; that is, the export policy has the final word.



Note:

The 7210 SAS implementation of ORF filtering is efficient. It takes less time to filter a large number of VPN routes with ORF than it does to reject non-matching VPN routes using a conventional BGP export policy.

Despite the advantages of ORF compared to manually configured BGP export policies, RTC is the better technology when it comes to dynamic filtering based on route target extended communities. See [RT constrained route distribution](#) for more information about RTC.

5.12.3.3 RT constrained route distribution

The RTC mechanism allows a router to advertise an RTC route, which is a special type of MP-BGP route, to specific peers; the associated AFI is 1 and the SAFI is 132. The NLRI of an RTC route encodes an origin AS and a route target extended community with prefix-type encoding (for example, if there is a prefix-length and host bits after the prefix-length are set to zero). A peer receiving RTC routes does not advertise VPN routes to the RTC-sending router unless they contain a route target extended community that matches one of the received RTC routes. As with any other type of BGP route, RTC routes are propagated loop-free throughout and between ASs. If multiple RTC routes exist for the same NLRI, the BGP decision process selects one as the best path. The propagation of the best path installs RIB-OUT filter rules as it travels from one router to the next, and this process creates an optimal VPN route distribution tree rooted at the source of the RTC route.



Note:

RTC and extended community-based ORF mechanisms are similar in that they both allow a router to signal to a peer the route target extended communities they want to receive in VPN routes from that peer. However, RTC has distinct advantages over extended community-based ORF because it is more widely supported, it is simpler to configure, and its distribution scope is not limited to a direct peer.

The capability to exchange RTC routes is advertised when the **route-target** keyword is added to the relevant **family** command. RTC is supported on eBGP and iBGP sessions of the base router instance. On a specific session, either ORF or RTC may be used, but not both; if RTC is configured, the ORF capability is not announced to the peer.

RTC is supported for the following BGP address families:

- VPN-IPv4
- VPN-IPv6
- L2-VPN (BGP-AD)
- EVPN

When RTC is negotiated with one or more peers, the software automatically originates and advertises to these peers one /96 RTC route (the origin AS and route target extended community are fully specified) for every route target imported by a locally-configured VPRN or BGP-based Layer 2 VPN. Route targets are supported for all BGP families in the preceding list.



Note:

When **route-target** is enabled, it is activated for all address families configured on the node under BGP. Per-family activation is not supported.

The 7210 SAS also supports the **group** or **neighbor** level **default-route-target** command, which causes routers to generate and send a 0:0:0/0 default RTC route to one or more peers. Sending the default RTC route to a peer conveys a request to receive all VPN routes from that peer. The **default-route-target** command is typically configured on sessions that a route reflector has established with its PE clients. A received default RTC route is never propagated to other routers.

The route reflector advertises RTC routes in accordance with the rules described in RFC 4684. These rules ensure that RTC routes for the same NLRI that are originated by different PE routers in the same AS are correctly distributed within the AS.

When a BGP session comes up and RTC is enabled on the session (both peers advertised the MP-BGP capability), routers delay sending any VPN-IPv4 and VPN-IPv6 routes until either the session has been up for 60 seconds or the end-of-RIB marker is received for the RTC address family. When the VPN-IPv4 and VPN-IPv6 routes are sent, they are filtered to include only those with a route target extended community that matches an RTC route from the peer. VPN-IP routes matching an RTC route originated in the local AS are advertised to any iBGP peer that advertises a valid path for the RTC NLRI. That is, route distribution is not constrained to only the iBGP peer advertising the best path. However, VPN-IP routes matching an RTC route originated outside the local AS are only advertised to the eBGP or iBGP peer that advertises the best path.



Note:

The 7210 SAS does not support an equivalent of BGP-Multipath for RTC routes. There is no way to distribute VPN routes across more than one "almost" equal set of inter-AS paths.

5.12.3.4 Minimum Route Advertisement Interval

In accordance with RFC 4271, a BGP router should not send updated NLRI reachability information to a BGP peer until a specific period of time (the minimum route advertisement interval (MRAI)) has elapsed since the last update. The RFC suggests that the MRAI should be configurable per peer, but does not propose a specific algorithm; consequently, MRAI implementation details vary from one router operating system to another.

On the 7210 SAS, the MRAI is configurable on a per-session basis using the **min-route-advertisement** command. This CLI command can be configured with any value between 1 and 255 seconds, and the configuration applies to all address families. The default value is 30 seconds, regardless of the session type (eBGP or iBGP). The MRAI timer is started at the configured value when the session is established and counts down continuously. When the timer reaches zero, it resets to the configured value and all pending RIB-OUT routes are sent to the peer.

To send Update messages that advertise new NLRI reachability information more frequently for some address families than others, use the **rapid-update** command to override the remaining time on a peer MRAI timer and immediately send routes belonging to specified address families (and all other pending updates) to the peers receiving these routes. The following address families support **rapid-update**:

- EVPN
- L2-VPN

In many cases, the default MRAI is appropriate for all address families (or at least those not included in the preceding list) when it applies to Update messages that advertise reachable NLRI, but it is not the best option for Update messages that advertise unreachable NLRI (route withdrawals). Fast reconvergence

after some types of failures requires route withdrawals to propagate to other routers as quickly as possible so that they can calculate and start using new best paths, which would be impeded by the effect of the MRAI timer at each router hop. This is facilitated by the **rapid-withdrawal** configuration command.

When **rapid-withdrawal** is configured, Update messages containing withdrawn NLRI are sent immediately to a peer without waiting for the MRAI timer to expire. Update messages containing reachable NLRI continue to wait for the MRAI timer to expire, or for a **rapid-update** trigger, if it applies. When **rapid-withdrawal** is enabled, it applies to all address families.

5.12.3.5 Advertise-inactive

BGP does not allow a route to be advertised unless it is the best path in the RIB and an export policy allows the advertisement.

In some cases, it may be useful to advertise the best BGP path to peers despite the fact that the BGP path is inactive, for example, if the path is inactive because there are one or more preferred non-BGP routes to the same destination and one of these other routes is the active route. The 7210 SAS supports this flexibility using the **advertise-inactive** command; other supported methods include [Add-path](#).

When the BGP **advertise-inactive** command is configured on a BGP session, it has the following effect on the IPv4, IPv6, label-IPv4, and label-IPv6 routes advertised to that peer.

- If the active route for the IP prefix is a BGP route, that route is advertised. If the active route for the IP prefix is a non-BGP route and there is at least one valid but inactive BGP route for the same destination, the best of the inactive and valid BGP routes is advertised unless the non-BGP active route is matched and accepted by an export policy applied to the session.
- If the active route for the IP prefix is a non-BGP route and there are no (valid) BGP routes for the same destination, no route is advertised for the prefix unless the non-BGP active route is matched and accepted by an export policy applied to the session.

5.12.3.6 Split-horizon

Split-horizon refers to the action taken by a router to avoid advertising a route back to the peer from which it was received. By default, the 7210 SAS applies split-horizon behavior only to routes received from iBGP non-client peers, and split-horizon only works for routes to non-imported routes within a RIB. Split-horizon functionality, which can never be disabled, prevents a route learned from a non-client iBGP peer from being advertised to the sending peer or any other non-client peer.

To apply split-horizon behavior to routes learned from RR clients or eBGP peers, configure the **split-horizon** command in either the global BGP, **group** or **neighbor** contexts. When **split-horizon** is enabled on these types of sessions, it only prevents the advertisement of a route back to its originating peer; for example, the software does not prevent the advertisement of a route learned from one eBGP peer back to a different eBGP peer in the same neighbor AS.

5.13 Add-path

5.13.1 Receiving multiple paths per prefix from a BGP peer

If the 7210 SAS receives an advertisement of an NLRI and path from a specific peer and that peer subsequently advertises the same NLRI with different path information (a different next-hop or different path attributes), the new path overwrites the existing path.

However, when the add-path has been negotiated with the peer, the newly advertised path is stored in the RIB-IN along with all paths previously advertised (and not withdrawn) by the peer.

For router A to receive multiple paths per NLRI from peer B for a specific address family (AFI x, SAFI y), the BGP capabilities advertisement during session setup must indicate that peer B must send multiple paths for (AFI x, SAFI y) and router A is willing to receive multiple paths for (AFI x, SAFI y).

When the add-path receive capability for (AFI x, SAFI y) has been negotiated with a peer, all advertisements and withdrawals of NLRI within that address family by that peer include a path identifier.


If the add-path has been negotiated with a peer and a path identifier is expected but missing, or if the add-path has not been negotiated with a peer and a path identifier is present but not expected, a Notification message is sent with the error subcode indicating Invalid Network Field, in accordance with standard BGP error handling procedures.

The path identifiers have no significance to the receiving peer. If the combination of NLRI and path identifier in an advertisement from a peer is unique (does not match an existing route in the RIB-IN from that peer), the route is added to the RIB-IN. If the combination of NLRI and path identifier in a received advertisement is the same as an existing route in the RIB-IN from the peer, the new route replaces the existing one. If the combination of NLRI and path identifier in a received withdrawal matches an existing route in the RIB-IN from the peer, that route is removed from the RIB-IN.

A BGP Update message from an add-path peer may advertise and withdraw more than one NLRI belonging to one or more address families. In this case, the add-path may be supported for some address families and not others. In this situation, the receiving BGP router should not require that all path identifiers in the Update message be the same.

The following figure shows an Update message carrying an IPv4 NLRI with a path identifier.

Figure 22: BGP Update message with path identifier for IPv4 NLRI



Currently, add-path is only supported for the iBGP sessions with other add-path capable peers. The add-path capability is not supported for eBGP sessions or for native IPv4 and IPv6 routes (that is, IPv4 and IPv6 routes advertised without a label) in iBGP sessions. The ability to receive multiple paths per prefix from an add-path peer is configurable per route type. The supported route types are the following:

- label-IPv4
- label-IPv6

5.13.2 Path selection with add-paths

The LOC-RIB may have multiple paths for a prefix. The path selection mode refers to the algorithm used to decide which of these paths to advertise to an add-path peer. SR OS supports the Add-N path selection algorithm described in *draft-ietf-idr-add-paths-guidelines*. The Add-N algorithm selects as candidates for advertisement the N best paths with unique BGP next-hops. In the SR OS implementation, the default value of N is configurable per address family at the BGP instance, group, and neighbor levels; however, this default value can be overridden for specific prefixes using route policies. The maximum number of paths to advertise for a prefix to an add-path neighbor is the value N assigned by a BGP import policy to the best path for P; otherwise, it defaults to the neighbor, group, or instance level configuration of N for the address family to which P belongs.

Add-paths allows non-best paths to be advertised to a peer, but it still complies with basic BGP advertisement rules, such as the iBGP split horizon rule that a route learned from an iBGP neighbor cannot be readvertised to another iBGP neighbor unless the router is configured as a route reflector.

5.13.3 BGP decision process with add-path

To use multiple paths per NLRI for forwarding and to advertise multiple paths per NLRI to add-path peers, a router implementing an add-path must run a modified version of the BGP decision process. The existing BGP decision algorithm selects the one best path for any particular NLRI. Paths that are second best or third best remain in the RIB-IN but are not installed in the LOC-RIB and not advertised to peers.

The system automatically changes its BGP decision process for routes belonging to a particular address family whenever either of the following applies:

- BGP edge PIC is enabled for the address family
- the add-path send capability is enabled for that address family on one or more peering sessions

When BGP PIC is enabled, the BGP decision process selects a backup path per prefix or NLRI to install in the LOC_RIB. The algorithm is summarized as follows.

1. Select the single best path based on a full evaluation of all the BGP tie-breaking rules, as described in the following examples.
 - a. Select the route with the highest route preference.
 - b. From all routes with an AIGP metric, select the route with the lowest sum of the AIGP metric value stored with the RIB-IN copy of the route and the iteratively resolved distance between the calculating router and the BGP NEXT_HOP of the route.
 - c. Select the route with the highest local preference (LOCAL_PREF).
 - d. Select the route with the shortest AS path.
 - e. Select the route with the lowest origin.

- f. Among routes advertised by the same neighbor AS (unless **always-compare-med** is configured).
Select the route with the lowest MED.
 - g. Prefer routes learned from eBGP peers over routes learned from iBGP peers.
 - h. Select the route with the lowest IGP cost (unless **ignore-nh-metric** is configured).
 - i. Select the route received by the peer with the lowest originator ID or BGP identifier.
 - j. Select the route with the shortest cluster list.
 - k. Select the route received from the lowest peer IP address.
2. Select up to one additional second best path among the paths remaining after removing from consideration all paths with a NEXT_HOP or BGP identifier (or originator ID) in common with any of the previously-selected best paths. A full evaluation of all the BGP tie-breaking rules is required to find this single second-best path, as shown in the following examples.
 - a. Select the route with the highest route preference.
 - b. Select the route with the highest local preference (LOCAL_PREF).
 - c. Select the route with the shortest AS path (unless **as-path-ignore** is configured).
 - d. Select the route with the lowest origin.
 - e. Among routes advertised by the same neighbor AS (unless **always-compare-med** is configured) select the route with the lowest MED.
 - f. Prefer routes learned from eBGP peers over routes learned from iBGP peers.
 - g. Select the route with the lowest IGP cost.
 - h. Select the route received by the peer with the lowest originator ID or BGP identifier.
 - i. Select the route with the shortest cluster list.
 - j. Select the route received from the lowest peer IP address.

5.13.4 Advertising multiple paths using add-path

For router A to send multiple paths per NLRI to peer B for a particular address family (AFI x, SAFI y), the BGP capability advertisement during session setup must indicate that router A must send multiple paths for (AFI x, SAFI y), and peer B is willing to receive multiple paths for (AFI x, SAFI y).

By default, unless changed through configuration, all paths for a particular NLRI in the LOC-RIB are advertised to all add-path peers with which the send capability has been negotiated. All such advertisements (and any subsequent withdrawals) include a path identifier. Each advertised path for a specific NLRI must have a unique path identifier. When a path is reflected or propagated from one peer to another, the path identifier is expected to change, even if there has been no change in the next-hop. A BGP Update message sent to an add-path peer may advertise and withdraw more than one NLRI belonging to one or more address families. In this case, the add-path may be supported for some address families and not others, and the path identifiers associated with different NLRI in the Update message may be the same or different.

In the current implementation, the add-path is only supported by the iBGP sessions it forms with other add-path capable peers. The add-path capability is not supported for eBGP sessions or for native IPv4 and IPv6 routes (that is, IPv4 and IPv6 routes advertised without a label) in iBGP sessions. The ability to receive multiple paths per prefix from an add-path peer is configurable per route type. The following route types are supported:

- label-IPv4
- label-IPv4

5.13.5 Limiting the number of paths per prefix

Advertising multiple paths per prefix to a peer means that the peer must maintain more entries in its RIB-IN than would be the case without add-path. The memory and CPU resources associated with these extra paths may not be justified if the peer cannot take advantage of them. Operators may therefore want precise control over the number of paths per prefix to send to particular peers.

The new add-paths CLI node (BGP, group or neighbor level) has address family-specific commands to set the maximum number of paths to send per prefix.

To ensure routing consistency in cases where an add-path speaking router has a mix of add-path and non add-path peers and where the number of paths to send for a particular prefix can vary by add-path peer, the following behavior should be enforced: if the advertising router advertises n paths for prefix XYZ to peer1 and m paths to peer2, and $n < m$, then all the paths advertised to peer1 must be included in the paths advertised to peer2. Suppose the LOC-RIB has N paths for prefix XYZ. The preceding behavior can be guaranteed if:

- the N paths are sorted in strict order of their preference by the BGP decision process: p_1 (overall best path found during step 1 of [BGP decision process with add-path](#)), p_2 , p_3 , ..., p_N (a path found during step 2 or 3 of [BGP decision process with add-path](#))
- p_1 (only) is advertised to non add-path peers, add-path peers that indicate a send-only capability and add-path peers for which the configured path-limit is 1
- (p_1, p_2) is advertised to add-path peers for which the configured path-limit is 2
- $(p_1, p_2, p_3, \dots, p_N)$ is advertised to add-path peers for which the configured path-limit is N , or else the path limit is configured as **max** (default)

5.14 AIGP metric

The accumulated IGP (AIGP) metric is an optional, non-transitive attribute that can be attached to selected routes using route policies. In networks that use AIGP, BGP paths with a lower end-to-end IGP cost are preferred, even if the compared paths span more than one AS or IGP instance. AIGP differs from MED in the following important ways:

- AIGP is not transitive between completely distinct autonomous systems. It is only transitive across internal AS boundaries.
- AIGP is always compared in paths that have the AIGP attribute, regardless of whether they are located in different neighbor ASs.
- AIGP is more important than MED in the BGP decision process.
- AIGP is automatically incremented every time there is a BGP next-hop change, so that the system can track the end-to-end IGP cost. All arithmetic operations on MED attributes must be performed manually, such as by using route policies.

On the 7210 SAS, AIGP is supported only in the base router BGP instance and only for label-IPv4 and 6PE routes. The AIGP attribute is only sent to peers configured using the **aigp** command. If the attribute is received from a peer that is not configured for AIGP, or if the attribute is received in a non-supported route

type, the attribute is discarded and not propagated to other peers. The AIGP attribute is still displayed in BGP **show** command output.

When the 7210 SAS receives a route with an AIGP attribute and it re-advertises the route to an AIGP-enabled peer without changes to the BGP next hop, the AIGP metric value is unchanged by the advertisement (RIB-OUT) process. However, if the route is re-advertised with a new BGP next hop, the AIGP metric value is automatically incremented, either by the route table or tunnel table cost to reach the received BGP next hop, or by a value configured using route policies.

5.15 Command interactions and dependencies

This section describes the BGP command interactions and dependencies that apply to the configuration or operational maintenance of 7210 SAS routers.

See the [BGP command reference](#) for detailed descriptions of the BGP configuration commands.

5.15.1 Changing the ASN

If the autonomous system number (ASN) is changed on a router with an active BGP instance, the new ASN is not used until the BGP instance is restarted, either by administratively disabling or enabling the BGP instance or by rebooting the system with the new configuration.

5.15.2 Changing the local ASN

Changing the local ASN of an active BGP instance:

- at the global level, causes the BGP instance to restart with the new local ASN
- at the group level, causes BGP to re-establish the peer relationships with all peers in the group with the new local ASN
- at the neighbor level, causes BGP to re-establish the peer relationship with the new local ASN

5.15.3 Changing the router ID at the configuration level

If you configure a new router ID in the **config>router** context, protocols are not automatically restarted with the new router ID. The updated router ID is only used the next time the protocol is initialized or reinitialized. An interim period can occur when the protocols use different router IDs.

5.15.4 Hold time and keep alive timer dependencies

The BGP hold time specifies the maximum time BGP will wait between successive messages (either keep alive or update) from its peer, before closing the connection. This configuration parameter can be set at three levels. The most specific value is used.

- Global level — applies to all peers
- Group level — applies to all peers in group
- Neighbor level — only applies to specified peer

Although the keep alive time can be user specified, the configured keep alive timer is overridden by the value of hold time under the following circumstances:

- If the hold time specified is less than the configured keep alive time, then the operational keep alive time is set to one third of the specified hold time; the configured keep alive time is unchanged.
- If the hold time is set to zero, then the operational value of the keep alive time is set to zero; the configured keep alive time is unchanged. This means that the connection with the peer will be up permanently and no keep alive packets are sent to the peer.

If the hold time or keep alive values are changed, the changed timer values take effect when the new peering relationship is established. Changing the values cause the peerings to restart. The changed timer values are used when renegotiating the peer relationship.

5.15.5 Import and export route policies

Import and export route policy statements are specified for BGP on the global, group, and neighbor level. Up to five unique policy statement names can be specified in the command line per level. The most specific command is applied to the peer. Defining the policy statement name is not required before being applied. Policy statements are evaluated in the order in which they are specified within the command context.

The import and export policies configured on different levels are not cumulative. The most specific value is used. An **import** or **export** policy command specified on the neighbor level takes precedence over the same command specified on the group or global level. An **import** or **export** policy command specified on the group level takes precedence over the same command specified on the global level.

5.15.6 Route damping and route policies

To prevent BGP systems from sending excessive route changes to peers, BGP route damping can be implemented. Damping can reduce the number of Update messages sent between BGP peers, to reduce the load on peers, without adversely affecting the route convergence time for stable routes.

The damping profile defined in the policy statement is applied to control route damping parameters. Route damping characteristics are specified in a route damping profile and are referenced in the action for the policy statement or in the action for a policy entry. Damping can be specified at the global, group, or neighbor level with the most specific command applied to the peer.

5.15.7 AS Override

The BGP-4 Explicit AS Override simplifies the use of the same ASN across multiple RFC 2547 VPRN sites.

The Explicit AS Override feature can be used in VPRN scenarios where a customer is running BGP as the PE-CE protocol and some or all of the CE locations are in the same Autonomous System (AS). With normal BGP, two sites in the same AS would not be able to reach each other directly since there is an apparent loop in the AS_PATH.

With AS Override enabled on an egress eBGP session, the Service Provider network can rewrite the customer ASN in the AS_PATH with its own ASN as the route is advertised to the other sites within the same VPRN.

5.16 Configuration guideline for BGP

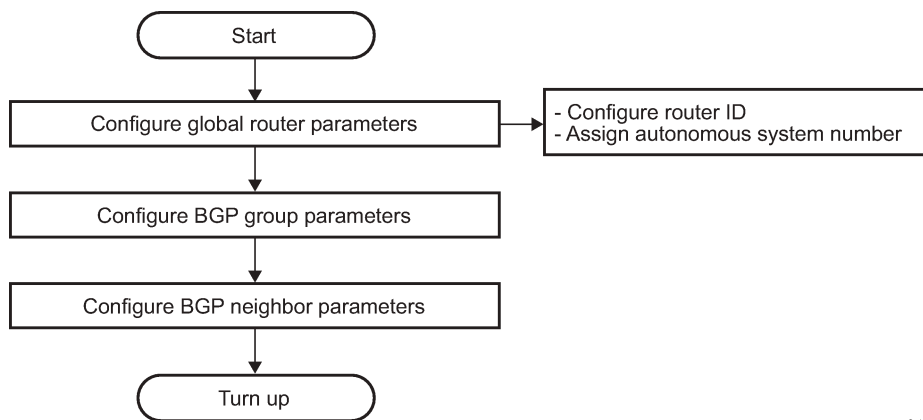
The following are the configuration guidelines for BGP:

- 7210 SAS can act only as a route reflector client.
- 7210 SAS support IPv4 family for PE-CE eBGP instance and for RFC 3107 labeled IPv4 routes. It supports IPv4 family in the base routing instance to exchange IPv4 routes.
- 7210 SAS support IPv6 family for PE-CE eBGP instance and for RFC 3107 labeled IPv6 routes. It supports IPv6 family in the base routing instance to exchange IPv6 routes.

5.17 BGP configuration process overview

The following figure shows the process to provision basic BGP parameters.

Figure 23: BGP configuration and implementation flow



sw0492

5.18 Configuration notes

This section describes BGP configuration caveats.

5.18.1 General

- Before BGP can be configured, the router ID (a valid host address, not the MAC address default) and autonomous system global parameters must be configured.
- BGP instances must be explicitly created on each BGP peer. There are no default BGP instances on a 7210 SAS.

5.18.1.1 BGP defaults

The following list summarizes the BGP configuration defaults:

- By default, the 7210 SAS is not assigned to an AS.
- A BGP instance is created in the administratively enabled state.
- A BGP group is created in the administratively enabled state.
- A BGP neighbor is created in the administratively enabled state.
- No BGP router ID is specified. If no BGP router ID is specified, BGP uses the router system interface address.
- The 7210 SAS BGP timer defaults are the values recommended in IETF drafts and RFCs (see [BGP MIB notes](#))
- If no import route policy statements are specified, all BGP routes are accepted.
- If no export route policy statements specified, all best and used BGP routes are advertised and non-BGP routes are not advertised.

5.18.1.2 BGP MIB notes

The 7210 SAS implementation of the RFC 1657 MIB variables listed in the following table differs from the IETF MIB specification.

Table 62: 7210 SAS and IETF MIB variations

MIB variable	Description	RFC 1657 allowed values	Allowed values
bgpPeerMinASOriginationInterval	Time interval in seconds for the MinASOriginationInterval timer. The suggested value for this timer is 15 seconds.	1 to 65535	2 to 255
bgpPeerMinRouteAdvertisementInterval	Time interval in seconds for the MinRouteAdvertisementInterval timer. The suggested value for this timer is 30.	1 to 65535	1 to 255 ¹⁴

¹⁴ A value of 0 is supported when the rapid-update command is applied to an address family that supports it.

If SNMP is used to set a value of X to the MIB variable in the following table, there are three possible results:

Table 63: MIB variable with SNMP

Condition	Result
X is within IETF MIB values and X is within 7210 SAS values	SNMP set operation does not return an error MIB variable set to X
X is within IETF MIB values and X is outside 7210 SAS values	SNMP set operation does not return an error MIB variable set to "nearest" 7210 SAS supported value (for example, 7210 SAS range is 2 - 255 and X = 65535, MIB variable will be set to 255) Log message generated
X is outside IETF MIB values and X is outside 7210 SAS values	SNMP set operation returns an error

When the value set using SNMP is within the IETF allowed values and outside the 7210 SAS values as specified in the preceding tables, a log message is generated. The log messages that display are similar to the following log messages:

Example: Log message for setting bgpPeerMinASOriginationInterval to 65535

```
576 2006/11/12 19:45:48 [Snmpd] BGP-4-
bgpVariableRangeViolation: Trying to set bgpPeerMinASOrigInt to 65535 -
valid range is [2-255] - setting to 255
```

Example: Log message for setting bgpPeerMinASOriginationInterval to 1

```
594 2006/11/12 19:48:05 [Snmpd] BGP-4-
bgpVariableRangeViolation: Trying to set bgpPeerMinASOrigInt to 1 -
valid range is [2-255] - setting to 2
```

Example: Log message for setting bgpPeerMinRouteAdvertisementInterval to 256

```
535 2006/11/12 19:40:53 [Snmpd] BGP-4-
bgpVariableRangeViolation: Trying to set bgpPeerMinRouteAdvInt to 256 -
valid range is [2-255] - setting to 255
```

Example: Log message for setting bgpPeerMinRouteAdvertisementInterval to 1

```
566 2006/11/12 19:44:41 [Snmpd] BGP-4-
bgpVariableRangeViolation: Trying to set bgpPeerMinRouteAdvInt to 1 -
valid range is [2-255] - setting to 2
```

5.19 Configuring BGP with CLI

This section provides information to configure BGP using the command line interface.

5.20 BGP configuration overview

5.20.1 Preconfiguration requirements

Before BGP can be implemented, the following entities must be configured:

- The autonomous system (AS) number for the router.

An ASN is a globally unique value which associates a router to a specific autonomous system. This number is used to exchange exterior routing information with neighboring ASs and as an identifier of the AS itself. Each router participating in BGP must have an ASN specified.

To implement BGP, the ASN must be specified in the **config>router** context.

- Router ID — The router ID is the IP address of the local router. The router ID identifies a packet's origin. The router ID must be a valid host address.

5.20.2 BGP hierarchy

BGP is configured in the **config>router>bgp** context. Three hierarchical levels are included in BGP configurations:

- Global level
- Group level
- Neighbor level

Commands and parameters configured on the global level are inherited to the group and neighbor levels although parameters configured on the group and neighbor levels take precedence over global configurations.

5.20.3 Internal and external BGP configurations

A BGP system is comprised of ASs which share network reachability information. Network reachability information is shared with adjacent BGP systems neighbors. Further logical groupings are established within BGP systems within ASs. BGP supports two types of routing information exchanges:

- External BGP (EBGP) is used between ASs.

EBGP speakers peer to different ASs and typically share a subnet. In an external group, the next hop is dependent upon the interface shared between the external peer and the specific neighbor. The **multihop** command must be specified if an EBGP peer is more than one hop away from the local router. The next hop to the peer must be configured so that the two systems can establish a BGP session.

- Internal BGP (IBGP) is used within an AS.

An IBGP speaker peers to the same AS and typically does not share a subnet. Neighbors do not have to be directly connected to each other. Since IBGP peers are not required to be directly connected, IBGP uses the IGP path (the IP next-hop learned from the IGP) to reach an IBGP peer for its peering connection.

5.21 Basic BGP configuration

This section provides information to configure BGP and configuration examples of common configuration tasks. The minimal BGP parameters that need to be configured are:

- An autonomous system number for the router.
- A router ID - Note that if a new or different router ID value is entered in the BGP context, then the new value takes precedence and overwrites the router-level router ID.
- A BGP peer group.
- A BGP neighbor with which to peer.
- A BGP peer-AS that is associated with the preceding peer.

The BGP configuration commands have three primary configuration levels: **bgp** for global configurations, **group name** for BGP group configuration, and **neighbor ip-address** for BGP neighbor configuration. Within the different levels, many of the configuration commands are repeated. For the repeated commands, the command that is most specific to the neighboring router is in effect, that is, neighbor settings have precedence over group settings which have precedence over BGP global settings.

Example

The following is a sample configuration that includes the preceding parameters. The following parameters are optional.

```
info
#-----
echo "IP Configuration"
#-----
...
    autonomous-system 200
    router-id 10.10.10.103
#-----
...
#-----
echo "BGP Configuration"
#-----
    bgp
        exit

        export "direct2bgp"
        router-id 10.0.0.12
        group "To_AS_10000"
            connect-retry 20
            hold-time 90
            keepalive 30
            local-preference 100
            remove-private
            peer-as 10000
            neighbor 10.0.0.8
                description "To_Router B - EBGP Peer"
                connect-retry 20
```



```
        hold-time 90
        keepalive 30
        local-address 10.0.0.12
        passive
        preference 99
        peer-as 10000
    exit
exit
group "To_AS_30000"
    connect-retry 20
    hold-time 90
    keepalive 30
    local-preference 100
    remove-private
    peer-as 30000
    neighbor 10.0.3.10
        description "To_Router C - EBGp Peer"
        connect-retry 20
        hold-time 90
        keepalive 30
        peer-as 30000
    exit
exit
group "To_AS_40000"
    connect-retry 20
    hold-time 30
    keepalive 30
    local-preference 100
    peer-as 65206
    neighbor 10.0.0.15

description "To_Router E - Sub Confederation AS 65205"
    connect-retry 20
    hold-time 90
    keepalive 30
    local-address 10.0.0.12
    peer-as 65205
    exit
exit
exit
#-----
....
A:ALA-48>config>router#
```

5.22 Common configuration tasks

About this task

This section provides a brief overview of the tasks that must be performed to configure BGP and provides the CLI commands. To enable BGP, one AS must be configured and at least one group must be configured which includes neighbor (system or IP address) and peering information ASN.

Configure BGP hierarchically, the global level (applies to all peers), the group level (applies to all peers in peer-group), or the neighbor level (only applies to specified peer). By default, group members inherit the group's configuration parameters although a parameter can be modified on a per-member basis without affecting the group-level parameters.

Many of the hierarchical BGP commands can be used on different levels. The most specific value is used. That is, a BGP group-specific command takes precedence over a global BGP command. A neighbor-specific statement takes precedence over a global BGP or group-specific command.

All BGP instances must be explicitly created on each node. When created, BGP is administratively enabled.

Configuration planning is essential to organize ASs and the 7210 SAS nodes within the ASs, and determine the internal and external BGP peering.

To configure a basic autonomous system, perform the following tasks:

Procedure

- Step 1.** Prepare a plan detailing the autonomous system, the 7210 SAS node belonging to each group, group names, and peering connections.
- Step 2.** Associate each 7210 SAS node with an autonomous system number.
- Step 3.** Configure each 7210 SAS node with a router ID.
- Step 4.** Associate each 7210 SAS node with a peer group name.
- Step 5.** Specify the local IP address that will be used by the group or neighbor when communicating with BGP peers.
- Step 6.** Specify neighbors.
- Step 7.** Specify the autonomous system number associated with each neighbor.

5.22.1 Creating an autonomous system

Before BGP can be configured, the autonomous system must be configured first. In BGP, routing reachability information is exchanged between autonomous systems (ASs). An AS is a group of networks that share routing information. The **autonomous-system** command associates an autonomous system number to the router being configured. A 7210 SAS device can only belong to one AS. The **autonomous-system** command is configured in the **config>router** context.

Use the following syntax to associate a 7210 SAS device to an autonomous system.

```
config>router# autonomous-system autonomous-system
```

The 7210 SAS device supports 4 bytes AS numbers by default. This means autonomous-system can have any value from 1 to 4294967295.

Example: Command usage to configure the autonomous system

```
config>router# autonomous-system 100
```

Example: Autonomous system configuration output

```
ALA-B>config>router# info
#-----
# IP Configuration
#-----
      interface "system"
        address 10.10.10.104/32
      exit
```

```
interface "to-103"
  address 10.0.0.104/24
  port 1/1/1
exit
autonomous-system 100

#-----
ALA-B>config>router#
```

5.22.2 Configuring a router ID

In BGP, routing information is exchanged between autonomous systems. The BGP router ID, expressed like an IP address, uniquely identifies the router. It can be set to be the same as the loopback address.

Note that if a new or different router ID value is entered in the BGP context, then the new router ID value is used instead of the router ID configured on the router level, system interface level, or inherited from the MAC address. The router-level router ID value remains intact. A router ID can be derived by:

- Defining the value in the **config>router router-id** context.
- Defining the system interface in the **config>router>interface ip-int-name** context.
- Inheriting the last four bytes of the MAC address.
- The BGP protocol level. The router ID can be defined in the **config>router>bgp router-id** context and is only used within BGP.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is (re) 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 CLI syntax to configure the router ID:

```
config>router# router-id router-id
```

Example: Command usage to configure router ID

```
config>router# router-id 10.10.10.104
```

Example: Router ID configuration output

```
ALA-B>config>router# info
#-----
# IP Configuration
#-----
interface "system"
  address 10.10.10.104/32
exit
interface "to-103"
  address 10.0.0.104/24
  port 1/1/1
exit
autonomous-system 100
router-id 10.10.10.104
#-----
...
ALA-B>config>router#
```

5.22.3 BGP components

The following section describes the syntax used to configure the BGP components.

5.22.4 Configuring BGP

When the BGP protocol instance is created, the **no shutdown** command is not required because BGP is administratively enabled upon creation. Minimally, to enable BGP on a router, you must associate an autonomous system number for the router, have a preconfigured router ID or system interface, create a peer group, neighbor, and associate a peer ASN. There are no default groups or neighbors. Each group and neighbor must be explicitly configured.

All parameters configured for BGP are applied to the group and are inherited by each peer, but a group parameter can be overridden on a specific basis. BGP command hierarchy consists of three levels:

- The global level
- The group level
- The neighbor level

```
config>router# bgp (global level)
               group (group level)
               neighbor (neighbor level)
```



Note:

Careful planning is essential to implement commands that can affect the behavior of global, group, and neighbor levels. Because the BGP commands are hierarchical, analyze the values that can disable features on a particular level.

Example: Basic BGP configuration output

```
ALA-B>config>router# info
#-----
# BGP Configuration
#-----
# BGP
#-----

      bgp
      exit

#-----
ALA-B>config>router#
```

5.22.5 Configuring group attributes

A group is a collection of related BGP peers. The group name should be a descriptive name for the group. Follow your group, name, and ID naming conventions for consistency and to help when troubleshooting faults.

All parameters configured for a peer group are applied to the group and are inherited by each peer (neighbor), but a group parameter can be overridden on a specific neighbor-level basis.

Example: BGP group configuration output

```
ALA-B>config>router>bgp# info
-----
...
    group "headquarters1"
      description "HQ execs"
      local-address 10.0.0.104
      disable-communities standard extended
      ttl-security 255
      exit
    exit
...
ALA-B>config>router>bgp#
```

5.22.6 Configuring neighbor attributes

After you create a group name and assign options, add neighbors within the same autonomous system to create IBGP connections and/or neighbors in different autonomous systems to create EBGP peers. All parameters configured for the peer group level are applied to each neighbor, but a group parameter can be overridden on a specific neighbor basis.

Example

The following is a sample output for neighbors configured in group "headquarters1".

```
ALA-B>config>router>bgp# info
-----
...
    group "headquarters1"
      description "HQ execs"
      local-address 10.0.0.104
      disable-communities standard extended
      ttl-security 255
      neighbor 10.0.0.5
        passive
        peer-as 300
      exit
      neighbor 10.0.0.106
        peer-as 100
      exit
      neighbor 17.5.0.2
        hold-time 90
        keepalive 30
        min-as-origination 15
        local-preference 170
        peer-as 10701
      exit
      neighbor 17.5.1.2
        hold-time 90
        keepalive 30
        min-as-origination 15
        local-preference 100
        min-route-advertisement 30
        preference 170
    exit
...
ALA-B>config>router>bgp#
```

```

        peer-as 10702
    exit
exit
...
-----
ALA-B>config>router>bgp#

```

5.22.7 Configuring AIGP

The AIGP metric is an optional, non-transitive attribute that can be attached to selected routes using route policies. In networks that use AIGP, BGP paths with a lower end-to-end IGP cost are preferred, even if the compared paths span more than one AS or IGP instance.

AIGP is supported only in the base router BGP instance and only for label-IPv4 and 6PE routes. The AIGP attribute is only sent to peers configured using the **configure>router>bgp>group>aigp** and **configure>router>bgp>group>neighbor>aigp** commands.

Example: BGP policy configuration output

The following is a sample BGP policy configuration output with AIGP attribute information included.

```

*A:Dut-C>config>router>policy-options# info
-----
    policy-statement "AIGP_ADD"
        description "Policy From bgp To bgp"
        entry 10
            description "Entry 10 - From Prot. bgp To bgp"
            from
                protocol bgp
            exit
            to
                protocol bgp
            exit
            action accept
                aigp-metric add 555
            exit
        exit
    exit
    policy-statement "AIGP_EXPORT_PLCY"
        description "Policy From bgp To bgp"
        entry 10
            description "Entry 10 - From Prot. bgp To bgp"
            from
                protocol bgp
            exit
            to
                protocol bgp
            exit
            action accept
                next-hop 10.20.1.3
            exit
        exit
    exit
-----

```

Example: BGP instance configuration output

The following is a sample BGP instance configuration output with AIGP attribute information included.

```

*A:Dut-C>config>router>bgp# info

```

```
-----
min-route-advertisement 1
router-id 10.20.1.3
group "PEER_TO_A"
    neighbor 10.10.1.1
        local-address 10.10.1.3
        peer-as 200
        advertise-label ipv4
    exit
exit
group "PEER_RR_TO_D_E_B"
    cluster 10.20.1.3
    aigp
    neighbor 10.20.1.2
        local-address 10.20.1.3
        med-out 100
        import "AIGP_ADD"
        peer-as 300
        advertise-label ipv4
    exit
    neighbor 10.20.1.4
        local-address 10.20.1.3
        med-out 100
        peer-as 300
        advertise-label ipv4
    exit
    neighbor 10.20.1.5
        local-address 10.20.1.3
        export "AIGP_EXPORT_PLCY"
        peer-as 300
        advertise-label ipv4
    exit
exit
no shutdown
-----
```

5.23 BGP configuration management tasks

This section describes the BGP configuration management tasks.

5.23.1 Modifying an ASN

You can modify an ASN on a 7210 SAS but the new ASN will not be used until the BGP instance is restarted either by administratively disabling or enabling the BGP instance or by rebooting the system with the new configuration.

Since the ASN is defined in the **config>router** context, not in the BGP configuration context, the BGP instance is not aware of the change. Re-examine the plan detailing the autonomous system the SRs belonging to each group, group names, and peering connections. Changing an ASN on a 7210 SAS could cause configuration inconsistencies if associated **peer-as** values are not also modified as required. At the group and neighbor levels, BGP will re-establish the peer relationships with all peers in the group with the new ASN.

Use the following syntax to change an ASN.

```
config>router# autonomous-system autonomous-system
```

```
config>router# bgp
  group name
  neighbor ip-addr
  peer-as asn
```

Example: Command usage to change an ASN

```
config>router# autonomous-system 400
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# neighbor 10.10.10.103
config>router>bgp>group# peer-as 400
config>router>bgp>group# exit
```

5.23.2 Modifying the BGP router ID

Changing the router ID number in the BGP context causes the new value to overwrite the router ID configured on the router level, system interface level, or the value inherited from the MAC address. Changing the router ID on a router could cause configuration inconsistencies if associated values are not also modified.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time BGP is (re) initialized the new router ID is used. To force the new router ID, issue the **shutdown** and **no shutdown** commands for BGP or restart the entire router.

Example: Command usage to configure a new router ID

```
config>router>bgp# router-id 10.0.0.104
config>router>bgp# shutdown
config>router>bgp# router-id 10.0.0.123
config>router>bgp# no shutdown
```

Example: BGP configuration output

The following is a sample BGP configuration output with the BGP router ID specified.

```
ALA-B>config>router>bgp# info detail
-----
no shutdown
no description
no always-compare-med
ibgp-multipath load-balance
. . .
router-id 10.0.0.123
-----
ALA-B>config>router>bgp#
```


5.23.3 Modifying the router-level router ID

Changing the router ID number in the **config>router** context causes the new value to overwrite the router ID configured on the protocol level, system interface level, or the value inherited from the MAC address. Changing the router ID on a router could cause configuration inconsistencies if associated values are not also modified.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is (re) 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 change a router ID.

```
config>router# router-id router-id
```

Example: Command usage to change a router ID

```
config>router# router-id 10.10.10.104
config>router# no shutdown
config>router>bgp# shutdown
config>router>bgp# no shutdown
```

Example: Router ID configuration output

```
ALA-A>config>router# info
#-----
# IP Configuration
#-----
    interface "system"
        address 10.10.10.104/32
    exit
    interface "to-103"
        address 10.0.0.104/24
        port 1/1/1
    exit
    autonomous-system 100
    router-id 10.10.10.104

#-----
ALA-B>config>router#
```

5.23.4 Deleting a neighbor

To delete a neighbor, you must shut down the neighbor before issuing the **no neighbor ip-addr** command.

Use the following syntax to delete a neighbor.

```
config>router# bgp
group name
no neighbor ip-address
shutdown
    no peer-as asn
shutdown
```

Example: Command usage to delete a neighbor

```
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# neighbor 10.0.0.103
config>router>bgp>group>neighbor# shutdown
config>router>bgp>group>neighbor# exit
config>router>bgp>group# no neighbor 10.0.0.103
```

Example: Configuration output

The following is a sample of the "headquarters1" configuration output with the neighbor 10.0.0.103 removed.

```
ALA-B>config>router>bgp# info
-----
description      group "headquarters1"
                  "HQ execs"
                  local-address 10.0.0.104
                  neighbor 10.0.0.5
                  passive
                  peer-as 300
                  exit
exit
-----
ALA-B>config>router>bgp#
```

5.23.5 Deleting groups

To delete a group, the neighbor configurations must be shut down first. After each neighbor is shut down, you must shut down the group before issuing the **no group name** command.

Use the following syntax to shut down a peer and neighbor and then delete a group.

```
config>router# bgp
no group name
shutdown
no neighbor ip-address
shutdown
shutdown
```

Example: Command usage to delete a group

```
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# neighbor 10.0.0.105
config>router>bgp>group>neighbor# shutdown
config>router>bgp>group>neighbor# exit
config>router>bgp>group# neighbor 10.0.0.103
config>router>bgp>group# shutdown
config>router>bgp>group# exit
config>router>bgp# no headquarters1
```

If you try to delete the group without shutting down the peer-group, the following message appears.

```
ALA-B>config>router>bgp# no group headquarters1
MINOR: CLI BGP Peer Group should be shutdown before deleted. BGP Peer Group not
```

```
deleted.
```

5.23.6 Editing BGP parameters

You can change existing BGP parameters in the CLI using the following syntax. The changes are applied immediately.

```
config>router# bgp
group name
. . .
neighbor ip-address
. . .
```

Example

```
config>router# bgp
```

See [BGP components](#) for a complete list of BGP parameters.

5.24 BGP command reference

5.24.1 Command hierarchies

5.24.1.1 Configuration commands

- [Global BGP commands](#)
- [Group BGP commands](#)
- [Neighbor BGP commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

5.24.1.1.1 Global BGP commands

```
config
- router [router-name]
- [no] bgp
- [no] add-paths
- ipv4 send send-limit receive [none]
- ipv4 send send-limit
- no ipv4
- ipv6 send send-limit receive [none]
- ipv6 send send-limit
- no ipv6
- [no] advertise-inactive
- [no] aggregator-id-zero
```

```
- no as-path-ignore
- authentication-key [authentication-key | hash-key] [hash | hash2]
- auth-keychain name
- no authentication-key
- [no] backup-path [ipv4] [ipv6]
- best-path-selection
  - always-compare-med {zero | infinity}
  - always-compare-med strict-as {zero | infinity}
  - no always-compare-med
  - as-path-ignore [ipv4] [vpn-ipv4]
  - no as-path-ignore
  - ignore-nh-metric
  - no ignore-nh-metric
  - ignore-router-id
  - no ignore-router-id
- [no] bfd-enable
- connect-retry seconds
- no connect-retry
- [no] damping
- description description-string
- no description
- [no] disable-4byte-asn
- disable-communities [standard] [extended]
- no disable-communities
- [no] disable-fast-external-failover
- [no] enable-peer-tracking
- export policy-name [policy-name...(up to 15 max)]
- no export
- family [ipv4] [vpn-ipv4] [ipv6] [vpn-ipv6] [l2-vpn] [route-target]
- no family
- hold-time seconds [strict]
- no hold-time
- import policy-name [policy-name ...(up to 15 max)]
- no import
- keepalive seconds
- no keepalive
- local-preference local-preference
- no local-preference
- loop-detect {drop-peer | discard-route | ignore-loop | off}
- no loop-detect
- med-out {number | igp-cost}
- no med-out
- min-as-origination seconds
- no min-as-origination
- min-route-advertisement seconds
- no min-route-advertisement
- multihop ttl-value
- no multihop
- next-hop-resolution
  - label-route-transport-tunnel
    - family family
      - resolution {any | filter | disabled}
      - resolution-filter
        - [no] ldp
        - [no] rsvp
        - [no] sr-isis
        - [no] sr-ospf
- [no] outbound-route-filtering
  - [no] extended-community
    - [no] accept-orf
    - send-orf [comm-id...(up to 32 max)]
    - no send-orf comm-id
- [no] path-mtu-discovery
- preference preference
```

```
- purge-timer
- no purge-timer
- no preference
- rapid-update [l2-vpn] [evpn]
- no rapid-update
- [no] rapid-withdrawal
- [no] remove-private {limited}
- router-id ip-address
- no router-id
- [no] shutdown
- [no] vpn-apply-export
- [no] vpn-apply-import
```

5.24.1.1.2 Group BGP commands

```
config
- router [router-name]
- [no] bgp
- [no] group name
- [no] add-paths
-   - ipv4 send send-limit receive [none]
-   - ipv4 send send-limit
-   - no ipv4
-   - ipv6 send send-limit receive [none]
-   - ipv6 send send-limit
-   - no ipv6
- [no] advertise-inactive
- [no] aggregator-id-zero
- [no] aigp
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- auth-keychain name
- [no] bfd-enable
- connect-retry seconds
- no connect-retry
- [no] damping
- [no] default-route-target
- description description-string
- no description
- [no] disable-4byte-asn
- [no] disable-capability-negotiation
- disable-communities [standard] [extended]
- no disable-communities
- [no] disable-fast-external-failover
- [no] enable-peer-tracking
- export policy-name [policy-name...(up to 15 max)]
- no export
- family [ipv4] [vpn-ipv4] [ipv6] [vpn-ipv6] [l2-vpn] [route-target]
- no family
- hold-time seconds [strict]
- no hold-time
- import policy-name [policy-name ...(up to 15 max)]
- no import
- keepalive seconds
- no keepalive
- local-address ip-address
- no local-address
- local-as as-number [private]
- no local-as
- local-preference local preference
- no local-preference
```

```
- loop-detect {drop-peer | discard-route | ignore-loop | off}
- no loop-detect
- med-out {number | igp-cost}
- no med-out
- min-as-origination seconds
- no min-as-origination
- min-route-advertisement seconds
- no min-route-advertisement
- multihop ttl-value
- no multihop
- [no] outbound-route-filtering
  - [no] extended-community
    - [no] accept-orf
    - send-orf [comm-id...(up to 32 max)]
    - no send-orf [comm-id]
  - [no] path-mtu-discovery
- peer-as as-number
- no peer-as
- preference preference
- no preference
- prefix-limit family limit [log-only] [threshold percentage]
- no prefix-limit
- [no] remove-private [limited]
- [no] shutdown
- type {internal | external}
- no type
- [no] vpn-apply-export
- [no] vpn-apply-import
```

5.24.1.1.3 Neighbor BGP commands

```
config
- router [router-name]
  - [no] bgp
    - [no] group name
      - [no] neighbor ip-address
        - [no] add-paths
          - ipv4 send send-limit receive [none]
          - ipv4 send send-limit
          - no ipv4
          - ipv6 send send-limit receive [none]
          - ipv6 send send-limit
          - no ipv6
        - [no] advertise-inactive
        - advertise-label [ipv4 [use-svc-routes] [implicit-null]] [ipv6]
        - [no] advertise-label
        - [no] aggregator-id-zero
        - [no] aigp
        - auth-keychain name
        - authentication-key [authentication-key | hash-key] [hash | hash2]
        - no authentication-key
        - [no] bfd-enable
        - connect-retry seconds
        - no connect-retry
        - [no] damping
        - [no] default-route-target
        - description description-string
        - no description
        - [no] disable-4byte-asn
        - [no] disable-capability-negotiation
        - disable-communities [standard] [extended]
```

```
- no disable-communities
- [no] disable-fast-external-failover
- [no] enable-peer-tracking
- export policy-name [policy-name...(up to 15 max)]
- no export
- family [ipv4] [vpn-ipv4] [ipv6] [vpn-ipv6] [l2-vpn] [route-target]
- no family
- hold-time seconds [strict]
- no hold-time
- import policy-name [policy-name ...(up to 15 max)]
- no import
- keepalive seconds
- no keepalive
- local-address ip-address
- no local-address
- local-as as-number [private]
- no local-as
- local-preference local-preference
- no local-preference
- loop-detect {drop-peer | discard-route | ignore-loop | off}
- no loop-detect
- med-out {number | igp-cost}
- no med-out
- min-as-origination seconds
- no min-as-origination
- min-route-advertisement seconds
- no min-route-advertisement
- multihop ttl-value
- no multihop
- [no] outbound-route-filtering
  - [no] extended-community
    - [no] accept-orf
      - send-orf [comm-id...(up to 32 max)]
      - no send-orf [comm-id]
  - [no] path-mtu-discovery
- peer-as as-number
- no peer-as
- preference preference
- no preference
- prefix-limit limit
- no prefix-limit
- [no] remove-private {limited}
- [no] shutdown
- type {internal | external}
- no type
- [no] vpn-apply-export
- [no] vpn-apply-import
```

5.24.1.1.4 Other BGP-related commands

```
config
- router [router-name]
  - autonomous-system as-number
  - no autonomous-system
```

5.24.1.2 Show commands

```
show
```

```
- router [router-instance]
- bgp
- auth-keychain keychain
- damping [ip-prefix [/ip-prefix-length]] [damp-type] [detail] ipv4
- damping [ip-prefix [/ip-prefix-length]] [damp-type] [detail] ipv6
- damping [ip-prefix [/ip-prefix-length]] [damp-type] [detail] vpn-ipv4
- damping [ip-prefix [/ip-prefix-length]] [damp-type] [detail] vpn-ipv6
- group [name] [detail]
- neighbor [ip-address [detail]]
- neighbor [as-number [detail]]
- neighbor ip-address [family] filter1 [brief]
- neighbor ip-address [family] filter2
- neighbor as-number [family] filter2
- neighbor ip-address orf filter3
- neighbor ip-address graceful-restart
- next-hop family [ip-address] [detail]
- paths
- routes [family] [brief]
- routes [family] prefix [detail | longer | hunt [brief]]
- routes [family [type mvpn-type]] community comm-id
- routes [family [type mvpn-type]] aspath-regex reg-ex
- routes ms-pw [rd rd] [aii-type2 aii-type2] [brief]
- routes l2-vpn l2vpn-type [{rd rd} | [siteid site-id] | [veid veid] [offset vpls-
base-offset]]
- routes evpn auto-disc [hunt | detail] [rd rd] [community comm-id] [tag tag]
[next-hop ip-address] [esi esi]
- routes evpn eth-seg [hunt | detail] [rd rd] [community comm-id] [originator-
ip ip-address] [next-hop ip-address] [esi esi]
- routes evpn inclusive-mcast [hunt | detail] [rd rd] [community comm-id]
[originator-ip ip-address] [next-hop ip-address] [esi esi]
- routes evpn inclusive-mcast [hunt | detail] [rd rd] [community comm-id]
[originator-ip ip-address] [next-hop ip-address] [tag tag]
- routes evpn mac [hunt | detail] [rd rd] [next-hop ip-address] [mac-address mac-
address] [community comm-id] [tag tag]
- summary [all]
- summary [family family] [neighbor ip-address]
```

5.24.1.3 Clear commands

```
clear
- router
- bgp
- damping [{ip-prefix/ip-prefix-length} [neighbor ip-address]] | {group name}
- flap-statistics [{ip-prefix/mask [neighbor ip-address]} | [group group-name] |
[regex reg-exp | policy policy-name]]
- neighbor {ip-address | as as-number | external | all} [soft | soft-inbound]
- neighbor {ip-address | as as-number | external | all} statistics
- neighbor ip-address end-of-rib
- protocol
```

5.24.1.4 Debug commands

```
debug
- router
- bgp
- events [neighbor ip-address | group name]
- no events
- keepalive [neighbor ip-address | group name]
```



```
- no keepalive
- notification [neighbor ip-address | group name]
- no notification
- open [neighbor ip-address | group name]
- no open
- [no] outbound-route-filtering
- packets [neighbor ip-address | group name]
- no packets
- route-refresh [neighbor ip-address | group name]
- no route-refresh
- rtm [neighbor ip-address | group name]
- no rtm
- socket [neighbor ip-address | group name]
- no socket
- timers [neighbor ip-address | group name]\
- no timers
- update [neighbor ip-address | group name]
- no update
```

5.24.2 Command descriptions

5.24.2.1 Configuration commands

bgp

Syntax

[no] bgp

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the BGP protocol instance and BGP configuration context. BGP is administratively enabled upon creation.

The **no** form of this command deletes the BGP protocol instance and removes all configuration parameters for the BGP instance. BGP must be **shutdown** before deleting the BGP instance. An error occurs if BGP is not **shutdown** first.

add-paths

Syntax

[no] add-paths

Context

```
config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure multiple paths for one or more families in a BGP instance, group, or neighbor. The BGP add-path capability allows the router to send and receive multiple paths per prefix to and from a peer.

The **no** form of this command removes the add-path capability from the BGP instance, group, or neighbor, causing sessions established using **add-paths** to go down and come back up without the add-path capability.

Default

no add-paths

ipv4

Syntax

```
ipv4 send send-limit receive [none]
ipv4 send send-limit
no ipv4
```

Context

```
config>router>bgp>add-paths
config>router>bgp>group>add-paths
config>router>bgp>group>neighbor>add-paths
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the add-path capability for IPv4 labeled routes; the add-path capability is disabled by default.



Note:

The add-path capability is not supported for IPv4 native routes (that is, IPv4 routes without a label).

The maximum number of paths to send per IPv4 NLRI is configured using the *send-limit* mandatory parameter. The capability to receive multiple paths per prefix from a peer is configured using the optional **receive** keyword. If the **receive** keyword is not specified, the receive capability is enabled by default.

The **no** form of this command disables add-path support for IPv4 routes, causing sessions established using add-paths for IPv4 to go down and come back up without the add-paths capability.

Default

no ipv4

Parameters

send-limit

Specifies the maximum number of paths per IPv4 NLRI that can be advertised to add-path peers. The actual number of advertised routes may be less depending on the next-hop diversity requirement, other configuration options, route policies, and route advertisement rules.

Values 1 to 16, none

Default max

receive

Specifies that the router negotiates the add-paths receive capability for IPv4 routes with its peers.

none

Specifies that the router does not negotiate the add-paths receive capability for IPv4 routes with its peers.

ipv6

Syntax

ipv6 send *send-limit* **receive** [none]

ipv6 send *send-limit*

no ipv6

Context

config>router>bgp>add-paths

config>router>bgp>group>add-paths

config>router>bgp>group>neighbor>add-paths

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the add-path capability for IPv6 labeled routes; the capability is disabled by default.



Note:

The add-path capability is not supported for IPv6 native routes (that is, IPv6 routes without a label).

The maximum number of paths to send per IPv6 NLRI is configured using the *send-limit* mandatory parameter. The capability to receive multiple paths per prefix from a peer is configured using the optional **receive** keyword. If the **receive** keyword is not specified, the receive capability is enabled by default.

The **no** form of this command disables add-path support for IPv6 routes, and causes sessions established using **add-paths** for IPv6 to go down and come back up without the add-path capability.

Default

no ipv6

Parameters

send-limit

Specifies the maximum number of paths per IPv6 NLRI that can be advertised to add-path peers. The actual number of advertised routes may be less depending on the next-hop diversity requirement, other configuration options, route policies, and route advertisement rules.

Values 1 to 16, none

Default max

receive

Specifies that the router negotiates the add-paths receive capability for IPv6 routes with its peers.

none

Specifies that the router does not negotiate the add-paths receive capability for IPv6 routes with its peers.

advertise-inactive

Syntax

[no] **advertise-inactive**

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the advertising of inactive BGP routes to other BGP peers. By default, BGP only advertises BGP routes to other BGP peers if a specific BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a specific destination.

The **no** form of this command disables the advertising of inactive BGP routers to other BGP peers.

Default

no advertise-inactive

advertise-label

Syntax

advertise-label [ipv4 [use-svc-routes] [implicit-null]] [ipv6]

no advertise-label

Context

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the IPv4 or IPv6 transport peers to exchange RFC 3107-labeled IPv4 or IPv6 routes.

If IPv4 is enabled, all IPv4 routes advertised to the remote BGP peer are sent with an RFC 3107-formatted label for the destination route.

If IPv6 is enabled, all IPv6 routes advertised to the remote BGP peer are sent with an RFC 3107-formatted label for the destination route.

The optional keyword **use-svc-routes** allows the user to limit the number of BGP 3107 IPv4 labeled routes that are installed in the MPLS FIB. If the keyword is specified, only BGP 3107 labeled routes that are required by services or required for establishing a BGP session with a configured neighbor are installed in the MPLS FIB. The following will trigger installation of the MPLS label into the MPLS FIB for the received BGP 3107 IPv4 labeled route:

- configuration of SDP to use BGP tunnel to the far-end
- dynamic creation of spoke-SDP binding when a route is received through BGP AD and the far-end of the SDP binding is reachable using the labeled route
- installation of VPN IPv4 routes received from the PE, which is reachable using the labeled route

- configuration of the BGP session to a BGP peer using the **bgp>neighbor** CLI command, and the BGP peer is reachable using the labeled route
- Other IP applications such as FTP, SSH, and other applications do not trigger installation of the IPv4 labeled routes into the MPLS FIB.

The **no** form of this command disables any or all configured options.

Default

no advertise-label

Parameters

ipv4

Specifies the advertisement label address family for core IPv4 routes. This keyword can be specified only for an IPv4 peer.

use-svc-routes

Optional keyword that allows the user to limit the number of BGP 3107 labeled routes that are installed in the MPLS FIB. If it is specified, only BGP 3107 labeled routes that are required by services configured in the system, or required for establishing a BGP session with a configured neighbor, are installed in the MPLS FIB.

implicit-null

Keyword to enable or disable the use of the implicit null label for BGP-LU routes advertised to the neighbor.

When enabled, BGP on the node advertises implicit null for all IPv4 /32 prefix routes, for which the node is the egress LER, along with the route advertised to the BGP neighbor.

When disabled, BGP does not advertise implicit null for any of the IPv4 routes.

ipv6

Keyword to specify the advertisement label address family for core IPv6 routes. This keyword can be specified only for an IPv6 peer.

aggregator-id-zero

Syntax

[no] aggregator-id-zero

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths.

When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the ASN and router ID to the aggregator path attribute.

When this command is enabled, BGP adds the router ID to the aggregator path attribute. This command is used at the group level to revert to the value defined under the global level. This command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of this command used at the global level reverts to the default, where BGP adds the ASN and router ID to the aggregator path attribute.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no aggregator-id-zero

aigp

Syntax

[no] aigp

Context

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables AIGP path attribute support with one or more BGP peers. BGP path selection among routes with an associated AIGP metric is based on the end-to-end IGP metrics of the different BGP paths, even when these BGP paths span more than one AS and IGP instance.

The **no** form of this command disables AIGP path attribute support, removes the AIGP attribute from advertised routes, and causes the AIGP attribute in received routes to be ignored.

Default

no aigp

auth-keychain

Syntax

auth-keychain *name*

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a TCP authentication keychain to use for the session. The keychain allows the rollover of authentication keys during the lifetime of a session.

Default

no auth-keychain

Parameters

name

Specifies the name of the keychain, up to 32 characters, to use for the specified TCP session or sessions.

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the BGP authentication key.

Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD-5 message based digest. MD5 authentication is disabled by default.

The authentication *key* can be any combination of ASCII characters up to 255 characters.

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

Parameters

authentication-key

Specifies the authentication key. The key can be any combination of ASCII characters up to 255 characters (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 342 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

backup-path

Syntax

[no] backup-path [ipv4] [ipv6]

Context

config>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables BGP Fast Reroute (FRR) with Prefix-Independent Convergence (PIC), allowing for the creation of a backup path for labeled IPv4 and IPv6 BGP learned prefixes belonging to the base router. Multiple paths must be received for a prefix to take advantage of this feature.

If a prefix has a backup path and its primary paths fail, the affected traffic is rapidly diverted to the backup path without waiting for control plane reconvergence to occur. If many prefixes share the same primary

paths, and in some cases also the same backup path, the time to failover traffic to the backup path is independent of the number of prefixes.

By default, IPv4 and IPv6 prefixes do not have a backup path installed in the IOM.

The **no** form of this command disables BGP FRR with PIC.

Default

no backup-path

Parameters

ipv4

Enables BGP fast reroute for labeled IPv4 routes.

ipv6

Enables BGP fast reroute for labeled IPv6 routes.

best-path-selection

Syntax

best-path-selection

Context

config>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables path selection configuration.

always-compare-med

Syntax

always-compare-med {zero | infinity}

no always-compare-med strict-as {zero | infinity}

no always-compare-med

Context

config>router>bgp>best-path-selection

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the comparison of BGP routes based on the MED attribute.

The default behavior of 7210 SAS is to compare two routes on the basis of MED only if they have the same neighbor AS (the first non-confed AS in the received AS_PATH attribute). By default, a route without a MED attribute is handled the same as though it had a MED attribute with the value 0.

The **always-compare-med** command without the **strict-as** keyword allows MED to be compared even if the paths have a different neighbor AS. In this case, if neither **zero** nor **infinity** is specified, the **zero** option is inferred, meaning a route without a MED is handled the same as though it had a MED attribute with the value 0. When the **strict-as** keyword is configured, MED is only compared between paths from the same neighbor AS, and in this case, **zero** or **infinity** is mandatory and tells BGP how to interpret paths without a MED attribute.

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

Default

no always-compare-med

Parameters

zero

Specifies that for routes learned without a MED attribute, a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred.

infinity

Specifies that for routes learned without a MED attribute, a value of infinity ($2^{32}-1$) is used in the MED comparison. This, in effect makes, these routes the least desirable.

strict-as

Specifies that the MEDs of two paths are compared even if they come from different neighboring AS.

as-path-ignore

Syntax

as-path-ignore [ipv4] [vpn-ipv4]

no as-path-ignore

Context

config>router>bgp>best-path-selection

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command determines whether the AS path is used to determine the best BGP route.

If this option is present, the AS paths of incoming routes are not used in the route selection process.

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

Default

no as-path-ignore

Parameters

ipv4

Specifies that the AS-path length will be ignored for all IPv4 routes.

vpn-ipv4

Specifies that the AS-path length will be ignored for all IPv4 VPRN routes.

ignore-nh-metric

Syntax

ignore-nh-metric

no ignore-nh-metric

Context

config>router>bgp>best-path-selection

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command instructs BGP to disregard the resolved distance to the BGP next-hop in its decision process for selecting the best route to a destination.

When configured in the **config>router>bgp>best-path-selection** context, this command applies to the comparison of two BGP routes with the same NLRI learned from base router BGP peers. When configured in the **config>service>vpn** context, this command applies to the comparison of two BGP-VPN routes for the same IP prefix imported into the VPRN from the base router BGP instance. When configured in the **config>service>vpn>bgp>best-path-selection** context, this command applies to the comparison of two BGP routes for the same IP prefix learned from VPRN BGP peers.

The **no** form of this command restores the default behavior where BGP factors the distance to the next-hop into its decision process.

Default

no ignore-nh-metric

ignore-router-id

Syntax

ignore-router-id

no ignore-router-id

Context

config>router>bgp>best-path-selection

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command modifies the route selection behavior. When this command is enabled and the current best path to a destination was learned from eBGP peer X with BGP identifier x, new paths that are received from eBGP peer Y with BGP identifier y and are equivalent will not change the best path even if y is less than x during BGP identifier comparison.

The **no** form of this command restores the default behavior of selecting the route with the lowest BGP identifier (Y) as best.

Default

no ignore-router-id

bfd-enable

Syntax

[no] bfd-enable

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables bidirectional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a specific protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

IPv4 BFD can be used for multihop or single hop MP-BGP sessions. For more information about the protocols and platforms that support BFD, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*.

The **no** form of this command removes BFD from the associated IGP/BGP protocol adjacency.

Default

no bfd-enable

connect-retry

Syntax

connect-retry *seconds*

no connect-retry

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the BGP connect retry timer value in seconds.

When this timer expires, BGP tries to reconnect to the configured peer. This configuration parameter can be set at three levels: global level (applies to all peers), peer-group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

connect-retry 120

Parameters

seconds

Specifies the BGP connect retry timer value, in seconds, expressed as a decimal integer.

Values 1 to 65535

damping

Syntax

[no] damping

Context

```
config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables BGP route damping for learned routes, which are defined within the route policy. Use damping to reduce the number of update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set through route policy definition.

When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters:

Half-life:	15 minutes
Max-suppress:	60 minutes
Suppress-threshold:	3000
Reuse-threshold:	750

The **no** form of this command used at the global level reverts route damping.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no damping

default-route-target

Syntax

[no] default-route-target

Context

```
config>router>bgp>group
config>router>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command originates the default RTC route (zero prefix length) toward the selected peers.

The **no** form of this command disables the advertisement of the default RTC route.

Default

no default-route-target

description

Syntax

description *description-string*

no description

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

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

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.

disable-4byte-asn

Syntax

[no] **disable-4byte-asn**

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command disables the use of 4-byte ASNs. It can be configured at all 3 levels of the hierarchy so it can be specified down to the per peer basis.

If this command is enabled, 4-byte ASN support should not be negotiated with the associated remote peers.

The **no** form of this command reverts to the default behavior, which is to enable the use of 4-byte ASN.

disable-capability-negotiation

Syntax

[no] disable-capability-negotiation

Context

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command disables the exchange of capabilities when the command is enabled, after the peering is flapped, any new capabilities are not negotiated and strictly support IPv4 routing exchanges with that peer.

The **no** form of this command removes this command from the configuration and restores the normal behavior.

Default

no disable-capability-negotiation

disable-communities

Syntax

disable-communities [standard] [extended]

no disable-communities

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures BGP to disable sending communities.

Parameters

standard

Specifies standard communities that existed before VPRNs or 2547.

extended

Specifies BGP communities used were expanded after the concept of 2547 was introduced, to include handling the VRF target.

disable-fast-external-failover

Syntax

[no] **disable-fast-external-failover**

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures BGP fast external failover.

disallow-igp

Syntax

[no] **disallow-igp**

Context

config>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or disables the use of the IGP next-hop to the BGP next-hop as the next-hop of the last resort.

```
enable-inter-as-vpn
```

Syntax

```
[no] enable-inter-as-vpn
```

Context

```
config>router>bgp
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies whether VPNs can exchange routes across autonomous system boundaries, providing model B connectivity.

The **no** form of this command disallows ASBRs to advertise VPRN routes to their peers in other autonomous systems.

Default

```
no enable-inter-as-vpn
```

```
enable-peer-tracking
```

Syntax

```
[no] enable-peer-tracking
```

Context

```
config>router>bgp
```

```
config>router>bgp>group
```

```
config>router>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables BGP peer tracking. BGP peer tracking allows a BGP peer to be dropped immediately if the route used to resolve the BGP peer address is removed from the IP routing table and there is no alternative available. The BGP peer will not wait for the holdtimer to expire; therefore, the BGP reconvergence process is accelerated.

The **no** form of this command disables peer tracking.

Default

no enable-peer-tracking

export

Syntax

export *policy-name* [*policy-name...upto 15 max*]

no export [*policy-name*]

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the export route policy used to determine which routes are advertised to peers.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific level is used.

When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be configured. The first policy that matches is applied.

When multiple export commands are issued, the last command entered overrides the previous command.

When no export policies are specified, BGP routes are advertised and non-BGP routes are not advertised by default.

The **no** form of this command removes the policy association with the BGP instance. To remove association of all policies, use the **no export** command without arguments.

Default

no export

Parameters

policy-name

Specifies the route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

family

Syntax

family [ipv4] [vpn-ipv4] [ipv6] [vpn-ipv6] [l2-vpn] [route-target]

no family

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the address families to be supported over BGP peerings in the base router. This command is additive, so issuing the **family** command adds the specified address family to the list.

The **no** form of this command removes the specified address family from the associated BGP peerings. If an address family is not specified, the system resets the supported address family back to the default.

Default

family ipv4

Parameters

ipv4

Exchanges IPv4 routing information.

vpn-ipv4

Exchanges IPv4 VPN routing information.

ipv6

Exchanges IPv6 routing information.

vpn-ipv6

Exchanges IPv6 VPN routing information.

l2-vpn

Exchanges Layer 2 VPN information.

route-target

Keyword to exchange RT constrained route information.

vpn-apply-export

Syntax

[no] vpn-apply-export

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command causes the base instance BGP export route policies to be applied to VPN-IPv4 routes.

The **no** form of this command disables the application of the base instance BGP route policies to VPN-IPv4 routes.

Default

no vpn-apply-export

vpn-apply-import

Syntax

[no] vpn-apply-import

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command applies base instance BGP import route policies to VPN-IPv4 routes.

The **no** form of this command disables the application of the base instance BGP import route policies to VPN-IPv4 routes.

Default

no vpn-apply-import

group

Syntax

[no] group *name*

Context

config>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure a BGP peer group.

The **no** form of this command deletes the specified peer group and all configurations associated with the peer group. The group must be **shutdown** before it can be deleted.

Parameters

name

Specifies the peer group 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.

hold-time

Syntax

hold-time *seconds* [strict]

no hold-time

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the BGP hold time, expressed in seconds.

The BGP hold time specifies the maximum amount of time that BGP waits between successive messages (either **keepalive** or **update**) from its peer before closing the connection.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

Even though the 7210 SAS implementation allows setting the **keepalive** time separately, the configured **keepalive** timer is overridden by the **hold-time** value under the following circumstances.

- If the specified hold-time is less than the configured **keepalive** time, the operational **keepalive** time is set to a third of the **hold-time**; the configured **keepalive** time is not changed.
- If the **hold-time** is set to zero, the operational value of the **keepalive** time is set to zero; the configured **keepalive** time is not changed. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

hold-time 90

Parameters

seconds

Specifies the hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently.

Values 0, 3 to 65535

strict

When this parameter is specified, the advertised BGP hold-time from the far-end BGP peer must be greater than or equal to the specified value.

import

Syntax

import *policy-name* [*policy-name*...up to 15 max]

no import [*policy-name*]

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the import route policy to be used to determine which routes are accepted from peers. Route policies are configured in the **config>router>policy-options** context.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific level is used.

When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be specified. The first policy that matches is applied.

When multiple **import** commands are issued, the last command entered will override the previous command.

When an import policy is not specified, BGP routes are accepted by default.

The **no** form of this command removes the policy association with the BGP instance. To remove association of all policies, use **no import** without arguments.

Default

no import

Parameters

policy-name

Specifies the route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

keepalive

Syntax

keepalive *seconds*

no keepalive

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires.

The **keepalive** parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **keepalive** value is generally one-third of the **hold-time** interval. Even though the 7210 SAS implementation allows the **keepalive** value and the **hold-time** interval to be independently set, under the following circumstances, the configured **keepalive** value is overridden by the **hold-time** value.

- If the specified **keepalive** value is greater than the configured **hold-time**, the specified value is ignored, and the **keepalive** value is set to one third of the current **hold-time** value.
- If the specified **hold-time** interval is less than the configured **keepalive** value, the **keepalive** value is reset to one third of the specified **hold-time** interval.
- If the **hold-time** interval is set to zero, the configured **keepalive** value is ignored. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

keepalive 30

Parameters

seconds

Specifies the keepalive timer, in seconds, expressed as a decimal integer.

Values 0 to 21845

local-address

Syntax

local-address *ip-address*

no local-address

Context

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the local IP address used by the group or neighbor when communicating with BGP peers.

Outgoing connections use the **local-address** as the source of the TCP connection when initiating connections with a peer.

When a local address is not specified, 7210 SAS uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGP peers. This command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of this command removes the configured local-address for BGP.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default
no local-address

Parameters
ip-address
Specifies the local IP address.

Values	
ipv4-address:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 to 32
ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D

local-as

Syntax
local-as *as-number* [**private**]
no local-as

Context
config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Platforms
Supported on all 7210 SAS platforms as described in this document

Description
This command configures a BGP virtual autonomous system (AS) number.

In addition to the ASN configured for BGP in the **config>router>autonomous-system** context, a virtual (local) ASN is configured. The virtual ASN is added to the as-path message before the router ASN makes the virtual AS the second AS in the as-path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). Therefore, by specifying this at each neighbor level, it is possible to have a separate as-number per EBGP session.

When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The **private** attribute can be added or removed dynamically by reissuing the command.

Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local ASN. Changing the local AS at the global level in an active BGP instance causes BGP to re-establish the peer relationships with all peers in the group with the new local ASN. Changing the local AS at the neighbor level in an active BGP instance causes BGP to re-establish the peer relationship with the new local ASN.

This is an optional command and can be used in the following situation.

Example: Provider router P is moved from AS1 to AS2. The customer router that is connected to P, however, is configured to belong to AS1. To avoid reconfiguring the customer router, the **local-as** value on router P can be set to AS1. Therefore, router P adds AS1 to the as-path message for routes it advertises to the customer router.

The **no** form of this command used at the global level will remove any virtual ASN configured.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no local-as

Parameters

as-number

Specifies the virtual autonomous system number expressed as a decimal integer.

Values 1 to 65535

private

Specifies the local AS is hidden in paths learned from the peering.

local-preference

Syntax

local-preference *local-preference*

no local-preference

Context

config>router>bgp

config>router>bgp>group

```
config>router>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute.

This value is used if the BGP route arrives from a BGP peer without the **local-preference** integer set.

The specified value can be overridden by any value set via a route policy. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no local-preference

Parameters

local-preference

Specifies the local preference value to be used as the override value expressed as a decimal integer.

Values 0 to 4294967295

loop-detect

Syntax

loop-detect {**drop-peer** | **discard-route** | **ignore-loop** | **off**}

no loop-detect

Context

```
config>router>bgp
```

```
config>router>bgp>group
```

```
config>router>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures how the BGP peer session handles loop detection in the AS path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

Dynamic configuration changes of **loop-detect** are not recognized.

The **no** form of this command used at the global level reverts to default, which is **loop-detect ignore-loop**.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

loop-detect ignore-loop

Parameters

drop-peer

Specifies to send a notification to the remote peer and drops the session.

discard-route

Specifies to discard routes received from a peer with the same ASN as the router. This option prevents routes looped back to the router from being added to the routing information base and consuming memory. When this option is changed, the change will not be active for an established peer until the connection is re-established for the peer.

ignore-loop

Specifies to ignore routes with loops in the AS path but maintains peering.

off

Disables loop detection.

med-out

Syntax

med-out {*number* | **igp-cost**}

no med-out

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set.

The specified value can be overridden by any value set through a route policy.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command used at the global level reverts to default where the MED is not advertised.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no med-out

Parameters

number

Specifies the MED path attribute value expressed as a decimal integer.

Values 0 to 4294967295

igp-cost

Specifies that the MED is set to the IGP cost of the specific IP prefix.

min-as-origination

Syntax

min-as-origination *seconds*

no min-as-origination

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the minimum interval, in seconds, at which a path attribute, originated by the local router, can be advertised to a peer.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command used at the global level reverts to default.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

min-as-origination 15

Parameters

seconds

Specifies the minimum path attribute advertising interval, in seconds, expressed as a decimal integer.

Values 2 to 255

min-route-advertisement

Syntax

min-route-advertisement *seconds*

no min-route-advertisement

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the minimum interval, in seconds, at which a prefix can be advertised to a peer.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command used at the global level reverts to the default.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

min-route-advertisement 30

Parameters

seconds

Specifies the minimum route advertising interval, in seconds, expressed as a decimal integer.

Values 1 to 255

multihop

Syntax

multihop *ttl-value*

no multihop

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGp peer multiple hops away.

The **no** form of this command is used to convey to the BGP instance that the EBGp peers are directly connected.

The **no** form of this command used at the global level reverts to default.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

multihop 1 (EBGP peers are directly connected)

multihop 64 (IBGP)

Parameters

ttl-value

Specifies the TTL value expressed as a decimal integer.

Values 1 to 255

next-hop-resolution

Syntax

next-hop-resolution

Context

config>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure next-hop resolution.

label-route-transport-tunnel

Syntax

label-route-transport-tunnel

Context

config>router>bgp>next-hop-resolution

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure the binding of BGP labeled routes to tunnels.

family

Syntax

family *family*

Context

config>router>bgp>next-hop-res>label-route-transport-tunnel

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the binding of BGP labeled routes to tunnels for a specific family.

Default

family ipv4

Parameters

family

Specifies the family.

- Values** **ipv4** — Specifies tunnels for the IPv4 family.
 ipv6 — Specifies tunnels for the IPv6 family.
 vpn — Specifies tunnels for the VPN family.

resolution

Syntax

resolution {any | filter | disabled}

Context

config>router>bgp>next-hop-res>lbl-rt-tunn>family

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the resolution state of BGP labeled routes using tunnels to BGP peers.

Default

resolution filter

Parameters

any

Keyword that enables binding to any supported tunnel type in the BGP labeled route context following the TTM preference.

filter

Keyword that enables binding to the subset of tunnel types configured under the **resolution-filter** context.

disabled

Keyword that disables the resolution of BGP labeled routes using tunnels to BGP peers.

resolution-filter

Syntax

resolution-filter

Context

config>router>bgp>next-hop-res>lbl-rt-tunn>family

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure the subset of tunnel types that can be used in the resolution of BGP label routes.

ldp

Syntax

[no] ldp

Context

config>router>bgp>next-hop-res>lbl-rt-tunn>family>resolution-filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures LDP tunneling for next-hop resolution.

rsvp

Syntax

[no] rsvp

Context

config>router>bgp>next-hop-res>lbl-rt-tunn>family>resolution-filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures RSVP tunneling for next-hop resolution.

Default

no rsvp

sr-isis

Syntax

[no] sr-isis

Context

config>router>bgp>next-hop-res>lbl-rt-tunn>family>resolution-filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command selects the SR tunnel type programmed by an IS-IS instance in the TTM for next-hop resolution and specifies SR tunnels (shortest path) to destinations reachable by the IS-IS protocol. This command allows BGP to use the SR tunnel in the tunnel table submitted by the lowest preference IS-IS instance or, in the case of IS-IS instances with the same lowest preference, the IS-IS instance with the lowest ID number.

The **no** form of this command removes the SR tunnel type.

sr-ospf

Syntax

[no] sr-ospf

Context

config>router>bgp>next-hop-res>lbl-rt-tunn>family>resolution-filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command selects the SR tunnel type programmed by an OSPF instance in the TTM for next-hop resolution and specifies SR tunnels (shortest path) to destinations reachable by the OSPF protocol. This command allows BGP to use the SR tunnel in the tunnel table submitted by the lowest preference OSPF instance or, in the case of IS-IS instances with the same lowest preference, the OSPF instance with the lowest ID number.

The **no** form of this command removes the SR tunnel type.

outbound-route-filtering

Syntax

[no] **outbound-route-filtering**

Context

```
config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command opens the configuration tree for sending or accepting BGP filter lists from peers (outbound route filtering).

Default

no outbound-route-filtering

extended-community

Syntax

[no] **extended-community**

Context

```
config>router>bgp>orf
config>router>bgp>group>orf
config>router>bgp>group>neighbor>orf
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command opens the configuration tree for sending or accepting extended community-based BGP filters.

In order for the **no** version of the command to work, all subcommands (**send-orf**, **accept-orf**) must be removed first.

accept-orf

Syntax

[no] accept-orf

Context

config>router>bgp>orf>ext-comm

config>router>bgp>group>orf>ext-comm

config>router>bgp>group>neighbor>orf>ext-comm

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command instructs the router to negotiate the receive capability in the BGP outbound route filtering (ORF) negotiation with a peer, and to accept filters that the peer wishes to send.

Accepting ORFs is not enabled by default.

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

send-orf

Syntax

send-orf [*comm-id...*(up to 32 max)]

no send-orf [*comm-id*]

Context

config>router>bgp>orf>ext-comm

config>router>bgp>group>orf>ext-comm

config>router>bgp>group>neighbor>orf>ext-comm

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command instructs the router to negotiate the send capability in the BGP ORF negotiation with a peer.

This command also causes the router to send a community filter, prefix filter, or AS path filter configured as an inbound filter on the BGP session to its peer as an ORF Action ADD.

If the *comm-id* parameters are not exclusively route target communities, the router will extract appropriate route targets and use those. If, for some reason, the *comm-id* parameters specified contain no route targets, the router will not send an ORF.

The **no** form of this command causes the router to remove the send capability in the BGP ORF negotiation with a peer. The **no** form also causes the router to send an ORF remove action for a community filter, prefix filter, or AS path filter configured as an inbound filter on the BGP session to its peer.

Default
no send-orf

Parameters

comm-id

Specifies a community policy that consists exclusively of route target extended communities. If it is not specified, the ORF policy is automatically generated from configured route target lists, accepted client route target ORFs, and locally configured route targets.

neighbor

Syntax
[no] neighbor ip-address

Context
config>router>bgp>group

Platforms
Supported on all 7210 SAS platforms as described in this document

Description

This command creates a BGP peer/neighbor instance within the context of the BGP group.

This command can be issued repeatedly to create multiple peers and their associated configuration.

The **no** form of this command is used to remove the specified neighbor and the entire configuration associated with the neighbor. The neighbor must be administratively **shutdown** before attempting to delete it. If the neighbor is not shutdown, the command will not result in any action except a warning message on the console indicating that neighbor is still administratively up.

Parameters

ip-address

Specifies the IP address of the BGP peer router.

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

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

peer-as

Syntax
peer-as *as-number*

Context
config>router>bgp>group
config>router>bgp>group>neighbor

Platforms
Supported on all 7210 SAS platforms as described in this document

Description
This command configures the autonomous system number for the remote peer. The peer ASN must be configured for each configured peer.
For eBGP peers, the peer ASN configured must be different from the autonomous system number configured for this router under the global level since the peer will be in a different autonomous system than this router
For iBGP peers, the peer ASN must be the same as the ASN of this router configured under the global level.
This is required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.

Parameters
as-number
Specifies the autonomous system number expressed as a decimal integer.
Values 1 to 4294967295

path-mtu-discovery

Syntax
[no] path-mtu-discovery

Context

```
config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables path MTU discovery for the associated TCP connections. In doing so, the MTU for the associated TCP session will be initially set to the egress interface MTU. The DF bit will also be set so that if a router along the path of the TCP connection cannot handle a packet of a particular size without fragmenting, it will send back an ICMP message to set the path MTU for the specific session to a lower value that can be forwarded without fragmenting.

The **no** form of this command disables path MTU discovery.

Default

no path-mtu-discovery

preference

Syntax

[no] **preference** *preference*

Context

```
config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the route preference for routes learned from the configured peers.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The lower the preference the higher the chance of the route being the active route. The 7210 SAS assigns BGP routes the highest default preference compared to routes that are direct, static, or learned through MPLS or OSPF.

The **no** form of this command used at the global level reverts to default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

preference 170

Parameters

preference

Specifies the route preference expressed as a decimal integer.

Values 1 to 255

purge-timer

Syntax

[no] **purge-timer** *minutes*

Context

config>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum time before stale routes are purged.

Parameters

minutes

Specifies the duration of the purge timer, in minutes.

Values 1 to 60

rapid-update

Syntax

rapid-update [I2-vpn] [evpn]

no rapid-update

Context

config>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables BGP rapid update for specified address families.

If rapid update is enabled for a set of address families, and a route belonging to a family in that set is received by the router and chosen for propagation to specific BGP peers, the remaining time on the MRAI timer of these peers is ignored and the route is transmitted immediately, along with all other pending routes for these peers, including routes of address families not specified in the **rapid-update** command.

The **rapid-update** command overrides the peer-level time and applies the minimum setting of 0 seconds to routes belonging to specified address families; routes of other address families continue to be advertised according to the session-level MRAI setting.

The **no** form of this command disables rapid update for all address families.

Default

no rapid-update

Parameters

l2-vpn

Keyword to enable the BGP rapid update for the 12-byte Virtual Switch Instance identifier (VSI-ID) value, which consists of the 8-byte route distinguisher (RD) followed by a 4-byte value.

evpn

Keyword to enable the BGP rapid update for the EVPN address family by including EVPN routes from the set of routes that can trigger rapid update.

rapid-withdrawal

Syntax

[no] rapid-withdrawal

Context

config>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command disables the delay (Minimum Route Advertisement) on sending BGP withdrawals. Normal route withdrawals may be delayed up to the minimum route advertisement to allow for efficient packing of BGP updates.

The **no** form of this command removes the configuration and reverts withdrawal processing to the normal behavior.

Default

no rapid-withdrawal

prefix-limit

Syntax

prefix-limit *family limit* [**log-only**] [**threshold percentage**]

no prefix-limit

Context

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum number of routes BGP can learn from a peer.

When the number of routes reaches 90% of this limit, an SNMP trap is generated. When the limit is exceeded, the BGP peering is dropped and disabled.

This command only applies to BGP routes learned for different families supported by the BGP protocol on the 7210 SAS. Different IP FIB limits are supported for different IPv4 and IPv6 address prefix lengths. There are two limits to consider: one is the value configured as part of the **prefix-limit** command and the second is the maximum IP FIB limit supported on the node. These two limits impact the behavior of the **prefix-limit** command as follows. When a BGP route for the configured *family* is received, the following comparison is completed:

- BGP peering is brought down if all of the following conditions are true:
 - if the number of routes in the FIB plus the received route is greater than the value configured for **prefix-limit**
 - if the number of routes is less than the maximum IP FIB limit
 - if **log-only** is not configured
- BGP peering remains up if all of the following conditions are true:
 - if the number of routes in the FIB plus the received route is greater than the value configured for **prefix-limit**
 - if the number of routes is less than the maximum IP FIB limit
 - if **log-only** is configured

A log is generated to report the addition of the route if the **prefix-limit** value is exceeded. Excess routes are added to the IP FIB.

- If the number of routes in the FIB plus the received route is greater than the maximum IP FIB limit, regardless of whether **log-only** is configured, the BGP peering session is brought down.

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

Default

no prefix-limit

Parameters

log-only

Keyword to enable the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. BGP peering is not dropped.

percent

Specifies the threshold value (as a percentage) that triggers a warning message.

Values 1 to 100

limit

Specifies the number of routes, expressed as a decimal integer, that can be learned from a peer.

Values 1 to 4294967295

family

Specifies the address family applied for the prefix limit.

Values ipv4, vpn-ipv4, ipv6, vpn-ipv6, mvpn-ipv4

remove-private

Syntax

[no] **remove-private** {limited}

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables private AS numbers to be removed from the AS path before advertising them to BGP peers.

When the **remove-private** command is configured at the global level, it applies to all peers, regardless of group or neighbor configuration. When the parameter is set at the group level, it applies to all peers in the group, regardless of the neighbor configuration.

7210 SAS software recognizes the set of AS numbers that are defined by the Internet Assignment Numbers Authority (IANA) as private. These are AS numbers in the range 64512 through 65535, inclusive.

The **no** form of this command used at the global level reverts to default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no remove-private

Parameters

limited

Optional keyword to remove private ASNs up to the first public ASN encountered. It then stops removing private ASNs.

router-id

Syntax

router-id *ip-address*

no router-id

Context

config>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the router ID to be used with this BGP instance.

Changing the BGP router ID on an active BGP instance causes the BGP instance to restart with the new router ID. The router ID must be set to a valid host address.

By default, the system interface IP address is used.

Parameters

ip-address

Specifies the router ID, expressed in dotted-decimal notation. The allowed value is a valid routable IP address on the router, either an interface or system IP address. It is highly recommended that this address be the system IP address.

Values a.b.c.d

shutdown

Syntax

[no] **shutdown**

Context

config>router>bgp

config>router>bgp>group

```
config>router>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

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

Default administrative states for services and service entities are described in Special Cases.

The **no** form of this command places an entity in an administratively enabled state.

Special Cases

BGP Global

The BGP protocol is created in the **no shutdown** state.

BGP Group

BGP groups are created in the **no shutdown** state.

BGP Neighbor

BGP neighbors/peers are created in the **no shutdown** state.

type

Syntax

```
[no] type {internal | external}
```

Context

```
config>router>bgp>group
```

```
config>router>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command designates the BGP peer as type internal or external.

The type **internal** indicates that the peer is an iBGP peer while the type **external** indicates that the peer is an eBGP peer.

By default, 7210 SAS derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, the peer is considered **external**.

The **no** form of this command used at the group level reverts to the default value.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no type

Parameters

internal

Keyword to configure the peer as internal.

external

Keyword to configure the peer as external.

5.24.2.2 Other BGP-related commands

autonomous-system

Syntax

autonomous-system *autonomous-system*

no autonomous-system

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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

If the ASN is changed on a router with an active BGP instance, the new ASN is not used until the BGP instance is restarted either by administratively disabling or enabling (**shutdown/no shutdown**) the BGP instance or rebooting the system with the new configuration.

Parameters

as-number

Specifies the ASN, expressed as a decimal integer.

Values 1 to 4294967295

router-id

Syntax

router-id *ip-address*

no router-id

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the router ID for the router instance.

The router ID is used by both OSPF and BGP routing protocols in this instance of the routing table manager.

When 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. 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.

By 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.

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

Parameters

ip-address

Specifies the router ID, expressed in dotted-decimal notation. The allowed value is a valid routable IP address on the router, either an interface or system IP address. It is highly recommended that this address be the system IP address.

Values a.b.c.d

5.24.2.3 Show commands

router

Syntax

router [*router-instance*]

Context

show

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays router instance information.

Parameters

router-instance

Specifies either the router name or service ID.

Values router-instance: Base, management

Default Base, Management

bgp

Syntax

bgp

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context display BGP related information.

auth-keychain

Syntax

auth-keychain *keychain*

Context

show>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays BGP sessions using a specific authentication keychain.

Parameters

keychain
Specifies an existing keychain name, up to 32 characters.

Output

The following output is an example of authentication keychain information, and [Table 64: Output fields: BGP auth-keychain](#) describes the output fields.

Sample output — auth-keychain

```
*A:ALA-48# show router 2 bgp auth-keychain
=====
Sessions using key chains
=====
Peer address          Group          Keychain name
-----
10.20.1.3             1             eta_keychain1
30.1.0.2              1             eta_keychain1
=====

*A:ALA-48#
*A:ALA-48>config>router>bgp# show router bgp group "To_AS_10000"
=====
BGP Group : To_AS_10000
-----
Group              : To_AS_10000
-----
Group Type        : No Type          State           : Up
Peer AS           : 10000             Local AS        : 200
Local Address     : n/a              Loop Detect     : Ignore
Import Policy     : None Specified / Inherited
Hold Time         : 90                Keep Alive      : 30
NLRI              : Unicast           Preference      : 170
TTL Security      : Disabled          Min TTL Value   : n/a
Graceful Restart  : Enabled           Stale Routes Time: 360
Auth key chain    : testname

List of Peers
- 10.0.0.8 :
  To_Router B - EBGP Peer
Total Peers    : 1                Established     : 0
-----
Peer Groups : 1
=====

*A:ALA-48>config>router>bgp#

*A:Dut-b>config>router>

*A:ALA-48>config>router>bgp# show router bgp neighbor 10.0.0.8
=====
BGP Neighbor
-----
Peer   : 10.0.0.8
Group  : To_AS_10000
-----
Peer AS           : 10000          Peer Port        : 0
Peer Address      : 10.0.0.8
```

```

Local AS      : 200      Local Port      : 0
Local Address : 0.0.0.0
Peer Type     : External
State        : Active    Last State     : Idle
Last Event    : stop
Last Error    : Cease
Local Family  : IPv4
Remote Family : Unused
Hold Time     : 90        Keep Alive      : 30
Active Hold Time : 0      Active Keep Alive : 0
Preference    : 99        Num of Flaps    : 0
Recd. Paths   : 0
IPv4 Recd. Prefixes : 0    IPv4 Active Prefixes : 0
IPv4 Suppressed Pfxs : 0    VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs : 0     VPN-IPv4 Active Pfxs : 0
Mc IPv4 Recd. Pfxs. : 0     Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs : 0
Input Queue   : 0        Output Queue    : 0
i/p Messages  : 0        o/p Messages    : 0
i/p Octets    : 0        o/p Octets      : 0
i/p Updates   : 0        o/p Updates     : 0
TTL Security  : Disabled  Min TTL Value    : n/a
Graceful Restart : Enabled  Stale Routes Time : 360
Advertise Inactive : Disabled Peer Tracking     : Disabled
Advertise Label : None
Auth key chain  : testname
Local Capability : RouteRefresh MP-BGP
Remote Capability :
Import Policy   : None Specified / Inherited

```

Neighbors : 1

=====

```
*A:ALA-48>config>router>bgp#
```

```
*A:ALA-48>config>router>bgp# show router bgp auth-keychain testname
```

=====

```
Sessions using key chain: keychain
```

```

=====
Peer address      Group          Keychain name
-----
10.0.0.8          To_AS_10000    testname
=====

```

```
*A:ALA-48>config>router>bgp#
```

Table 64: Output fields: BGP auth-keychain

Label	Description
Peer address	Displays the IP address of the peer
Group	Displays the BGP group name
Keychain name	Displays the authentication keychain associated with the session, if applicable

damping

Syntax

```
damping [ip-prefix [/ip-prefix-length]] [damp-type] [detail] [ipv4]
damping [ip-prefix [/ip-prefix-length]] [damp-type] [detail] [ipv6]
damping [ip-prefix [/ip-prefix-length]] [damp-type] [detail] [vpn-ipv4]
damping [ip-prefix [/ip-prefix-length]] [damp-type] [detail] [vpn-ipv6]
```

Context

show>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays BGP routes that have been dampened due to route flapping. This command can be entered with or without a route parameter. If no parameters are included, all dampened routes are listed.

When the keyword **detail** is included, more detailed information is displayed.

If a *damp-type* is specified, only those types of dampened routes (**decayed**, **history**, or **suppressed**) are displayed. Routes that have a state of **decayed** have gained penalties for flapping but have not yet reached the suppression limit. Routes that have a state of **history** have had a route flap and have been withdrawn. Routes that have a state of **suppressed** have reached the suppression limit and are not considered in BGP path selection.

Parameters

ip-prefix/ip-prefix-length

Displays damping information for the specified IP address.

Values	
ipv4-prefix:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 to 32
ipv6-prefix:	x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6-prefix-length:	0 to 128

damp-type

Displays damping information for routes with the specified damp type.

Values decayed, history, suppressed.

detail

Displays detailed information.

ipv4

Displays dampened routes for the IPv4 address family.

ipv6

Displays dampened routes for the IPv6 address family.

vpn-ipv4

Displays dampened routes for the VPN-IPv4 address family.

vpn-ipv6

Displays dampened routes for the VPN-IPv6 address family.

Output

The following output is an example of BGP damping information, and [Table 65: Output fields: router BGP damping](#) describes the output fields.

Sample output

```
A:ALA-12# show router bgp damping
=====
BGP Router ID : 10.0.0.14      AS : 65206   Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes
=====
```

Flag	Network	From	Reuse	AS-Path
ud*i	10.149.7.0/24	10.0.28.1	00h00m00s	60203 65001 19855 3356 1239 22406
si	10.155.6.0/23	10.0.28.1	00h43m41s	60203 65001 19855 3356 2914 7459
si	10.155.8.0/22	10.0.28.1	00h38m31s	60203 65001 19855 3356 2914 7459
si	10.155.12.0/22	10.0.28.1	00h35m41s	60203 65001 19855 3356 2914 7459
si	10.155.22.0/23	10.0.28.1	00h35m41s	60203 65001 19855 3356 2914 7459
si	10.155.24.0/22	10.0.28.1	00h35m41s	60203 65001 19855 3356 2914 7459
si	10.155.28.0/22	10.0.28.1	00h34m31s	60203 65001 19855 3356 2914 7459
si	10.155.40.0/21	10.0.28.1	00h28m24s	60203 65001 19855 3356 7911 7459
si	10.155.48.0/20	10.0.28.1	00h28m24s	60203 65001 19855 3356 7911 7459
ud*i	10.8.140.0/24	10.0.28.1	00h00m00s	60203 65001 19855 3356 4637 17447
ud*i	10.8.141.0/24	10.0.28.1	00h00m00s	60203 65001 19855 3356

```

                                4637 17447
ud*i 10.9.0.0/18      10.0.28.1      00h00m00s 60203 65001 19855 3356
                                3561 9658 6163
. . .
ud*i 10.213.184.0/23 10.0.28.1      00h00m00s 60203 65001 19855 3356
                                6774 6774 9154
-----
A:ALA-12#

```

A:ALA-12# show router bgp damping detail

```

=====
BGP Router ID : 10.0.0.14      AS : 65206      Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * -
valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes
=====
-----
Network : 10.149.7.0/24
-----
Network      : 10.149.7.0/24      Peer      : 10.0.28.1
NextHop      : 10.0.28.1          Reuse time : 00h00m00s
Peer AS      : 60203              Peer Router-Id : 32.32.27.203
Local Pref   : none
Age          : 00h22m09s          Last update : 02d00h58m
FOM Present  : 738                FOM Last upd. : 2039
Number of Flaps : 2                Flags       : ud*i
Path         : 60203 65001 19855 3356 1239 22406
Applied Policy : default-damping-profile
-----
Network : 10.142.48.0/20
-----
Network      : 10.142.48.0/20      Peer      : 10.0.28.1
NextHop      : 10.0.28.1          Reuse time : 00h00m00s
Peer AS      : 60203              Peer Router-Id : 32.32.27.203
Local Pref   : none
Age          : 00h00m38s          Last update : 02d01h20m
FOM Present  : 2011               FOM Last upd. : 2023
Number of Flaps : 2                Flags       : ud*i
Path         : 60203 65001 19855 3356 3561 5551 1889
Applied Policy : default-damping-profile
-----
Network : 10.200.128.0/19
-----
Network      : 10.200.128.0/19      Peer      : 10.0.28.1
NextHop      : 10.0.28.1          Reuse time : 00h00m00s
Peer AS      : 60203              Peer Router-Id : 32.32.27.203
Local Pref   : none
Age          : 00h00m38s          Last update : 02d01h20m
FOM Present  : 2011               FOM Last upd. : 2023
Number of Flaps : 2                Flags       : ud*i
Path         : 60203 65001 19855 1299 702 1889
Applied Policy : default-damping-profile
-----
Network : 10.203.192.0/18
-----
Network      : 10.203.192.0/18      Peer      : 10.0.28.1
NextHop      : 10.0.28.1          Reuse time : 00h00m00s
Peer AS      : 60203              Peer Router-Id : 32.32.27.203

```



```
Local Pref      : none
Age             : 00h00m07s      Last update      : 02d01h20m
FOM Present     : 1018           FOM Last upd.    : 1024
Number of Flaps : 1              Flags              : ud*i
Path            : 60203 65001 19855 1299 702 1889
Applied Policy  : default-damping-profile
-----
A:ALA-12#
A:ALA-12# show router bgp damping 10.203.192.0/18 detail
=====
      BGP Router ID : 10.0.0.14      AS : 65206      Local AS : 65206
=====
      Legend -
      Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
      Origin codes  : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes 15.203.192.0/18
=====
Network : 10.203.192.0/18
-----
Network      : 10.203.192.0/18      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h00m00s
Peer AS      : 60203               Peer Router-Id : 32.32.27.203
Local Pref   : none
Age          : 00h00m42s           Last update : 02d01h20m
FOM Present  : 2003                FOM Last upd. : 2025
Number of Flaps : 2                 Flags        : ud*i
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Paths : 1
=====
A:ALA-12#

A:ALA-12# show router bgp damping suppressed detail
=====
      BGP Router ID : 10.0.0.14      AS : 65206      Local AS : 65206
=====
      Legend -
      Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
      Origin codes  : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes (Suppressed)
=====
Network : 10.142.48.0/20
-----
Network      : 10.142.48.0/20      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h29m22s
Peer AS      : 60203               Peer Router-Id : 32.32.27.203
Local Pref   : none
Age          : 00h01m28s           Last update : 02d01h20m
FOM Present  : 2936                FOM Last upd. : 3001
Number of Flaps : 3                 Flags        : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Network : 10.200.128.0/19
-----
Network      : 10.200.128.0/19      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h29m22s
Peer AS      : 60203               Peer Router-Id : 32.32.27.203
```

```

Local Pref      : none
Age             : 00h01m28s      Last update      : 02d01h20m
FOM Present     : 2936           FOM Last upd.    : 3001
Number of Flaps : 3              Flags              : si
Path            : 60203 65001 19855 3356 702 1889
Applied Policy  : default-damping-profile
-----
Network : 10.203.240.0/20
-----
Network      : 10.203.240.0/20    Peer            : 10.0.28.1
NextHop      : 10.0.28.1          Reuse time      : 00h29m22s
Peer AS      : 60203              Peer Router-Id  : 32.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update     : 02d01h20m
FOM Present  : 2936               FOM Last upd.   : 3001
Number of Flaps : 3               Flags            : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Network : 10.206.0.0/17
-----
Network      : 10.206.0.0/17      Peer            : 10.0.28.1
NextHop      : 10.0.28.1          Reuse time      : 00h29m22s
Peer AS      : 60203              Peer Router-Id  : 32.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update     : 02d01h20m
FOM Present  : 2936               FOM Last upd.   : 3001
Number of Flaps : 3               Flags            : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
A:ALA-12#

```

Table 65: Output fields: router BGP damping

Label	Description
BGP Router ID	Displays the local BGP router ID
The local BGP router ID.	Displays the configured ASN
Local AS	Displays the configured or inherited local AS for the specified peer group. If not configured, it is the same value as the AS.
Network	Displays the route IP prefix and mask length for the route
Flag(s)	Legend: Status codes: u- used, s-suppressed, h-history, d-decayed, *-valid. If a * is not present, then the status is invalid. Origin codes: i-IGP, e-EGP, ?-incomplete, >-best
From	Displays the originator ID path attribute value
Reuse time	Displays the time when a suppressed route can be used again
From	Displays the originator ID path attribute value

Label	Description
Reuse time	Displays the time when a suppressed route can be used again
AS Path	Displays the BGP AS path for the route
Peer	Displays the router ID of the advertising router
NextHop	Displays the BGP next hop for the route
Peer AS	Displays the ASN of the advertising router
Peer Router-Id	Displays the router ID of the advertising router
Local Pref	Displays the BGP local preference path attribute for the route
Age	Displays the length of time in the hour/minute/second (HH:MM:SS) format.
Last update	Displays the time when BGP was updated last in the day/hour/minute (DD:HH:MM) format
FOM Present	Displays the current Figure of Merit (FOM) value
Number of Flaps	Displays the number of route flaps in the neighbor connection
Path	Displays the BGP AS path for the route
Applied Policy	Displays the applied route policy name

group

Syntax

group [*name*] [*detail*]

Context

show>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays group information for a BGP peer group. This command can be entered with or without parameters.

When this command is entered without a group name, information about all peer groups displays.

When the command is issued with a specific group name, information pertaining to only that specific peer group displays.

The 'State' field displays the BGP group operational state. Valid states are the following:

- **Up**
The BGP global process is configured and running.
- **Down**
Te BGP global process is administratively shutdown and not running.
- **Disabled**
The BGP global process is operationally disabled. The process must be restarted by the operator.

Parameters

- name**
Displays information for the specified BGP group, up to 32 characters.
- detail**
Displays detailed information.

Output

The following outputs are examples of BGP group information, and [Table 66: Output fields: router BGP group](#) describes the output fields:

- [Sample output](#)
- [Sample detailed output](#)

Sample output

```
A:ALA-12# show router bgp group
=====
BGP Groups
-----
Group           : To_AS_40000
-----
Description      : Not Available
Group Type       : No Type           State           : Up
Peer AS          : 40000             Local AS         : 65206
Local Address    : n/a               Loop Detect      : Ignore
Export Policy    : direct2bgp
Hold Time        : 90                Keep Alive       : 30
NLRI             : Unicast           Preference       : 170

List of Peers
- 10.0.0.1       : To_Jukebox
- 10.0.0.12      : Not Available
- 10.0.0.13      : Not Available
- 10.0.0.14      : To_SR1
- 10.0.0.15      : To_H-215

Total Peers      : 5                  Established      : 2
=====
A:ALA-12#
```

Sample detailed output

```
A:ALA-12# show router bgp group detail
=====
BGP Groups (detail)
-----
Group           : To_AS_40000
```

```

-----
Description      : Not Available
Group Type       : No Type
Peer AS          : 40000
Local Address    : n/a
Connect Retry    : 20
Local Pref       : 100
Multihop         : 0 (Default)
Min Route Advt.  : 30
Prefix Limit     : No Limit
Next Hop Self    : Disabled
Remove Private   : Disabled
Export Policy    : direct2bgp
Hold Time        : 90
NLRI             : Unicast
State            : Up
Local AS         : 65206
Loop Detect      : Ignore
Authentication   : None
MED Out          : 0
Min AS Originate : 15
Passive          : Disabled
Aggregator ID 0  : Disabled
Damping          : Disabled
Keep Alive       : 30
Preference       : 170

List of Peers
- 10.0.0.1       : To_Jukebox
- 10.0.0.12      : Not Available
- 10.0.0.13      : Not Available
- 10.0.0.14      : To_SR1
- 10.0.0.15      : To_H-215

Total Peers      : 5
Established      : 2
=====
A:ALA-12#

A:SetupCLI>show>router>bgp# group
=====
BGP Group
-----
Group            : bgp_group_1 34567890123456789012
-----
Description      : Testing the length of the group value for the DESCRIPTION
                  : parameter of BGP
Group Type       : No Type
Peer AS          : n/a
Local Address    : n/a
Import Policy    : test i1
                  : test i2
                  : test i3
                  : test i4
                  : test i5 890123456789012345678901
Export Policy    : test e1
                  : test e2
                  : test e3
                  : test e4
                  : test e5 890123456789012345678901
Hold Time        : 120
NLRI             : Unicast
TTL Security     : Disabled
Graceful Restart : Disabled
Auth key chain   : n/a
State            : Up
Local AS         : 100
Loop Detect      : Ignore
Keep Alive       : 30
Preference       : 101
Min TTL Value    : n/a
Stale Routes Time: n/a
Bfd Enabled      : Yes

List of Peers
- 10.3.3.3 :
  Testing the length of the neighbor value for the DESCRIPTION parameter of
  BGP
Total Peers      : 1
Established      : 0
-----
Peer Groups : 1
=====
A:SetupCLI>show>router>bgp#

```

Table 66: Output fields: router BGP group

Label	Description
Group	Displays the BGP group name
Group Type	No Type — Peer type not configured External — Peer type configured as external BGP peers Internal — Peer type configured as internal BGP peers
State	Disabled — The BGP peer group has been operationally disabled Down — The BGP peer group is operationally inactive Up — The BGP peer group is operationally active
Peer AS	Displays the configured or inherited peer AS for the specified peer group
Local AS	Displays the configured or inherited local AS for the specified peer group
Local Address	Displays the configured or inherited local address for originating peering for the specified peer group
Loop Detect	Displays the configured or inherited loop detect setting for the specified peer group
Connect Retry	Displays the configured or inherited connect retry timer value
Authentication	None — No authentication is configured MD5 — MD5 authentication is configured
Bfd	Yes — BFD is enabled No — BFD is disabled
Local Pref	Displays the configured or inherited local preference value
MED Out	Displays the configured or inherited MED value assigned to advertised routes without a MED attribute
Min Route Advt.	Displays the minimum amount of time that must pass between route updates for the same IP prefix
Min AS Originate	Displays the minimum amount of time that must pass between updates for a route originated by the local router
Multihop	Displays the maximum number of router hops a BGP connection can traverse
Prefix Limit	No Limit — No route limit assigned to the BGP peer group

Label	Description
	1 to 4294967295 — The maximum number of routes BGP can learn from a peer
Passive	Disabled — BGP attempts to establish a BGP connection with neighbor in the specified peer group Enabled — BGP will not actively attempt to establish a BGP connection with neighbor in the specified peer group
Next Hop Self	Disabled — BGP is not configured to send only its own IP address as the BGP next hop in route updates to neighbors in the peer group Enabled — BGP sends only its own IP address as the BGP next hop in route updates to neighbors in the specified peer group
Aggregator ID 0	Disabled — BGP is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group Enabled — BGP is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group
Remove Private	Disabled — BGP will not remove all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group Enabled — BGP removes all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group
Damping	Disabled — The peer group is configured not to dampen route flaps Enabled — The peer group is configured to dampen route flaps
Export Policy	Displays the configured export policies for the peer group
Import Policy	Displays the configured import policies for the peer group
Hold Time	Displays the configured hold time setting
Keep Alive	Displays the configured keepalive setting
Client Reflect	Disabled — The BGP route reflector will not reflect routes to this neighbor Enabled — The BGP route reflector is configured to reflect routes to this neighbor
NLRI	Displays the type of NLRI information that the specified peer group can accept Unicast — IPv4 unicast routing information can be carried

Label	Description
Preference	Displays the configured route preference value for the peer group
List of Peers	Displays a list of BGP peers configured under the peer group
Total Peers	Displays the total number of peers configured under the peer group
Established	Displays the total number of peers that are in an established state

neighbor

Syntax

neighbor [*ip-address* [*detail*]]
neighbor [*as-address* [*detail*]]
neighbor *ip-address* [**family**] [*filter1*] [**brief**]
neighbor *ip-number* [**family**] *filter2*
neighbor *as-number* [**family**] *filter2*
neighbor *ip-address* **orf** [*filter3*]
neighbor *ip-address* **graceful-restart**

Context

show>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays BGP neighbor information. This command can be entered with or without any parameters.

When this command is issued without any parameters, information about all BGP peers displays.

When the command is issued with a specific IP address or ASN, information regarding only that specific peer or peers with the same AS displays.

When either **received-routes** or **advertised-routes** is specified, the routes received from or sent to the specified peer are listed (see [Sample output for BGP neighbor received routes](#)). This information is not available with SNMP.

When either **history** or **suppressed** is specified, the routes learned from those peers that either have a history or are suppressed are listed.

The State field displays the BGP peer protocol state. In addition to the standard protocol states, this field can also display the Disabled operational state, which indicates that the peer is operationally disabled and must be restarted by the operator.

Parameters

ip-address

Displays information for 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[- <i>interface</i>] x:x:x:x:x:x:d.d.d.d[- <i>interface</i>] x: [0 to FFFF]H d: [0 to 255]D <i>interface</i> : 32 characters maximum, mandatory for link local addresses.

as-number

Displays information for the specified ASN.

Values	1 to 65535
--------	------------

family

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

Values	ipv4 — Displays only those BGP peers that have the IPv4 family enabled. vpn-ipv4 — Displays only those BGP peers that have the VPN-IPv4 family enabled. ipv6 — Displays only those BGP peers that have the IPv6 family enabled. vpn-ipv6 — Displays only those BGP peers that have the VPN-IPv4 family enabled.
--------	--

filter1

Displays information for the specified IP address

Values	received-routes — Displays the number of routes received from this peer. advertised-routes — Displays the number of routes advertised by this peer. history —Displays statistics for dampened routes. suppressed — Displays the number of paths from this peer that have been suppressed by damping.
--------	---

detail — Displays detailed information pertaining to *filter1*.

filter2

Displays information for the specified ASN.

- Values**
- history** — Display statistics for dampened routes.
 - suppressed** — Display the number of paths from this peer that have been suppressed by damping.
 - detail** — Displays detailed information pertaining to *filter2*.

brief

Displays information in a brief format. This parameter is only supported with received-routes and advertised-routes.

orf

Displays outbound route filtering for the BGP instance. ORF (Outbound Route Filtering) is used to inform a neighbor of targets (using target-list) that it is willing to receive. This mechanism helps lessen the update exchanges between neighbors and saves CPU cycles to process routes that could have been received from the neighbor only to be dropped/ignored.

filter3

Displays path information for the specified IP address.

- Values**
- send** — Displays the number of paths sent to this peer.
 - receive** — Displays the number of paths received from this peer.

graceful-restart

Displays neighbors configured for graceful restart.

Output

The following outputs are examples of BGP neighbor information. The associated tables describe the output fields.

- [Sample output](#), [Sample detailed output](#), [Table 67: Output fields: router BGP neighbor](#)
- [Sample output for BGP neighbor received routes](#), [Table 68: Output fields: router BGP neighbor received-routes](#)
- [Sample output — add-path](#), [Table 69: Output fields: show neighbor add-path](#)

Sample output

```
*A:7210-SAS>show>router>bgp# neighbor

=====
BGP Neighbor
=====
-----
Peer   : 1.1.1.1
Group  : sample
-----
Peer AS      : 12345      Peer Port      : 0
Peer Address : 1.1.1.1
Local AS     : 143        Local Port     : 0
Local Address: 0.0.0.0
```

```

Peer Type      : External
State          : Active      Last State      : Connect
Last Event     : openFail
Last Error     : Cease
Local Family   : IPv4 VPN-IPv4
Remote Family  : Unused
Hold Time      : 10000 (strict) Keep Alive       : 21845
Active Hold Time : 0         Active Keep Alive : 0
Cluster Id     : None
Preference     : 10          Num of Flaps    : 0
Recd. Paths    : 0
IPv4 Recd. Prefixes : 0      IPv4 Active Prefixes : 0
IPv4 Suppressed Pfxs : 0      VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs : 0      VPN-IPv4 Active Pfxs : 0
Mc IPv4 Recd. Pfxs. : 0      Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs : 0      IPv6 Suppressed Pfxs : 0
IPv6 Recd. Prefixes : 0      IPv6 Active Prefixes : 0
VPN-IPv6 Recd. Pfxs : 0      VPN-IPv6 Active Pfxs : 0
VPN-IPv6 Suppr. Pfxs : 0
MVPN-IPv4 Suppr. Pfxs: 0      MVPN-IPv4 Recd. Pfxs : 0
MVPN-IPv4 Active Pfxs: 0
Input Queue    : 0           Output Queue     : 0
i/p Messages   : 0           o/p Messages    : 1
i/p Octets     : 0           o/p Octets      : 0
i/p Updates    : 0           o/p Updates     : 0
TTL Security   : Disabled    Min TTL Value    : n/a
Graceful Restart : Enabled    Stale Routes Time : 3600
Advertise Inactive : Enabled  Peer Tracking    : Enabled
Advertise Label : None
Auth key chain  : keychain-one
Bfd Enabled     : Disabled    L2 VPN Cisco Interop : Disabled
Local Capability : RtRefresh MPBGP ORFSendExComm ORFRecvExComm
Remote Capability :
Import Policy   : abcd
Export Policy    : abcd

```

```

-----
Peer : 1.1.3.4
Group : test
-----

```

```

Peer AS        : 0           Peer Port       : 0
Peer Address    : 1.1.3.4
Local AS        : 12345      Local Port      : 0
Local Address   : 0.0.0.0
Peer Type       : External
State          : Idle       Last State      : Idle
Last Event      : none
Last Error      : Unrecognized Error
Local Family    : VPN-IPv4
Remote Family   : Unused
Hold Time       : 0 (strict) Keep Alive       : 0
Active Hold Time : 0         Active Keep Alive : 0
Preference      : 10          Num of Flaps    : 0
Recd. Paths     : 0
IPv4 Recd. Prefixes : 0      IPv4 Active Prefixes : 0
IPv4 Suppressed Pfxs : 0      VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs : 0      VPN-IPv4 Active Pfxs : 0
Mc IPv4 Recd. Pfxs. : 0      Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs : 0      IPv6 Suppressed Pfxs : 0
IPv6 Recd. Prefixes : 0      IPv6 Active Prefixes : 0
VPN-IPv6 Recd. Pfxs : 0      VPN-IPv6 Active Pfxs : 0
VPN-IPv6 Suppr. Pfxs : 0
MVPN-IPv4 Suppr. Pfxs: 0      MVPN-IPv4 Recd. Pfxs : 0
MVPN-IPv4 Active Pfxs: 0

```

```

Input Queue      : 0          Output Queue      : 0
i/p Messages    : 0          o/p Messages    : 0
i/p Octets      : 0          o/p Octets      : 0
i/p Updates     : 0          o/p Updates     : 0
TTL Security    : Disabled   Min TTL Value    : n/a
Graceful Restart : Enabled    Stale Routes Time : 100
Advertise Inactive : Enabled  Peer Tracking     : Enabled
Advertise Label  : None
Auth key chain   : n/a
Bfd Enabled      : Enabled    L2 VPN Cisco Interop : Disabled
Local Capability : RtRefresh MPBGP
Remote Capability :
Import Policy    : abcd
Export Policy    : abcd

```

*A:7210-SAS>

*A:SetupCLI>show>router>bgp# neighbor

=====

BGP Neighbor

=====

Peer : 3.3.3.3

Group : bgp_group_1 34567890123456789012

Peer AS	: 20	Peer Port	: 0
Peer Address	: 3.3.3.3		
Local AS	: 100	Local Port	: 0
Local Address	: 0.0.0.0		
Peer Type	: Internal		
State	: Active	Last State	: Idle
Last Event	: stop		
Last Error	: Cease		
Local Family	: IPv4		
Remote Family	: Unused		
Hold Time	: 10	Keep Alive	: 30
Active Hold Time	: 0	Active Keep Alive	: 0
Preference	: 101	Num of Flaps	: 0
Recd. Paths	: 0		
IPv4 Recd. Prefixes	: 0	IPv4 Active Prefixes	: 0
IPv4 Suppressed Pfxs	: 0	VPN-IPv4 Suppr. Pfxs	: 0
VPN-IPv4 Recd. Pfxs	: 0	VPN-IPv4 Active Pfxs	: 0
Mc IPv4 Recd. Pfxs	: 0	Mc IPv4 Active Pfxs	: 0
Mc IPv4 Suppr. Pfxs	: 0		
Input Queue	: 0	Output Queue	: 0
i/p Messages	: 0	o/p Messages	: 0
i/p Octets	: 0	o/p Octets	: 0
i/p Updates	: 0	o/p Updates	: 0
TTL Security	: Disabled	Min TTL Value	: n/a
Graceful Restart	: Enabled	Stale Routes Time	: 360
Advertise Inactive	: Disabled	Peer Tracking	: Enabled
Advertise Label	: None	Bfd Enabled	: Yes
Auth key chain	: n/a		
Local Capability	: RouteRefresh MP-BGP		
Remote Capability	:		
Import Policy	: test i1		
	: test i2		
	: test i3		
	: test i4		
	: test i5 890123456789012345678901		
Export Policy	: test e1		
	: test e2		

```

: test e3
: test e4
: test e5 890123456789012345678901
-----
Neighbors : 1
=====

*A:Dut-B>config>service# show router bgp neighbor
=====
BGP Neighbor
=====
-----
Peer          : 10.20.1.3
Description    : (Not Specified)
Group          : PEER_TO_C
-----
Peer AS       : 300
Peer Address   : 10.20.1.3
Local AS      : 300
Local Address  : 10.20.1.5
Peer Type     : Internal
State         : Established
Last Event    : recvKeepAlive
Last Error    : Cease (Other Configuration Change)
Local Family  : IPv4
Remote Family : IPv4
Hold Time     : 90
Min Hold Time : 0
Active Hold Time : 90
Cluster Id    : None
Preference    : 170
Recd. Paths   : 5
IPv4 Recd. Prefixes : 10
IPv4 Suppressed Pfxs : 0
VPN-IPv4 Recd. Pfxs : 0
Mc IPv4 Recd. Pfxs. : 0
Mc IPv4 Suppr. Pfxs : 0
IPv6 Recd. Prefixes : 0
VPN-IPv6 Recd. Pfxs : 0
VPN-IPv6 Suppr. Pfxs : 0
Mc IPv6 Recd. Pfxs. : 0
Mc IPv6 Suppr. Pfxs : 0
L2-VPN Recd. Pfxs : 0
MVPN-IPv4 Suppr. Pfxs : 0
MVPN-IPv4 Active Pfxs : 0
MDT-SAFI Recd. Pfxs : 0
Flow-IPv4 Suppr. Pfxs : 0
Flow-IPv4 Active Pfxs : 0
Rte-Tgt Recd. Pfxs : 0
Backup IPv4 Pfxs : 0
Mc Vpn Ipv4 Recd. Pf* : 0
Mc Vpn Ipv4 Suppr. P* : 0
Backup Vpn IPv4 Pfxs : 0
Input Queue : 0
i/p Messages : 30
i/p Octets : 1321
i/p Updates : 8
Flow-IPv6 Suppr. Pfxs : 0
Flow-IPv6 Active Pfxs : 0
Evpn Suppr. Pfxs : 0
Evpn Active Pfxs : 0
MS-PW Suppr. Pfxs : 0

Peer Port      : 179
Local Port     : 49635
Last State     : Active
Keep Alive     : 30
Active Keep Alive : 30
Num of Update Flaps : 20
IPv4 Active Prefixes : 10
VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Active Pfxs : 0
Mc IPv4 Active Pfxs. : 0
IPv6 Suppressed Pfxs : 0
IPv6 Active Prefixes : 0
VPN-IPv6 Active Pfxs : 0
Mc IPv6 Active Pfxs. : 0
L2-VPN Suppr. Pfxs : 0
L2-VPN Active Pfxs : 0
MVPN-IPv4 Recd. Pfxs : 0
MDT-SAFI Suppr. Pfxs : 0
MDT-SAFI Active Pfxs : 0
Flow-IPv4 Recd. Pfxs : 0
Rte-Tgt Suppr. Pfxs : 0
Rte-Tgt Active Pfxs : 0
Backup IPv6 Pfxs : 0
Mc Vpn Ipv4 Active P* : 0
Backup Vpn IPv6 Pfxs : 0
Output Queue : 0
o/p Messages : 26
o/p Octets : 470
o/p Updates : 0
Flow-IPv6 Recd. Pfxs : 0
Evpn Recd. Pfxs : 0
MS-PW Recd. Pfxs : 0

```

```

MS-PW Active Pfxs      : 0
TTL Security           : Disabled      Min TTL Value       : n/a
Graceful Restart       : Disabled      Stale Routes Time    : n/a
Restart Time           : n/a
Advertise Inactive     : Disabled      Peer Tracking        : Disabled
Advertise Label        : ipv4
Auth key chain         : n/a
Disable Cap Nego       : Disabled      Bfd Enabled          : Disabled
Flowspec Validate      : Disabled      Default Route Tgt    : Disabled
Aigp Metric            : Enabled        Split Horizon         : Disabled
Damp Peer Oscillatio*  : Disabled      Update Errors         : 0
GR Notification        : Disabled      Fault Tolerance       : Disabled
Rem Idle Hold Time     : 00h00m00s
Next-Hop Unchanged    : None
L2 VPN Cisco Interop   : Disabled
Local Capability       : RtRefresh MPBGP 4byte ASN
Remote Capability      : RtRefresh MPBGP 4byte ASN
Local AddPath Capabi*  : Disabled
Remote AddPath Capab*  : Send - None
                       : Receive - None
Import Policy          : None Specified / Inherited
Export Policy          : None Specified / Inherited
Origin Validation      : N/A
EBGP Link Bandwidth    : n/a
IPv4 Rej. Pfxs         : 0
VPN-IPv4 Rej. Pfxs     : 0
Mc IPv4 Rej. Pfxs      : 0
MVPN-IPv4 Rej. Pfxs    : 0
Flow-IPv4 Rej. Pfxs    : 0
L2-VPN Rej. Pfxs       : 0
Rte-Tgt Rej. Pfxs      : 0
Mc Vpn Ipv4 Rej. Pfxs  : 0
IPv6 Rej. Pfxs         : 0
VPN-IPv6 Rej. Pfxs     : 0
Mc IPv6 Rej. Pfxs      : 0
MVPN-IPv6 Rej. Pfxs    : 0
Flow-IPv6 Rej. Pfxs    : 0
MDT-SAFI Rej. Pfxs     : 0
MS-PW Rej. Pfxs        : 0
Evpn Rej. Pfxs         : 0

```

Neighbors : 1

=====

```

* indicates that the corresponding row element may have been truncated.
*A:Dut-B>config>service#

```

A:ALA-48# show router 2 bgp neighbor 10.20.1.3

=====

BGP Neighbor

=====

Peer : 10.20.1.3
Group : 1

```

-----
Peer AS           : 100      Peer Port         : 49725
Peer Address      : 10.20.1.3
Local AS          : 100      Local Port        : 179
Local Address     : 10.20.1.2
Peer Type         : Internal
State             : Established
Last Event        : recvKeepAlive
Last Error        : Cease
Local Family      : IPv4
Remote Family     : IPv4
Hold Time         : 3        Keep Alive         : 1
Active Hold Time  : 3        Active Keep Alive  : 1
Preference        : 170     Num of Flaps       : 0
Recd. Paths       : 1
IPv4 Recd. Prefixes : 11    IPv4 Active Prefixes : 10
IPv4 Suppressed Pfxs : 0    VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs : 0    VPN-IPv4 Active Pfxs : 0

```

```

Mc IPv4 Recd. Pfxs. : 0          Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs : 0
Input Queue         : 0          Output Queue         : 0
i/p Messages        : 471        o/p Messages         : 473
i/p Octets           : 3241       o/p Octets           : 3241
i/p Updates          : 4          o/p Updates          : 4
TTL Security        : Disabled    Min TTL Value         : n/a
Advertise Inactive   : Disabled    Peer Tracking         : Disabled
Advertise Label      : None
Auth key chain       : eta_keychain1
Local Capability     : RouteRefresh MP-BGP
Remote Capability    : RouteRefresh MP-BGP
Import Policy        : None Specified / Inherited
Export Policy        : static2bgp
-----
Neighbors : 1
=====
A:ALA-48#

A:ALA-12# show router bgp neighbor 10.0.0.11 orf
=====
BGP Neighbor 10.0.0.11 ORF
=====
Send List (Automatic)
-----
target:65535:10
target:65535:20
=====
A:ALA-12

A:ALA-22 show router bgp neighbor 10.0.0.1 orf
=====
BGP Neighbor 10.0.0.1 ORF
=====
Receive List
-----
target:65535:10
target:65535:20
=====
A:ALA-22

```

Sample detailed output

```

A:ALA-12# show router bgp neighbor detail
=====
BGP Neighbor (detail)
-----
Peer : 10.0.0.15          Group : To_AS_40000
-----
Peer AS      : 65205          Peer Port     : 0
Peer Address : 10.0.0.15      Local Port    : 0
Local AS     : 65206          Local Address : 10.0.0.16
Peer Type    : External
State        : Active         Last State     : Connect
Last Event   : openFail
Last Error   : Hold Timer Expire
Connect Retry : 20             Local Pref.    : 100
Min Route Advt. : 30          Min AS Orig.   : 15
Damping      : Disabled       Loop Detect     : Ignore

```

```

MED Out      : No MED Out      Authentication : None
Next Hop Self : Disabled       AggregatorID Zero: Disabled
Remove Private : Disabled      Passive         : Disabled
Prefix Limit  : No Limit
Hold Time     : 90              Keep Alive      : 30
Active Hold Time : 0           Active Keep Alive: 0
Preference    : 170            Num of Flaps    : 0
Recd. Prefixes : 0             Active Prefixes : 0
Recd. Paths    : 0             Suppressed Paths : 0
Input Queue    : 0             Output Queue    : 0
i/p Messages   : 0             o/p Messages   : 0
i/p Octets     : 0             o/p Octets     : 0
i/p Updates    : 0             o/p Updates    : 0
Export Policy  : direct2bgp
=====

```

A:ALA-12#

Table 67: Output fields: router BGP neighbor

Label	Description
Peer	Displays the IP address of the configured BGP peer
Group	Displays the BGP peer group to which this peer is assigned
Peer AS	Displays the configured or inherited peer AS for the peer group
Peer Address	Displays the configured address for the BGP peer
Peer Port	Displays the TCP port number used on the far-end system
Local AS	Displays the configured or inherited local AS for the peer group
Local Address	Displays the configured or inherited local address for originating peering for the peer group
Local Port	Displays the TCP port number used on the local system
Peer Type	External — Peer type configured as external BGP peers Internal — Peer type configured as internal BGP peers
Bfd	Yes — BFD is enabled No — BFD is disabled
State	Idle — The BGP peer is not accepting connections Active — BGP is listening for and accepting TCP connections from this peer Connect — BGP is attempting to establish a TCP connections from this peer Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION

Label	Description
	Established — BGP has successfully established a peering and is exchanging routing information
Last State	<p>Idle — The BGP peer is not accepting connections</p> <p>Active — BGP is listening for and accepting TCP connections from this peer</p> <p>Connect — BGP is attempting to establish a TCP connections from this peer</p> <p>Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer</p> <p>Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION</p>
Last Event	<p>start — BGP has initialized the BGP neighbor</p> <p>stop — BGP has disabled the BGP neighbor</p> <p>open — BGP transport connection opened</p> <p>close — BGP transport connection closed</p> <p>openFail — BGP transport connection failed to open</p> <p>error — BGP transport connection error</p> <p>connectRetry — Connect retry timer expired</p> <p>holdTime — Hold time timer expired</p> <p>keepAlive — Keepalive timer expired</p> <p>recvOpen — Receive an OPEN message</p> <p>revKeepalive — Receive a KEEPALIVE message</p> <p>recvUpdate — Receive an UPDATE message</p> <p>recvNotify — Receive a NOTIFICATION message</p> <p>None — No events have occurred</p>
Last Error	Displays the last BGP error and subcode to occur on the BGP neighbor
Connect Retry	Displays the configured or inherited connect retry timer value
Local Pref.	Displays the configured or inherited local preference value
Min Route Advt.	Displays the minimum amount of time that must pass between route updates for the same IP prefix
Min AS Originate	Displays the minimum amount of time that must pass between updates for a route originated by the local router
Multihop	Displays the maximum number of router hops a BGP connection can traverse

Label	Description
Damping	Disabled — BGP neighbor is configured not to dampen route flaps Enabled — BGP neighbor is configured to dampen route flaps
Loop Detect	Ignore — The BGP neighbor is configured to ignore routes with an AS loop Drop — The BGP neighbor is configured to drop the BGP peering if an AS loop is detected Off — AS loop detection is disabled for the neighbor
MED Out	Displays the configured or inherited MED value assigned to advertised routes without a MED attribute
Authentication	None — No authentication is configured MD5 — MD5 authentication is configured
Next Hop Self	Disabled — BGP is not configured to send only its own IP address as the BGP nexthop in route updates to the specified neighbor Enabled — BGP will send only its own IP address as the BGP nexthop in route updates to the neighbor
AggregatorID Zero	Disabled — The BGP Neighbor is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates Enabled — The BGP Neighbor is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates
Remove Private	Disabled — BGP will not remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor Enabled — BGP will remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor
Passive	Disabled — BGP will actively attempt to establish a BGP connection with the specified neighbor Enabled — BGP will not actively attempt to establish a BGP connection with the specified neighbor
Prefix Limit	No Limit — No route limit assigned to the BGP peer group 1 — 4294967295 — The maximum number of routes BGP can learn from a peer
Hold Time	Displays the configured hold time setting
Keep Alive	Displays the configured keepalive setting
Active Hold Time	Displays the negotiated hold time if the BGP neighbor is in an established state

Label	Description
Active Keep Alive	Displays the negotiated keepalive time if the BGP neighbor is in an established state
Client Reflect	Disabled — The BGP route reflector is configured not to reflect routes to this neighbor Enabled — The BGP route reflector is configured to reflect routes to this neighbor
Preference	Displays the configured route preference value for the peer group
Num of Flaps	Displays the number of route flaps in the neighbor connection
Recd. Prefixes	Displays the number of routes received from the BGP neighbor
Active Prefixes	Displays the number of routes received from the BGP neighbor and active in the forwarding table
Recd. Paths	Displays the number of unique sets of path attributes received from the BGP neighbor
Suppressed Paths	Displays the number of unique sets of path attributes received from the BGP neighbor and suppressed due to route damping
Input Queue	Displays the number of BGP messages to be processed
Output Queue	Displays the number of BGP messages to be transmitted
i/p Messages	Displays the total number of packets received from the BGP neighbor
o/p Messages	Displays the total number of packets sent to the BGP neighbor
i/p Octets	Displays the total number of octets received from the BGP neighbor
o/p Octets	Displays the total number of octets sent to the BGP neighbor
Export Policy	Displays the configured export policies for the peer group
Import Policy	Displays the configured import policies for the peer group

Sample output for BGP neighbor received routes

```

A:ALA-12# show router bgp neighbor 10.0.0.16 received-routes
=====
BGP Router ID : 10.0.0.16      AS : 65206   Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP IPv4 Routes
=====
Flag  Network                Nexthop                LocalPref  MED      As-Path

```

```

-----
?   10.0.0.16/32      10.0.0.16      100      none      No As-Path
?   10.0.6.0/24       10.0.0.16      100      none      No As-Path
?   10.0.8.0/24       10.0.0.16      100      none      No As-Path
?   10.0.12.0/24      10.0.0.16      100      none      No As-Path
?   10.0.13.0/24      10.0.0.16      100      none      No As-Path
?   10.0.204.0/24     10.0.0.16      100      none      No As-Path
=====
A:ALA-12#

```

Table 68: Output fields: router BGP neighbor received-routes

Label	Description
BGP Router ID	Displays the local BGP router ID
AS	Displays the configured autonomous system number
Local AS	Displays the configured local AS setting If not configured, then it is the same value as the AS
Flag	u - used s - suppressed h - history d - decayed * - valid i - igp e - egp ? - incomplete > - best
Network	Displays the route IP prefix and mask length for the route
Next Hop	Displays the BGP next hop for the route
LocalPref	Displays the BGP local preference path attribute for the route
MED	Displays the BGP Multi-Exit Discriminator (MED) path attribute for the route
AS Path	Displays the BGP AS path for the route

Sample output — add-path

```

*A:Dut-A# show router bgp neighbor 10.2.2.2
=====
BGP Neighbor
=====
-----
Peer           : 10.2.2.2
Description    : (Not Specified)
Group          : Metro

```

```

-----
Peer AS          : 100          Peer Port          : 179
Peer Address     : 2.2.2.2
Local AS         : 100          Local Port         : 50239
Local Address    : 1.1.1.1
Peer Type        : Internal
State            : Established   Last State         : Active
Last Event       : recvKeepAlive
Last Error       : Cease (Other Configuration Change)
Local Family     : IPv4
Remote Family    : IPv4
Hold Time        : 60           Keep Alive         : 30
Min Hold Time    : 0
Active Hold Time : 60           Active Keep Alive  : 20
Cluster Id       : None
Preference       : 170          Num of Update Flaps : 0
Recd. Paths      : 1
IPv4 Recd. Prefixes : 1        IPv4 Active Prefixes : 1
IPv4 Suppressed Pfxs : 0        VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs : 0        VPN-IPv4 Active Pfxs : 0
Mc IPv4 Recd. Pfxs. : 0        Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs : 0        IPv6 Suppressed Pfxs : 0
IPv6 Recd. Prefixes : 0        IPv6 Active Prefixes : 0
VPN-IPv6 Recd. Pfxs : 0        VPN-IPv6 Active Pfxs : 0
VPN-IPv6 Suppr. Pfxs : 0
Mc IPv6 Recd. Pfxs. : 0        Mc IPv6 Active Pfxs. : 0
Mc IPv6 Suppr. Pfxs : 0        L2-VPN Suppr. Pfxs  : 0
L2-VPN Recd. Pfxs  : 0        L2-VPN Active Pfxs  : 0
MVPN-IPv4 Suppr. Pfxs : 0      MVPN-IPv4 Recd. Pfxs : 0
MVPN-IPv4 Active Pfxs : 0      MDT-SAFI Suppr. Pfxs : 0
MDT-SAFI Recd. Pfxs : 0        MDT-SAFI Active Pfxs : 0
Flow-IPv4 Suppr. Pfxs : 0      Flow-IPv4 Recd. Pfxs : 0
Flow-IPv4 Active Pfxs : 0      Rte-Tgt Suppr. Pfxs  : 0
Rte-Tgt Recd. Pfxs  : 0        Rte-Tgt Active Pfxs  : 0
Backup IPv4 Pfxs     : 0        Backup IPv6 Pfxs     : 0
Mc Vpn Ipv4 Recd. Pf* : 0       Mc Vpn Ipv4 Active P* : 0
Mc Vpn Ipv4 Suppr. P* : 0
Backup Vpn IPv4 Pfxs : 0       Backup Vpn IPv6 Pfxs : 0
Input Queue          : 0        Output Queue         : 0
i/p Messages         : 40        o/p Messages         : 32
i/p Octets            : 411       o/p Octets            : 488
i/p Updates           : 1        o/p Updates           : 1
Flow-IPv6 Suppr. Pfxs : 0      Flow-IPv6 Recd. Pfxs : 0
Flow-IPv6 Active Pfxs : 0
Evpn Suppr. Pfxs     : 0        Evpn Recd. Pfxs      : 0
Evpn Active Pfxs     : 0
MS-PW Suppr. Pfxs    : 0        MS-PW Recd. Pfxs     : 0
MS-PW Active Pfxs    : 0
TTL Security         : Disabled   Min TTL Value         : n/a
Graceful Restart      : Disabled   Stale Routes Time     : n/a
Restart Time          : n/a
Advertise Inactive    : Disabled   Peer Tracking          : Enabled
Advertise Label       : ipv4
Auth key chain         : n/a
Disable Cap Nego       : Disabled   Bfd Enabled            : Enabled
Flowspec Validate     : Disabled   Default Route Tgt      : Disabled
Aigp Metric           : Disabled   Split Horizon          : Disabled
Damp Peer Oscillatio* : Disabled   Update Errors          : 0
GR Notification       : Disabled   Fault Tolerance        : Disabled
Rem Idle Hold Time    : 00h00m00s
Next-Hop Unchanged    : None
L2 VPN Cisco Interop  : Disabled
Local Capability       : RtRefresh MPBGP 4byte ASN
Remote Capability      : RtRefresh MPBGP 4byte ASN

```

```

Local AddPath Capabi*: Send - IPv4 (4)
                        : Receive - IPv4
Remote AddPath Capab*: Send - IPv4
                        : Receive - IPv4
Import Policy          : None Specified / Inherited
Export Policy          : from_prefix_to_bgp
Origin Validation      : N/A
EBGP Link Bandwidth    : n/a
IPv4 Rej. Pfxs         : 0
VPN-IPv4 Rej. Pfxs     : 0
Mc IPv4 Rej. Pfxs      : 0
MVPN-IPv4 Rej. Pfxs    : 0
Flow-IPv4 Rej. Pfxs    : 0
L2-VPN Rej. Pfxs       : 0
Rte-Tgt Rej. Pfxs      : 0
Mc Vpn Ipv4 Rej. Pfxs  : 0
IPv6 Rej. Pfxs         : 0
VPN-IPv6 Rej. Pfxs     : 0
Mc IPv6 Rej. Pfxs      : 0
MVPN-IPv6 Rej. Pfxs    : 0
Flow-IPv6 Rej. Pfxs    : 0
MDT-SAFI Rej. Pfxs     : 0
MS-PW Rej. Pfxs        : 0
Evpn Rej. Pfxs         : 0
-----
Neighbors : 1
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-A#

```

Table 69: Output fields: show neighbor add-path

Label	Description
Peer	Displays the IP address of the configured BGP peer
Group	Displays the BGP peer group to which this peer is assigned
Peer AS	Displays the configured or inherited peer AS for the peer group
Peer Address	Displays the configured address for the BGP peer
Peer Port	Displays the TCP port number used on the far-end system
Local AS	Displays the configured or inherited local AS for the peer group
Local Address	Displays the configured or inherited local address for originating peering for the peer group
Local Port	Displays the TCP port number used on the local system
Peer Type	External — peer type configured as external BGP peers Internal — peer type configured as internal BGP peers
State	Idle — The BGP peer is not accepting connections (Shutdown) is also displayed if the peer is administratively disabled Active — BGP is listening for and accepting TCP connections from this peer Connect — BGP is attempting to establish a TCP connection with this peer

Label	Description
	<p>Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer</p> <p>Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION</p> <p>Established — BGP has successfully established a peering session and is exchanging routing information</p>
Last State	<p>Idle — The BGP peer is not accepting connections</p> <p>Active — BGP is listening for and accepting TCP connections from this peer</p> <p>Connect — BGP is attempting to establish a TCP connections with this peer</p> <p>Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer</p> <p>Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION</p>
Last Event	<p>start — BGP has initialized the BGP neighbor</p> <p>stop — BGP has disabled the BGP neighbor</p> <p>open — BGP transport connection is opened</p> <p>close — BGP transport connection is closed</p> <p>openFail — BGP transport connection failed to open</p> <p>error — BGP transport connection error</p> <p>connectRetry — the connect retry timer expired</p> <p>holdTime — the hold time timer expired</p> <p>keepAlive — the keepalive timer expired</p> <p>recvOpen — BGP has received an OPEN message</p> <p>revKeepalive — BGP has received a KEEPALIVE message</p> <p>recvUpdate — BGP has received an UPDATE message</p> <p>recvNotify — BGP has received a NOTIFICATION message</p> <p>None — no events have occurred</p>
Last Error	Displays the last BGP error and subcode to occur on the BGP neighbor
Local Family	Displays the configured local family value
Remote Family	Displays the configured remote family value
Hold Time	Displays the configured hold-time setting

Label	Description
Keep Alive	Displays the configured keepalive setting
Min Hold Time	Displays the configured minimum hold-time setting
Active Hold Time	Displays the negotiated hold time, if the BGP neighbor is in an established state
Active Keep Alive	Displays the negotiated keepalive time if the BGP neighbor is in an established state
Cluster Id	Displays the configured route reflector cluster ID None — no cluster ID is configured
Preference	Displays the configured route preference value for the peer group
Num of Flaps	Displays the number of route flaps in the neighbor connection
Recd. Prefixes	Displays the number of routes received from the BGP neighbor
Recd. Paths	Displays the number of unique sets of path attributes received from the BGP neighbor
IPv4 Recd. Prefixes	Displays the number of unique sets of IPv4 path attributes received from the BGP neighbor
IPv4 Active Prefixes	Displays the number of IPv4 routes received from the BGP neighbor and active in the forwarding table
IPv4 Suppressed Pfxs	Displays the number of unique sets of IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
VPN-IPv4 Suppr. Pfxs	Displays the number of unique sets of VPN-IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
VPN-IPv4 Recd. Pfxs	Displays the number of unique sets of VPN-IPv4 path attributes received from the BGP neighbor
VPN-IPv4 Active Pfxs	Displays the number of VPN-IPv4 routes received from the BGP neighbor and active in the forwarding table
IPv6 Recd. Prefixes	Displays the number of unique sets of IPv6 path attributes received from the BGP neighbor
IPv6 Active Prefixes	Displays the number of IPv6 routes received from the BGP neighbor and active in the forwarding table
VPN-IPv6 Recd. Pfxs	Displays the number of unique sets of VPN-IPv6 path attributes received from the BGP neighbor

Label	Description
VPN-IPv6 Active Pfxs	Displays the number of VPN-IPv6 routes received from the BGP neighbor and active in the forwarding table
VPN-IPv6 Suppr. Pfxs	Displays the number of unique sets of VPN-IPv6 path attributes received from the BGP neighbor and suppressed due to route damping
Backup IPv4 Pfxs	Displays the number of BGP FRR backup path IPv4 prefixes
Backup IPv6 Pfxs	Displays the number of BGP FRR backup path IPv6 prefixes
Backup Vpn IPv4 Pfxs	Displays the number of BGP FRR backup path VPN IPv4 prefixes
Backup Vpn IPv6 Pfxs	Displays the number of BGP FRR backup path VPN IPv6 prefixes
Input Queue	Displays the number of BGP messages to be processed
Output Queue	Displays the number of BGP messages to be transmitted
i/p Messages	Displays the total number of packets received from the BGP neighbor
o/p Messages	Displays the total number of packets sent to the BGP neighbor
i/p Octets	Displays the total number of octets received from the BGP neighbor
o/p Octets	Displays the total number of octets sent to the BGP neighbor
i/p Updates	Displays the total number of updates received from the BGP neighbor
o/p Updates	Displays the total number of updates sent to the BGP neighbor
TTL Security	Enabled — TTL security is enabled Disabled — TTL security is disabled
Min TTL Value	Displays the minimum TTL value configured for the peer
Graceful Restart	Displays the state of graceful restart
Stale Routes Time	Displays the length of time that stale routes are kept in the route table
Advertise Inactive	Displays the state of advertising inactive BGP routes to other BGP peers (enabled or disabled)

Label	Description
Peer Tracking	Displays the state of tracking a neighbor IP address in the routing table for a BGP session
Advertise Label	Displays the enabled address family for supporting RFC 3107 BGP label capability
Auth key chain	Displays the value for the authentication key chain
Bfd Enabled	Enabled — BFD is enabled Disabled — BFD is disabled
Local Capability	Displays the capability of the local BGP speaker; for example, route refresh, MP-BGP, ORF
Remote Capability	Displays the capability of the remote BGP peer; for example, route refresh, MP-BGP, ORF
Local AddPath Capabi*	Displays the state of the local BGP add-paths capabilities The add-paths capability allows the router to send and receive multiple paths per prefix to or from a peer
Remote AddPath Capab*	Displays the state of the remote BGP add-paths capabilities
Import Policy	Displays the configured import policies for the peer group
Export Policy	Displays the configured export policies for the peer group

next-hop

Syntax

next-hop [*family*] [*ip-address*] [*detail*]

Context

show>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays BGP next-hop information.

Parameters

family

Displays the type of routing information to be distributed by the BGP instance.

- Values
- ipv4

— Displays only those BGP peers that have the IPv4 family enabled.
- vpn-ipv4

—Displays only those BGP peers that have the VPN-IPv4 family enabled.
- ipv6

— Displays only those BGP peers that have the IPv6 family enabled.
- vpn-ipv6

— Displays only those BGP peers that have the VPN-IPv4 family enabled.

ip-address

Displays next-hop information for 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[-interface]

x:x:x:x:x:x.d.d.d.d[-interface]

x:

[0 to FFFF]H

d:

[0 to 255]D

interface:

32 characters maximum,
mandatory for link local addresses.

detail

Displays detailed information.

Output

The following output is an example of BGP next-hop information, and [Table 70: Output fields: router BGP next-hop](#) describes the output fields.

Sample output

```
*A:Dut-C# show router bgp next-hop
=====
BGP Router ID:10.20.1.3      AS:5000      Local AS:5000
=====

=====
BGP Next Hop
=====
Next Hop                                Pref Owner
  Resolving Prefix                      Metric
  Resolved Next Hop                     Ref. Count
-----
10.20.1.1                                7    RSVP
  10.20.1.1/32                          1000
  10.10.2.1                              2
10.20.1.2                                7    RSVP
  10.20.1.2/32                          1000
  10.10.3.2                              2
10.20.1.4                                7    RSVP
  10.20.1.4/32                          1000
```

```
10.10.11.4                                     2
-----
Next Hops : 3

A:ALA-49>show>router>bgp# next-hop 192.168.2.194
-----
BGP Router ID : 10.10.10.104      AS : 200      Local AS : 200
=====
BGP Next Hop
=====
Next Hop      Resolving      Owner  Preference Reference  Resolved
              Prefix                               Count      Next Hop
-----
A:ALA-49>show>router>bgp# next-hop 10.10.10.104
```

Table 70: Output fields: router BGP next-hop

Label	Description
BGP ID	Displays the local BGP router ID
AS	Displays the configured ASN
Local AS	Displays the configured local AS setting. If not configured, then the value is the same as the AS.
Next Hop	Displays the next-hop address
Resolving Prefix	Displays the prefix of the best next hop.
Owner	Displays the routing protocol used to derive the best next hop
Preference	Displays the BGP preference attribute for the routes
Reference Count	Displays the number of routes using the resolving prefix
Resolved Next Hop	Displays the IP address of the next hop

paths

Syntax
paths

Context
show>router>bgp

Platforms
Supported on all 7210 SAS platforms as described in this document

Description

This command displays a summary of BGP path attributes.

Output

The following output is an example of BGP path attributes information, and [Table 71: Output fields: router BGP paths](#) describes the output fields.

Sample output

```
=====
BGP Router ID : 10.0.0.14      AS : 65206   Local AS : 65206
=====
BGP Paths
=====
Path: 60203 65001 19855 3356 15412
-----
Origin       : IGP                Next Hop      : 10.0.28.1
MED          : 60203              Local Preference : none
Refs         : 4                  ASes         : 5
Segments     : 1
Flags        : EBGp-learned
Aggregator   : 15412 62.216.140.1
-----
Path: 60203 65001 19855 3356 1 1236 1236 1236 1236
-----
Origin       : IGP                Next Hop      : 10.0.28.1
MED          : 60203              Local Preference : none
Refs         : 2                  ASes         : 9
Segments     : 1
Flags        : EBGp-learned
```

Table 71: Output fields: router BGP paths

Label	Description
BGP Router ID	Displays the local BGP router ID
AS	Displays the configured autonomous system number
Local AS	Displays the configured local AS setting. If not configured, then the value is the same as the AS.
Path	Displays the AS path attribute
Origin	EGP - The NLRI is learned by an EGP protocol IGP - The NLRI is interior to the originating AS INCOMPLETE - NLRI was learned another way
Next Hop	Displays the advertised BGP next hop
MED	Displays the MED value
Local Preference	Displays the local preference value. This value is used if the BGP route arrives from a BGP peer without the Local Pref attribute set. It is overridden by any value set via a route policy.

Label	Description
Refs	Displays the number of routes using a specified set of path attributes
ASes	Displays the number of autonomous system numbers in the AS path attribute
Segments	Displays the number of segments in the AS path attribute
Flags	EBGP-learned - Path attributes learned by an EBGP peering IBGP-Learned - Path attributes learned by an IBGP peering
Aggregator	Displays the route aggregator ID
Community	Displays the BGP community attribute list
Originator ID	Displays the originator ID path attribute value

routes

Syntax

routes [*family*] [*brief*]

routes [*family*] *prefix* [*detail* | *longer* | *hunt*] [*brief*]

routes [*family*] [*type mvpn-type*] **community** *comm-id*

routes [*family*] [*type mvpn-type*] **aspath-regex** *reg-ex*

routes ms-pw [*rd rd*] [*aii-type2 aii-type2*] [*brief*]

routes l2-vpn *l2vpn-type* {[*rd rd*] | [*siteid site-id*] | [*veid veid*] [*offset vpls-base-offset*]}

routes evpn auto-disc [*hunt* | *detail*] [*rd rd*] [**community** *comm-id*] [**tag** *tag*] [**next-hop** *ip-address*] [**esi** *esi*]

routes evpn eth-seg [*hunt* | *detail*] [*rd rd*] [**community** *comm-id*] [**originator-ip** *ip-address*] [**next-hop** *ip-address*] [**esi** *esi*]

routes evpn inclusive-mcast [*hunt* | *detail*] [*rd rd*] [**community** *comm-id*] [**originator-ip** *ip-address*] [**next-hop** *ip-address*] [**esi** *esi*]

routes evpn inclusive-mcast [*hunt* | *detail*] [*rd rd*] [**community** *comm-id*] [**originator-ip** *ip-address*] [**next-hop** *ip-address*] [**tag** *tag*]

routes evpn mac [*hunt* | *detail*] [*rd rd*] [**next-hop** *ip-address*] [**mac-address** *mac-address*] [**community** *comm-id*] [**tag** *tag*]

Context

show>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays BGP route information.

When this command is issued without any parameters, the entire BGP routing table displays.

When this command is issued with an IP prefix/mask or IP address, the best match for the parameter displays.

Parameters

family

Specifies the type of routing information to be distributed by the BGP instance

- Values
- ipv4

— Displays only those BGP peers that have the IPv4 family enabled.
- vpn-ipv4

— Displays only those BGP peers that have the VPN-IPv4 family enabled.
- ipv6

— Displays only those BGP peers that have the IPv6 family enabled.
- vpn-ipv6

— Displays only those BGP peers that have the VPN-IPv4 family enabled.
- evpn

— Displays only BGP peers that have the EVPN family enabled.

received

Displays the BGP routes received from the neighbor.

prefix

Specifies the type of routing information to display.

Values	
	<i>rd:[ip-address[/mask]]</i>
rd	<i>ip-address:number1</i> <i>as-number1:number2</i> <i>as-number2:number3</i>
number1	1 to 65535
as-number1	1 to 65535
number2	0 to 4294967295
as-number2	1 to 4294967295
number3	0 to 65535
ip-address	a.b.c.d
mask	0 to 32

filter

Specifies route criteria.

Values **hunt** — Displays entries for the specified route in the RIB-In, RIB-Out, and RTM.

longer — Displays the specified route and subsets of the route.

detail — Displays the longer, more detailed version of the output.

aspath-regex *reg-exp*

Displays all routes with an AS path, up to 32 characters, matching the specified regular expression *reg-exp*.

community *comm-id*

Displays all routes with the specified BGP community.

Values *[as-number1:comm-val1 | ext-comm | well-known-comm]*

ext-comm	type:{ip-address:comm-val1 as-number1:comm-val2 as-number2:comm-val1}
as-number1	0 to 65535
comm-val1	0 to 65535
type	target, origin
ip-address	a.b.c.d
comm-val2	0 to 4294967295
as-number2	0 to 4294967295
well-known-comm	no-export, no-export-subconfed, no-advertise

brief

Provides a summarized display of the set of peers to which a BGP route is advertised.

rd

Allows more precise definition of the RD and prefix for VPN-IPv6 routes.

Values ip-addr:comm-val

 2byte-asnumber:ext-comm-val

 4byte-asnumber:comm-val

veid

Specifies a two-byte identifier that represents the local bridging instance in a VPLS and is advertised through the BGP NLRI. This value must be lower than or equal to the *max-ve-id*.

Values 0 to 4294967295

vpls-base-offset

Specifies a two-byte identifier advertised through the NLRI that is used to indicate which VE ID should use the advertised NLRI at the receiving PE according to the following rule:

If the $\text{offset} \leq \text{local VE-ID} \leq \text{offset} + \text{VBS} - 1$ (VBS = virtual block size = 8 in our implementation), the NLRI is processed. Otherwise it is ignored.

The NLRI with this offset is generated as soon as the first VE ID value between (offset, offset + VBS-1) is advertised in the network.

Values 0 to 4294967295

l2vpn-type

Specifies a 12-byte Virtual Switch Instance identifier (VSI-ID) type.

Values bgp-ad | bgp-vpls | multi-homing

ms-pw [rd rd][aii-type2 aii-type2][brief]

Displays routes for the ms-pw family.

Output

The following outputs are examples of BGP route information, and the associated tables describe the output fields.

- [Sample output, Table 72: Output fields: router BGP routes](#)
- [Sample output — BGP PIC, Table 73: Output fields: BGP routes IPv4](#)
- [Sample output — EVPN auto-disc routes](#)
- [Sample output — EVPN eth-seg routes](#)
- [Sample output — EVPN inclusive mcast routes](#)
- [Sample output — EVPN MAC, Table 74: Output fields: BGP EVPN routes](#)

Sample output

```
*A:Dut-C# show router bgp routes hunt 10.1.1.1/32
=====
BGP Router ID:10.20.1.3      AS:5000      Local AS:5000
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * -
valid Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP IPv4 Routes
=====
-----
RIB In Entries
-----
Network       : 10.1.1.1/32
Nexthop       : 10.20.1.1
From          : 10.20.1.1
Res. Nexthop  : 10.20.1.1 (RSVP LSP: 1)
Local Pref.   : 100
Aggregator AS : None
Atomic Aggr.  : Not Atomic
Community     : No Community Members
Originator Id : None
Interface Name : ip-10.10.2.3
Aggregator    : None
MED           : None
Peer Router Id : 10.20.1.1
```

```

Flags          : Used Valid Best Incomplete
AS-Path        : No As-Path

-----
RIB Out Entries
-----
Routes : 1
=====

A:ALA-12>config>router>bgp# show router bgp routes family ipv4
=====
BGP Router ID : 10.10.10.103      AS : 200      Local AS : 200
=====
Legend -
Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Routes
=====
Flag  Network                Nexthop      LocalPref  MED
      VPN Label              As-Path
-----
No Matching Entries Found
=====
A:ALA-12>config>router>bgp#

A:ALA-12>config>router>bgp# show router bgp routes 10.1.0.0/24 de
=====
BGP Router ID : 10.128.0.161 AS : 65535 Local AS : 65535
=====
Legend - Status codes : u - used, s - suppressed, h - history, d - decayed, * -
valid Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Routes
=====
Original Attributes
Network : 10.1.0.0/24 Nexthop :10.20.1.20
Route Dist. : 10070:100 VPN Label :152784
From : 10.20.1.20 Res. Nexthop:10.130.0.2
Local Pref. :100
Aggregator AS:none Aggregator : none
Atomic Aggr.:Not Atomic MED :none
Community :target:10070:1
Originator Id:None Peer Router Id:10.20.1.20
Flags :Used Valid Best IGP
AS-Path :10070 {14730}

Modified Attributes

Network :10.1.0.0/24 Nexthop :10.20.1.20
Route Dist.: 10001:100 VPN Label :152560
From :10.20.1.20 Res. Nexthop :10.130.0.2
Local Pref.:100
Aggregator AS: none Aggregator:none
Atomic Aggr.:Not Atomic MED :none
Community :target:10001:1
Originator Id:None Peer Router Id:10.20.1.20
Flags :Used Valid Best IGP
AS-Path :No As-Path
-----
...

```

```
=====
A:ALA-12>config>router>bgp#

A:7210-12# show router bgp routes 10.0.0.0/30 hunt
=====
BGP Router ID : 10.20.1.1   AS : 100Local AS : 100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Routes
=====
RIB In Entries
-----
Network       : 10.0.0.0/30
Nexthop       : 10.20.1.2
Route Dist.   : 10.20.1.2:1VPN Label: 131070
From          : 10.20.1.2
Res. Nexthop  : 10.10.1.2
Local Pref.   : 100Interface Name: to-sr7
Aggregator AS : noneAggregator: none
Atomic Aggr.  : Not AtomicMED: none
Community     : target:10.20.1.2:1
Originator Id : NonePeer Router Id: 10.20.1.2
Flags         : Used Valid Best IGP
AS-Path       : No As-Path
VPRN Imported : 1 2 10 12
-----
RIB Out Entries
-----
Routes : 1
=====
A:7210-12#

*A:Dut-C>config>router>policy-options# show router bgp routes 10.10.0.0/24 hunt
=====
BGP Router ID:10.20.1.3      AS:300      Local AS:300
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes : i - IGP, e - EGP, ? - incomplete
=====
BGP IPv4 Routes
=====
RIB In Entries
-----
Network       : 10.10.0.0/24
Nexthop       : 10.20.1.2
Path Id       : None
From          : 10.20.1.2
Res. Nexthop  : 10.10.11.2 (LDP)
Local Pref.   : 100
Aggregator AS : None
Atomic Aggr.  : Not Atomic
AIGP Metric   : 555
Connector     : None
Community     : No Community Members
Cluster       : No Cluster Members
Interface Name : INT_T0_C3_D_1
Aggregator     : None
MED            : None
=====
```

```

Originator Id : None                      Peer Router Id : 10.20.1.2
IPv4 Label    : 131065
Flags         : Used Valid Best IGP
Route Source  : Internal
AS-Path       : 400 500
Route Tag     : 0
Neighbor-AS   : 400
Orig Validation: NotFound
Add Paths Send : Default
Last Modified  : 00h15m47s

```

```

Network       : 10.10.0.0/24
Nexthop       : 10.20.1.4
Path Id       : None
From          : 10.20.1.4
Res. Nexthop  : 10.10.5.4 (LDP)
Local Pref.   : 100
Aggregator AS : None                    Interface Name : INT_TO_C4_E_1
Atomic Aggr.  : Not Atomic              Aggregator     : None
AIGP Metric   : None                    MED            : None
Connector     : None
Community     : No Community Members
Cluster       : No Cluster Members

```

```

Originator Id : None                      Peer Router Id : 10.20.1.4
IPv4 Label    : 131065
Flags         : Valid IGP
TieBreakReason : AIGP
Route Source  : Internal
AS-Path       : 400 500
Route Tag     : 0
Neighbor-AS   : 400
Orig Validation: NotFound
Add Paths Send : Default
Last Modified  : 00h15m49s

```

```

Network       : 10.10.0.0/24
Nexthop       : 10.10.1.1
Path Id       : None
From          : 10.10.1.1
Res. Nexthop  : 10.10.1.1
Local Pref.   : None
Aggregator AS : None                    Interface Name : INT_TO_C1_A
Atomic Aggr.  : Not Atomic              Aggregator     : None
AIGP Metric   : None                    MED            : None
Connector     : None
Community     : No Community Members
Cluster       : No Cluster Members

```

```

Originator Id : None                      Peer Router Id : 10.20.1.1
IPv4 Label    : 131071
Flags         : Invalid IGP AS-Loop
Route Source  : External
AS-Path       : 200 300 400 500
Route Tag     : 0
Neighbor-AS   : 200
Orig Validation: NotFound
Add Paths Send : Default
Last Modified  : 00h15m48s

```

RIB Out Entries

```

Network       : 10.10.0.0/24
Nexthop       : 10.20.1.2
Path Id       : None

```

```

To                : 10.20.1.4
Res. Nexthop      : n/a
Local Pref.       : 100
Aggregator AS     : None
Atomic Aggr.      : Not Atomic
AIGP Metric       : 555
Connector         : None
Community         : No Community Members
Cluster           : 10.20.1.3
Originator Id     : 10.20.1.2
IPv4 Label        : 131065
Origin            : IGP
AS-Path           : 400 500
Route Tag         : 0
Neighbor-AS       : 400
Orig Validation   : NotFound

Interface Name    : NotAvailable
Aggregator       : None
MED              : 100

Peer Router Id   : 10.20.1.4

Network          : 10.10.0.0/24
Nexthop          : 10.20.1.2
Path Id          : None
To               : 10.20.1.2
Res. Nexthop     : n/a
Local Pref.      : 100
Aggregator AS    : None
Atomic Aggr.     : Not Atomic
AIGP Metric      : 555
Connector        : None
Community        : No Community Members
Cluster          : 10.20.1.3
Originator Id    : 10.20.1.2
IPv4 Label       : 131065
Origin           : IGP
AS-Path          : 400 500
Route Tag        : 0
Neighbor-AS      : 400
Orig Validation  : NotFound

Interface Name    : NotAvailable
Aggregator       : None
MED              : 100

Peer Router Id   : 10.20.1.2

Network          : 10.10.0.0/24
Nexthop          : 10.20.1.2
Path Id          : None
To               : 10.20.1.5
Res. Nexthop     : n/a
Local Pref.      : 100
Aggregator AS    : None
Atomic Aggr.     : Not Atomic
AIGP Metric      : 555
Connector        : None
Community        : No Community Members
Cluster          : 10.20.1.3
Originator Id    : 10.20.1.2
IPv4 Label       : 131065
Origin           : IGP
AS-Path          : 400 500
Route Tag        : 0
Neighbor-AS      : 400
Orig Validation  : NotFound

Interface Name    : NotAvailable
Aggregator       : None
MED              : None

Peer Router Id   : 10.20.1.5

Network          : 10.10.0.0/24
Nexthop          : 10.10.1.3
Path Id          : None
To               : 10.10.1.1
Res. Nexthop     : n/a
Local Pref.      : n/a
Aggregator AS    : None

Interface Name    : NotAvailable
Aggregator       : None

```

```

Atomic Aggr.   : Not Atomic           MED           : None
AIGP Metric    : None
Connector      : None
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None                 Peer Router Id : 10.20.1.1
IPv4 Label     : 131067
Origin         : IGP
AS-Path        : 300 400 500
Route Tag      : 0
Neighbor-AS    : 300
Orig Validation: NotFound

-----
Routes : 7
=====
*A:Dut-C>config>router>policy-options#

*A:dut-a# show router bgp routes mvpn-ipv4 type source-join source-as 200 source-
ip 150.100.1.2 group-ip 226.0.0.0 detail
=====
BGP Router ID:10.20.1.3      AS:200      Local AS:200
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP MVPN-IPv4 Routes
=====
Route Type      : Source-Join
Route Dist.     : 1:1
Source AS       : 200
Source IP       : 10.100.1.2
Group IP        : 239.0.0.0
Nexthop         : 10.20.1.4
From            : 10.20.1.4
Res. Nexthop    : 0.0.0.0
Local Pref.     : 100
Interface Name  : NotAvailable
Aggregator AS   : None
Aggregator      : None
Atomic Aggr.    : Not Atomic
MED             : 0
Community       : target:10.20.1.3:2
Originator Id   : None                 Peer Router Id : 10.20.1.4
Flags           : Used Valid Best IGP
AS-Path         : No As-Path

-----
Routes : 1
=====
*A:dut-a#

```

Table 72: Output fields: router BGP routes

Label	Description
BGP Router ID	Displays the local BGP router ID
AS	Displays the configured autonomous system number

Label	Description
Local AS	Displays the configured local AS setting If not configured, then the value is the same as the AS
Route Dist.	Displays the route distinguisher identifier attached to routes that distinguishes the VPN it belongs
VPN Label	Displays the label generated by the PE label manager
Network	Displays the IP prefix and mask length
Nexthop	Displays the BGP next hop
From	Displays the advertising BGP neighbor IP address
Res. Nexthop	Displays the resolved next hop
Local Pref.	Displays the local preference value This value is used if the BGP route arrives from a BGP peer without the Local Pref attribute set It is overridden by any value set via a route policy
Flag	u - used s - suppressed h - history d - decayed * - valid i - igp e - egp ? - incomplete > - best
Aggregator AS	Displays the aggregator AS value none - Aggregator AS attributes are not present
Aggregator	Displays the aggregator attribute value none - Aggregator attributes are not present
Atomic Aggr.	Atomic - The atomic aggregator flag is set Not Atomic - The atomic aggregator flag is not set
MED	Displays the MED metric value none - MED metrics are present
Community	Displays the BGP community attribute list

Label	Description
Originator Id	Displays the originator ID path attribute value none - The originator ID attribute is not present
Peer Router Id	Displays the router ID of the advertising router
AS-Path	Displays the BGP AS path attribute
VPRN Imported	Displays the VPRNs where a particular BGP-VPN received route has been imported and installed

Sample output — BGP PIC

```
*A:Dut-A# show router bgp routes ipv4
=====
BGP Router ID:1.1.1.1      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
=====
BGP IPv4 Routes
=====
Flag  Network                LocalPref  MED
      Nexthop (Router)      Path-Id    Label
      As-Path
-----
u*>i  10.77.77.77/32          100        None
      2.2.2.2                1          131066
      200
ub*i  10.77.77.77/32          100        None
      3.3.3.3                None       131068
      200
-----
Routes : 2
=====
*A:Dut-A#
```

Table 73: Output fields: BGP routes IPv4

Label	Description
BGP Router ID	Displays the local BGP router ID
AS	Displays the configured autonomous system number
Local AS	Displays the configured local AS setting. If not configured, then the value is the same as the AS
BGP IPv4 Routes	
Flag	u - used s - suppressed h - history

Label	Description
	d - decayed * - valid i - igp e - egp ? - incomplete > - best
Network	Displays the IP prefix and mask length
Nexthop	Displays the BGP next hop
AS-Path	Displays the BGP AS path attribute
Local Pref.	Displays the local preference value This value is used if the BGP route arrives from a BGP peer without the Local Pref attribute set It is overridden by any value set via a route policy
MED	Displays the MED metric value none - MED metrics are present
Path-Id	Displays the path ID None - The path ID is not present
Label	Displays the MPLS label associated with the BGP route
Routes	Displays the number of routes

Sample output — EVPN auto-disc routes

```
*A:Dut-B# show router bgp routes evpn auto-disc
=====
BGP Router ID:10.20.1.2      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
=====
BGP EVPN Auto-Disc Routes
=====
Flag  Route Dist.      ESI                               NextHop
      Tag                               Label
-----
u*>i  10.20.1.3:1      00:bc:01:00:00:00:00:00:01  10.20.1.3
      0                                LABEL 131069
u*>i  10.20.1.3:1      00:bc:01:00:00:00:00:00:01  10.20.1.3
      MAX-ET                             LABEL 0
u*>i  10.20.1.4:1      00:de:01:00:00:00:00:00:01  10.20.1.4
      0                                LABEL 131069
u*>i  10.20.1.4:1      00:de:01:00:00:00:00:00:01  10.20.1.4
```

```

MAX-ET                                LABEL 0
u*>i 10.20.1.5:1                      00:de:01:00:00:00:00:00:01 10.20.1.5
0                                     LABEL 131059
u*>i 10.20.1.5:1                      00:de:01:00:00:00:00:00:01 10.20.1.5
MAX-ET                                LABEL 0
-----
Routes : 6
=====

```

Sample output — EVPN eth-seg routes

```

*A:Dut-B# show router bgp routes evpn eth-seg
=====
BGP Router ID:10.20.1.2      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
=====
BGP EVPN Eth-Seg Routes
=====
Flag Route Dist.      ESI                      NextHop
     OrigAddr
-----
u*>i 10.20.1.3:0      00:bc:01:00:00:00:00:00:01 10.20.1.3
     10.20.1.3
-----
Routes : 1
=====

```

Sample output — EVPN inclusive mcast routes

```

*A:Dut-B# show router bgp routes evpn inclusive-mcast
=====
BGP Router ID:10.20.1.2      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
=====
BGP EVPN Inclusive-Mcast Routes
=====
Flag Route Dist.      OrigAddr      NextHop
     Tag
-----
u*>i 10.20.1.3:1      10.20.1.3      10.20.1.3
     0
u*>i 10.20.1.4:1      10.20.1.4      10.20.1.4
     0
u*>i 10.20.1.5:1      10.20.1.5      10.20.1.5
     0
-----
Routes : 3
=====

```

Sample output — EVPN MAC

```

*A:Dut-B# show router bgp routes evpn mac
=====
BGP Router ID:10.20.1.2      AS:100      Local AS:100
=====

```

```

=====
Legend -
Status codes   : u - used, s - suppressed, h - history, d - decayed, * - valid
                  l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes   : i - IGP, e - EGP, ? - incomplete
=====
BGP EVPN MAC Routes
=====
Flag  Route Dist.      MacAddr      ESI
      Tag              Mac Mobility  Ip Address
                               NextHop
                               Label1
-----
u*>i  10.20.1.3:1      00:00:00:00:00:03  00:bc:01:00:00:00:00:00:01
      0                Seq:0                N/A
                               10.20.1.3
                               LABEL 131069
-----
Routes : 1
=====

```

Table 74: Output fields: BGP EVPN routes

Label	Description
Flag	u - used s - suppressed h - history d - decayed * - valid i - igp e - egp ? - incomplete > - best
ESI	Displays the Ethernet segment ID value
Route Dist. Tag	Displays the route distinguisher tag
OrigAddr	Displays the BGP originator address
Nexthop	Displays the BGP next hop
MacAddr	Displays the MAC address
Routes	Displays the number of routes

summary

Syntax

```
summary [all]  
summary [family family] [neighbor ip-address]
```

Context

```
show>router>bgp
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays a summary of BGP neighbor information.

The State field displays the global BGP operational state. The valid values are the following:

- **Up**
The BGP global process is configured and running.
- **Down**
The BGP global process is administratively shutdown and not running.
- **Disabled**
The BGP global process is operationally disabled. The process must be restarted by the operator.

For example, if a BGP peer is operationally disabled, the state in the summary table shows the state Disabled.

Parameters

family

Specifies the type of routing information to be distributed by the BGP instance

- Values
- ipv4** — Displays only those BGP peers that have the IPv4 family enabled.
 - vpn-ipv4** — Displays only those BGP peers that have the VPN-IPv4 family enabled.
 - ipv6** — Displays only those BGP peers that have the IPv6 family enabled.
 - vpn-ipv6** — Displays only those BGP peers that have the VPN-IPv4 family enabled.

ip-address

Displays information for entries received from the BGP neighbor.

- 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:x:d.d.d.d[-interface]

 x: [0 to FFFF]H

 d: [0 to 255]D

 interface: 32 characters maximum,
 mandatory for link local addresses.

Output

The following output is an example of summary BGP neighbor information, and [Table 75: Output fields: router BGP summary](#) describes the output fields.

Sample output

```
A:Dut-C# show router bgp summary neighbor 3FFE::A0A:1064
=====
BGP Router ID : 10.20.1.3      AS : 100      Local AS : 100
=====
BGP Admin State      : Up      BGP Oper State      : Up
Number of Peer Groups : 4      Number of Peers      : 5
Total BGP Paths       : 8      Total Path Memory    : 1212
Total BGP Active Rts. : 0      Total BGP Rts.       : 0
Total Suppressed Rts. : 0      Total Hist. Rts.     : 0
Total Decay Rts.      : 0
Total VPN Peer Groups : 0      Total VPN Peers      : 0
Total VPN Local Rts.  : 0
Total VPN Remote Rts. : 0      Total VPN Remote Active Rts.: 0
Total VPN Supp. Rts.  : 0      Total VPN Hist. Rts.  : 0
Total VPN Decay Rts.  : 0
Total IPv6 Remote Rts. : 5      Total IPv6 Rem. Active Rts. : 4
=====
BGP Summary
=====
Neighbor
AS      PktRcvd InQ Up/Down  State|Rcv/Act/Sent (IPv4)
      PktSent OutQ                Rcv/Act/Sent (VpnIPv4)
      Rcv/Act/Sent (IPv6)
      Rcv/Act/Sent (MCastIPv4)
-----
      103      489      0 00h40m28s IPv4 Incapable
      569      0      VPN-IPv4 Incapable
      1/1/3
=====
A:Dut-C#

A:SetupCLI>show>router# bgp summary
=====
BGP Router ID : 10.3.4.5      AS : 35012      Local AS : 100
=====
BGP Admin State      : Up      BGP Oper State      : Up
Confederation AS      : 40000
Member Confederations : 35012 65205 65206 65207 65208
```

```

Rapid Withdrawal      : Disabled
Bfd Enabled           : Yes

Number of Peer Groups : 1      Number of Peers           : 1
Total BGP Paths        : 3      Total Path Memory         : 396
Total BGP Active Rts.  : 0      Total BGP Rts.            : 0
Total Suppressed Rts.  : 0      Total Hist. Rts.          : 0
Total Decay Rts.       : 0

Total VPN Peer Groups  : 1      Total VPN Peers           : 1
Total VPN Local Rts.   : 0      Total VPN Remote Active Rts.: 0
Total VPN Remote Rts.  : 0      Total VPN Hist. Rts.       : 0
Total VPN Supp. Rts.   : 0
Total VPN Decay Rts.   : 0

=====
BGP Summary
=====
Neighbor
      AS      PktRcvd InQ Up/Down  State|Rcv/Act/Sent (IPv4)
              PktSent OutQ              Rcv/Act/Sent (VpnIPv4)
Rcv/Act/Sent (MCastIPv4)
-----
3.3.3.3      20      0    0    01h55m56s Active
              0      0
=====
A:SetupCLI>show>router#

```

Table 75: Output fields: router BGP summary

Label	Description
BGP Router ID	Displays the local BGP router ID
AS	Displays the configured autonomous system number
Local AS	Displays the configured local AS setting. If not configured, then the value is the same as the AS.
BGP Admin State	Down - BGP is administratively disabled Up - BGP is administratively enabled
BGP Oper State	Down - BGP is operationally disabled Up - BGP is operationally enabled
Bfd	Yes - BFD is enabled No - BFD is disabled
Number of Peer Groups	Displays the total number of configured BGP peer groups
Number of Peers	Displays the total number of configured BGP peers
Total BGP Active Routes	Displays the total number of BGP routes used in the forwarding table

Label	Description
Total BGP Routes	Displays the total number of BGP routes learned from BGP peers
Total BGP Paths	Displays the total number of unique sets of BGP path attributes learned from BGP peers
Total Path Memory	Displays the total amount of memory used to store the path attributes
Total Suppressed Routes	Displays the total number of suppressed routes due to route damping.
Total History Routes	Displays the total number of routes with history due to route damping
Total Decayed Routes	Displays total number of decayed routes due to route damping
Total VPN Peer Groups	Displays the total number of configured VPN peer groups
Total VPN Peers	Displays the total number of configured VPN peers
Total VPN Local Rts	Displays the total number of configured local VPN routes
Total VPN Remote Rts	Displays the total number of configured remote VPN routes
Total VPN Remote Active Rts.	Displays the total number of active remote VPN routes used in the forwarding table
Total VPN Supp.Rts.	Displays the total number of suppressed VPN routes due to route damping
Total VPN Hist. Rts.	Displays the total number of VPN routes with history due to route damping
Total VPN Decay Rts.	Displays the total number of decayed routes due to route damping
Neighbor	Displays the BGP neighbor address
AS (Neighbor)	Displays the BGP neighbor autonomous system number
PktRcvd	Displays the total number of packets received from the BGP neighbor
PktSent	Displays the total number of packets sent to the BGP neighbor
InQ	Displays the number of BGP messages to be processed
OutQ	Displays the number of BGP messages to be transmitted

Label	Description
Up/Down	Displays the amount of time that the BGP neighbor has either been established or not established depending on its current state
State Recv/Actv/Sent	Displays the BGP neighbor's current state (if not established) or the number of received routes, active routes and sent routes (if established)

5.24.2.4 Clear commands

damping

Syntax

damping *[[ip-prefix/ip/mask] [neighbor ip-address]]* | *[group name]*

Context

clear>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears or resets the route damping information for received routes.

Parameters

ip-prefix/mask

Clears damping information for entries for the specified IP address.

Values

ipv4-prefix:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 to 32
ipv6-prefix:	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D
ipv6-prefix-length:	0 to 128

ip-address

Clears damping information for entries received from the BGP neighbor.

Values a.b.c.d

group name
Clears damping information for entries received from any BGP neighbors in the peer group, up to 32 characters.

flap-statistics

Syntax
flap-statistics *[[ip-prefix/mask] [neighbor ip-address]] | [group group-name] | [regex reg-exp] | [policy policy-name]*

Context
clear>router>bgp

Platforms
Supported on all 7210 SAS platforms as described in this document

Description
This command clears route flap statistics.

Parameters
ip-prefix/mask
Clears route flap statistics for entries that match the specified IP address.

Values	
ip-prefix:	a.b.c.d (host bits must be 0)
mask:	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]H d: [0 to 255]D
ipv6-prefix-length:	0 to 128

ip-address
Clears route flap statistics for entries received from the specified BGP neighbor.

Values a.b.c.d

group group-name
Clears route flap statistics for entries received from any BGP neighbors in the specified peer group, up to 32 characters.

regex *reg-exp*

Clears route flap statistics for all entries which have the regular expression and the AS path that matches the regular expression, up to 80 characters.

policy *policy-name*

Clears route flap statistics for entries that match the specified route policy, up to 32 characters.

neighbor

Syntax

neighbor {*ip-address* | **as** *as-number* | **external** | **all**} [**soft** | **soft-inbound**]

neighbor{*ip-address* | **as** *as-number* | **external** | **all**} **statistics**

neighbor *ip-address* **end-of-rib**

Context

clear>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command resets the specified BGP peers. This can cause existing BGP connections to be shut down and restarted.

Parameters

ip-address

Resets the BGP neighbor with the specified IP address.

Values ipv4-address: a.b.c.d

as as-number

Resets all BGP neighbors with the specified peer AS.

Values 1 to 65535

external

Keyword to reset all eBGP neighbors.

all

Keyword to reset all BGP neighbors.

soft

The specified BGP neighbors reevaluate all routes in the Local-RIB against the configured export policies.

soft-inbound

The specified BGP neighbors reevaluate all routes in the RIB-In against the configured import policies.

statistics

Keyword to clear the BGP neighbor statistics.

end-of-rib

Keyword to clear the Routing Information Base (RIB).

protocol

Syntax

protocol

Context

clear>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command resets the entire BGP protocol.

5.24.2.5 Debug commands

events

Syntax

events [*neighbor ip-address* | **group name**]

no events

Context

debug>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command logs all events changing the state of a BGP peer.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

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

group *name*

Debugs only events affecting the specified peer group and associated neighbors, up to 32 characters.

keepalive

Syntax

keepalive [**neighbor *ip-addr*** | **group *name***]

no keepalive

Context

debug>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command decodes and logs all sent and received keepalive messages in the debug log.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

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

group *name*

Debugs only events affecting the specified peer group and associated neighbors, up to 32 characters.

notification

Syntax

notification [**neighbor *ip-address*** | **group *name***]

no notification

Context

debug>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command decodes and logs all sent and received notification messages in the debug log.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

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

group *name*

Debugs only events affecting the specified peer group and associated neighbors, up to 32 characters.

open

Syntax

open [**neighbor *ip-address*** | **group *name***]

no open

Context

debug>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command decodes and logs all sent and received open messages in the debug log.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

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

group *name*

Debugs only events affecting the specified peer group and associated neighbors, up to 32 characters.

outbound-route-filtering

Syntax

[no] outbound-route-filtering

Context

debug>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for all BGP outbound route filtering (ORF) packets. ORF is used to inform a neighbor of targets (using target-list) that it is willing to receive.

packets

Syntax

packets [**neighbor** *ip-address* | **group** *name*]

packets

Context

debug>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command decodes and logs all sent and received BGP packets in the debug log.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

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

group *name*

Debugs only events affecting the specified peer group and associated neighbors, up to 32 characters.

route-refresh

Syntax

route-refresh [*neighbor ip-address* | **group name**]

no route-refresh

Context

debug>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables and disables debugging for BGP route-refresh.

Parameters

neighbor ip-address

Debugs only events affecting the specified BGP neighbor.

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

group name

Debugs only events affecting the specified peer group and associated neighbors, up to 32 characters.

rtm

Syntax

rtm [*neighbor ip-address* | **group name**]

no rtm

Context

debug>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command logs RTM changes in the debug log.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

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

group *name*

Debugs only events affecting the specified peer group and associated neighbors, up to 32 characters.

socket

Syntax

socket [**neighbor *ip-address*** | **group *name***]

no socket

Context

debug>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command logs all TCP socket events to the debug log.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

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

group *name*

Debugs only events affecting the specified peer group and associated neighbors, up to 32 characters.

timers

Syntax

timers [**neighbor *ip-address*** | **group *name***]

no timers

Context

debug>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command logs all BGP timer events to the debug log.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

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

group *name*

Debugs only events affecting the specified peer group and associated neighbors, up to 32 characters.

update

Syntax

update [**neighbor *ip-address*** | **group *name***]

no update

Context

debug>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command decodes and logs all sent and received update messages in the debug log.

Parameters

neighbor *ip-address*

Debugs only events affecting the specified BGP neighbor.

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

group *name*

Debugs only events affecting the specified peer group and associated neighbors, up to 32 characters.

6 Route policies

This chapter provides information about configuring route policies.

6.1 Configuring route policies

The 7210 SAS supports two databases for routing information. The routing database is composed of the routing information learned by the routing protocols. The forwarding database is composed of the routes actually used to forward traffic through a router. In addition, link state databases are maintained by interior gateway protocols (IGPs), such as IS-IS and OSPF.

Routing protocols calculate the best route to each destination and place these routes in a forwarding table. The routes in the forwarding table are used to forward routing protocol traffic, sending advertisements to neighbors and peers.

A routing policy can be configured that will not place routes associated with a specific origin in the routing table. Those routes will not be used to forward data packets to the intended destinations and the routes are not advertised by the routing protocol to neighbors and peers.

Routing policies control the size and content of the routing tables, the routes that are advertised, and the best route to take to reach a destination. Careful planning is essential to implement route policies that can affect the flow of routing information or packets in and traversing through the router. Before configuring and applying a route policy, develop an overall plan and strategy to accomplish your intended routing actions.

There are no default route policies. Each policy must be created explicitly and applied to a routing protocol or to the forwarding table. Policy parameters are modifiable.

6.1.1 Policy statements

Route policies contain policy statements containing ordered entries containing match conditions and actions you specify. The entries should be sequenced from the most explicit to least explicit. Packet forwarding and routing can be implemented according to your defined policies. Policy-based routing allows you to dictate where traffic can be routed, through specific paths, or whether to forward or drop the traffic. Route policies can match a specific route policy entry and continue searching for other matches within either the same route policy or the next route policy.

The process can stop when the first complete match is found and executes the action defined in the entry, either to accept or reject packets that match the criteria or proceed to the next entry or the next policy. You can specify matching criteria based on source, destination, or particular properties of a route. Route policies can be constructed to support multiple stages to the evaluation and setting various route attributes. You can also provide more matching conditions by specifying criteria, such as:

- prefix list - a named list of prefixes
- To and From criteria - a route's source and destination

6.1.1.1 Default action behavior

The default action specifies how packets are to be processed when a policy related to the route is not explicitly configured. The following default actions are applied in the event that:

- A route policy does not specify a matching condition, all the routes being compared with the route policy are considered to be matches.
- A packet does not match any policy entries, then the next policy is evaluated. If a match does not occur then the last entry in the last policy is evaluated.
- If no default action is specified, the default behavior of the protocol controls whether the routes match or not.

If a default action is defined for one or more of the configured route policies, then the default action is handled as follows:

- The default action can be set to all available action states including accept, reject, next-entry, and next-policy.
- If the action states accept or reject, then the policy evaluation terminates and the appropriate result is returned.
- If a default action is defined and no matches occurred with the entries in the policy, then the default action is used.
- If a default action is defined and one or more matches occurred with the entries of the policy, then the default action is not used.

6.1.1.2 Denied IP prefixes

The following IP address prefixes are not allowed by the routing protocols and the Route Table Manager and are not be populated within the forwarding table:

- 0.0.0.0/8 or longer
- 127.0.0.0/8 or longer
- 224.0.0.0/4 or longer
- 240.0.0.0/4 or longer

Any other prefixes that need to be filtered can be filtered explicitly using route policies.

6.1.1.3 Controlling route flapping

Route damping is a controlled acceptance of unstable routes from BGP peers so that any ripple effect caused by route flapping across BGP AS border routers is minimized. The motive is to delay the use of unstable routes (flapping routes) to forward data and advertisements until the route stabilizes.

Nokia implementation of route damping is based on the following parameters:

- **Figure of Merit**

A route is assigned a Figure of Merit (FoM), proportional to the frequency of flaps. FoM should be able to characterize a route's behavior over a period of time.

- **route flap**

A route flap is not limited to the withdrawn route. It also applies to any change in the AS path or the next hop of a reachable route. A change in AS path or next hop indicates that the intermediate AS or the route-advertising peer is not suppressing flapping routes at the source or during the propagation. Even if the route is accepted as a stable route, the data packets destined for the route could experience unstable routing because of the unstable AS path or next hop.

- **suppress threshold**

The threshold is a configured value that, when exceeded, the route is suppressed and not advertised to other peers. The state is considered to be down from the perspective of the routing protocol.

- **reuse threshold**

When FoM value falls below a configured reuse threshold and the route is still reachable, the route is advertised to other peers. The FoM value decays exponentially after a route is suppressed. This requires the BGP implementation to decay thousands of routes from a misbehaving peer.

Events that could trigger the route flapping algorithm are:

- **route flapping**

If a route flap is detected within a configured maximum route flap history time, the route's FoM is initialized and the route is marked as a potentially unstable route. Every time a route flaps, the FoM is increased and the route is suppressed if the FoM crosses the suppress threshold.

- **route reuse timer trigger**

A suppressed route's FoM decays exponentially. When it crosses the reuse threshold, the route is eligible for advertisement if it is still reachable.

If the route continues to flap, the FoM, with respect to time scale, looks like a sawtooth waveform with the exponential rise and decay of FoM. To control flapping, the following parameters can be configured:

- **half-life**

The half life value is the time, expressed in minutes, required for a route to remain stable in order for one half of the FoM value to be reduced. For example, if the half life value is 6 (minutes) and the route remains stable for 6 minutes, then the new FoM value is 3. After another 6 minutes passes and the route remains stable, the new FoM value is 1.5.

- **max-suppress**

The maximum suppression time, expressed in minutes, is the maximum amount of time that a route can remain suppressed.

- **suppress**

If the FoM value exceeds the configured integer value, the route is suppressed for use or inclusion in advertisements.

- **reuse**

If the suppress value falls below the configured **reuse** value, then the route can be reused.

6.2 Regular expressions

The ability to perform a filter match on confederations in the AS-PATH is supported. This feature allows customers to configure match criteria for specific confederation sets and sequences within the AS path so that they can be filtered out before cluttering the service provider's routing information base (RIB).

7210 SAS uses regular expression strings to specify match criteria for:

- an AS path string; for example, "100 200 300"
- a community string; for example, "100:200" where 100 is the ASN, and 200 is the community-value
- any AS path beginning with a confederation SET or SEQ containing 65001 and 65002 only: for example "< 65001 65002 >*"
- any AS path containing a confederation SET or SEQ, regardless of the contents: for example, ".*<.*>.*"

A regular expression is expressed in terms of terms and operators. A term for an AS path regular expression is:

1. Regular expressions should always be enclosed in quotes.
2. An elementary term; for example, an ASN "200".
3. A range term composed of two elementary terms separated by the '-' character like "200-300".
4. The '.' dot wild-card character which matches any elementary term.
5. A regular expression enclosed in parenthesis "()".
6. A regular expression enclosed in square brackets used to specify a set of choices of elementary or range terms; for example, [100-300 400] matches any ASN between 100 and 300 or the ASN 400.

A term for a community string regular expression is a string that is evaluated character by character and is composed of:

1. an elementary term which for a community string is any single digit like "4"
2. a range term composed of two elementary terms separated by the '-' character like "2-3"
3. a colon ':' to delimit the ASN from the community value
4. the '.' dot wild-card character which matches any elementary term or ':'
5. a regular expression enclosed in parenthesis "()"
6. a regular expression enclosed in square brackets used to specify a set of choices of elementary or range terms; for example, [1-37] matches any single digit between 1 and 3 or the digit 7

The regular expression operators are listed in the following table.

Table 76: Regular expression operators

Operator	Description
	Matches the term on alternate sides of the pipe.
*	Matches multiple occurrences of the term.
?	Matches 0 or 1 occurrence of the term.
+	Matches 1 or more occurrence of the term.
()	Used to parenthesize so a regular expression is considered as one term.
[]	Used to demarcate a set of elementary or range terms.
-	Used between the start and end of a range.

Operator	Description
{m, n}	Matches least m and at most n repetitions of the term.
{m}	Matches exactly m repetitions of the term.
{m, }	Matches m or more repetitions of the term.
^	Matches the beginning of the string - only allowed for communities.
\$	Matches the end of the string - only allowed for communities.
\	An escape character to indicate that the following character is a match criteria and not a grouping delimiter.

Examples of AS path and community string regular expressions are listed in the following table.

Table 77: AS path and community regular expression examples

AS path to match criteria	Regular expression	Example matches
Null AS path	<code>null</code> ¹⁵	Null AS path
AS path is 11	<code>11</code>	11
AS path is 11 22 33	<code>11 22 33</code>	11 22 33
Zero or more occurrences of ASN 11	<code>11*</code>	Null AS path 11 11 11 11 11 11 11 ... 11
Path of any length that begins with AS numbers 11, 22, 33	<code>11 22 33 .*</code>	11 22 33 11 22 33 400 500 600
Path of any length that ends with AS numbers 44, 55, 66	<code>.* 44 55 66</code>	44 55 66 100 44 55 66 100 200 44 55 66 100 200 300 44 55 66 100 200 300 ... 44 55 66
One occurrence of the AS numbers 100 and 200, followed by one or more occurrences of the number 33	<code>100 200 33+</code>	100 200 33 100 200 33 33 100 200 33 33 33

¹⁵ The `null` keyword matches an empty AS path.

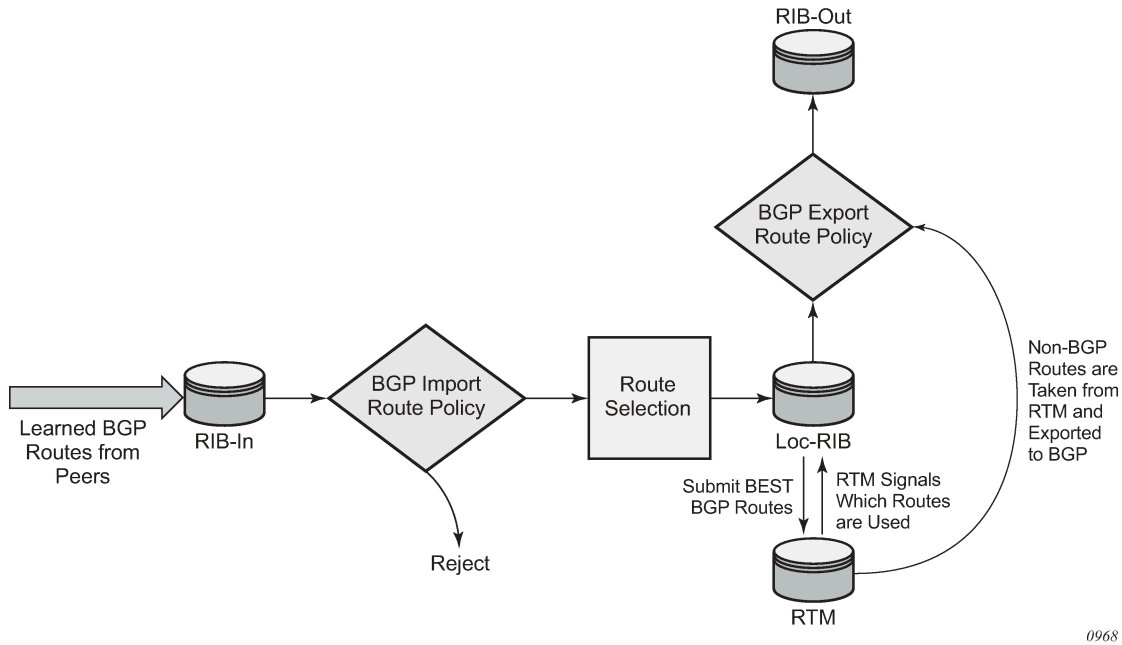
AS path to match criteria	Regular expression	Example matches
		100 200 33 33 33 ... 33
One or more occurrences of ASN 11, followed by one or more occurrences of ASN 22, followed by one or more occurrences of ASN 33	11+ 22+ 33+	11 22 33 11 11 22 33 11 11 22 22 33 11 11 22 22 33 33 11 ... 11 22 ... 22 33 ...33
Path whose second ASN must be 11 or 22	(. 11) (. 22) .* or . (11 22) .*	100 11 200 22 300 400 ...
Path of length one or two whose second ASN might be 11 or 22	. (11 22)?	100 200 11 300 22
Path whose first ASN is 100 and second ASN is either 11 or 22	100 (11 22) .*	100 11 100 22 200 300
Either AS path 11, 22, or 33	[11 22 33]	11 22 33
Range of AS numbers to match a single ASN	10 - 14	10 or 11 or 12 or 13 or 14
	[10-12]*	Null AS path 10 or 11 or 12 10 10 or 10 11 or 10 12 11 10 or 11 11 or 11 12 12 10 or 12 11 or 12 12 ...
Zero or one occurrence of ASN 11	11? or 11{0,1}	Null AS path 11
One through four occurrences of ASN 11	11{1,4}	11 11 11 11 11 11 11 11 11 11
One through four occurrences of ASN 11 followed by one occurrence of ASN 22	11{1,4} 22	11 22 11 11 22

AS path to match criteria	Regular expression	Example matches
		11 11 11 22 11 11 11 11 22
Path of any length, except nonexistent, whose second ASN can be anything, including nonexistent	<code>. .* or . {0,}</code>	100 100 200 11 22 33 44 55
ASN is 100. Community value is 200.	<code>^100:200\$</code>	100:200
ASN is 11 or 22. Community value is any number.	<code>^((11) (22)): (.*)\$</code>	11:100 22:100 11:200 ...
ASN is 11. Community value is any number that starts with 1.	<code>^11:(1.*)\$</code>	11:1 11:100 11:1100 ...
ASN is any number. Community value is any number that ends with 1, 2, or 3.	<code>^(.*):(.*[1-3])\$</code>	11:1 100:2002 333:55553 ...
ASN is 11 or 22. Community value is any number that starts with 3 and ends with 4, 5 or 9.	<code>^((11) (22)):(3.*[459])\$</code>	11:34 22:3335 11:3777779 ...
ASN is 11 or 22. Community value ends in 33 or 44.	<code>[^((11 22)):(.*((33) (44)))\$</code>	11:33 22:99944 22:555533 ...

6.2.1 BGP and OSPF route policy support

BGP and OSPF requires route policy support. [Figure 24: BGP route policy diagram](#) and [Figure 25: OSPF route policy diagram](#) show where route policies are evaluated in the protocol. [Figure 24: BGP route policy diagram](#) shows BGP which applies a route policy as an internal part of the BGP route selection process. [Figure 25: OSPF route policy diagram](#) shows OSPF which applies routing policies at the edge of the protocol, to control only the routes that are announced to or accepted from the Route Table Manager (RTM).

Figure 24: BGP route policy diagram



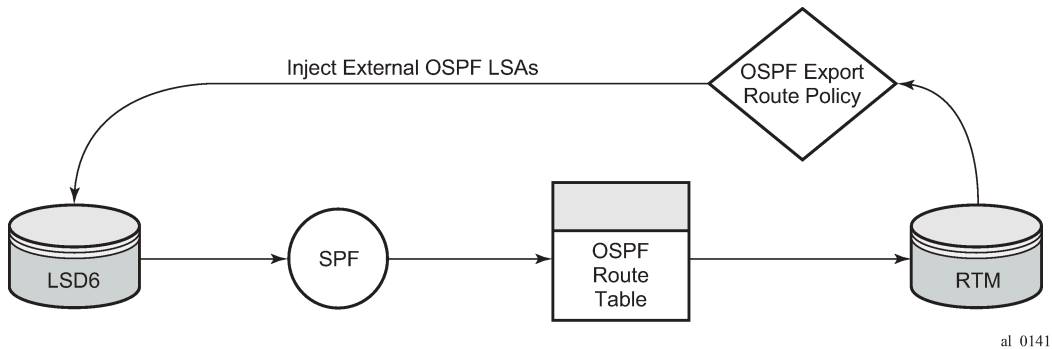
6.2.1.1 BGP route policies

The Nokia implementation of BGP uses route policies extensively. The implied or default route policies can be overridden by customized route policies. The default BGP properties, with no route policies configured, behave as follows:

- Accept all BGP routes into the RTM for consideration.
- Announce all used BGP learned routes to other BGP peers
- Announce none of the IGP, static or local routes to BGP peers.

The following figure shows the OSPF route policy.

Figure 25: OSPF route policy diagram



6.2.1.2 Re-advertised route policies

Occasionally, BGP routes may be readvertised from BGP into OSPF, IS-IS. OSPF export policies control which routes are exported to OSPF) are not handled by the main OSPF task but are handled by a separate task or an RTM task that filters the routes before they are presented to the main OSPF task.

6.2.2 When to use route policies

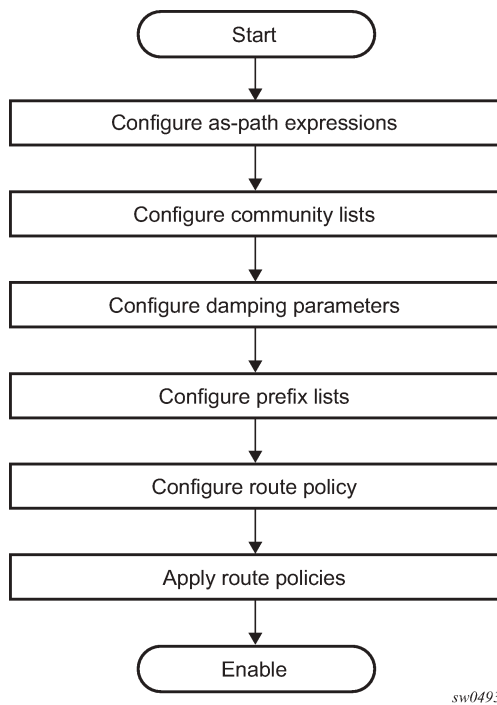
The following are examples of circumstances of when to configure and apply unique route policies:

- When you want to control the protocol to allow all routes to be imported into the routing table. This enables the routing table to learn about particular routes to enable packet forwarding and redistributing packets into other routing protocols.
- When you want to control the exporting of a protocol's learned active routes.
- When you want a routing protocol to announce active routes learned from another routing protocol, which is sometimes called route redistribution.
- Route policies can be used to filter IGMP membership reports from specific hosts and/or specific multicast groups.
- When you want unique behaviors to control route characteristics. For example, change the route preference.
- When you want unique behaviors to control route characteristics. For example, change the route preference, AS path, or community values to manipulate the control the route selection.
- When you want to control BGP route flapping (damping).

6.3 Route policy configuration process overview

The following figure shows the process to provision basic route policy parameters.

Figure 26: Route policy configuration and implementation flow



6.4 Configuration notes

This section describes route policy configuration caveats.

6.4.1 General

When configuring policy statements, the policy statement name must be unique.

6.5 Configuring route policies with CLI

This section provides information to configure route policies using the command line interface.

6.6 Route policy configuration overview

Route policies allow you to configure routing according to specifically defined policies. You can create policies and entries to allow or deny paths based on various parameters such as destination address.

Policies can be as simple or complex as required. A simple policy can block routes for a specific location or IP address. More complex policies can be configured using numerous policy statement entries containing

matching conditions to specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

6.6.1 When to create routing policies

Route policies are created in the **config>router** context. There are no default route policies. Each route policy must be explicitly created and applied. Applying route policies can introduce more efficiency as well as more complexity to 7210 SAS routers' capabilities.

A route policy impacts the flow of routing information or packets within and through the router. A routing policy can be specified to prevent a particular customer's routes to be placed in the route table which causes those routes to not forward traffic to various destinations and the routes are not advertised by the routing protocol to neighbors.

Route policies can be created to control the following:

- a protocol to export all the active routes learned by that protocol
- route characteristics to control which route is selected to act as the active route to reach a destination and advertise the route to neighbors
- protocol to import all routes into the routing table; a routing table must learn about particular routes to be able to forward packets and redistribute to other routing protocols
- to filter IGMP membership reports from specific hosts and/or specific multicast groups
- damping

Before a route policy is applied, analyze the policy's purpose and be aware of the results (and consequences) when packets match the specified criteria and the associated actions and default actions, if specified, are executed. Membership reports can be filtered based on a specific source address.

6.6.2 Default route policy actions

Each routing protocol has default behaviors for the import and export of routing information. The following table describes the default behavior for each routing protocol.

Table 78: Default route policy actions

Protocol	Import	Export
OSPF	Not applicable. All OSPF routes are accepted from OSPF neighbors and cannot be controlled via route policies.	<ul style="list-style-type: none"> • Internal routes: All OSPF routes are automatically advertised to all neighbors. • External routes: By default all non-OSPF learned routes are not advertised to OSPF neighbors
IS-IS	Not applicable. All IS-IS routes are accepted from IS-IS neighbors and can not be controlled via route policies	<ul style="list-style-type: none"> • Internal routes: All IS-IS routes are automatically advertised to all neighbors. • External routes: By default all non-IS-IS learned routes are not advertised to IS-IS peers.
BGP	By default, all routes from BGP.	<ul style="list-style-type: none"> • Internal routes: By default all active BGP routes are advertised to BGP peers

Protocol	Import	Export
		<ul style="list-style-type: none">External routes: By default all non-BGP learned routes are not advertised to BGP peers.

6.6.3 Policy evaluation

Routing policy statements can consist of as few as one or several entries. The entries specify the matching criteria. A route is compared to the first entry in the policy statement. If it matches, the specified entry action is taken, either accepted or rejected. If the action is to accept or reject the route, that action is taken and the evaluation of the route ends.

If the route does not match the first entry, the route is compared to the next entry (if more than one is configured) in the policy statement. If there is a match with the second entry, the specified action is taken. If the action is to accept or reject the route, that action is taken and the evaluation of the route ends, and so on.

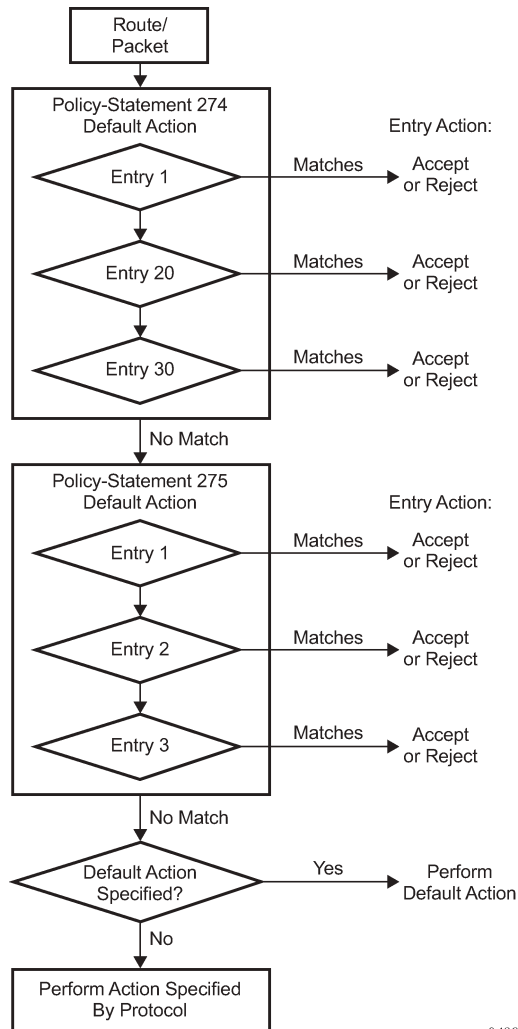
Each route policy statement can have a default-action clause defined. If a default-action is defined for one or more of the configured route policies, then the default actions should be handled in the following ways:

- The process stops when the first complete match is found and executes the action defined in the entry.
- If the packet does not match any of the entries, the system executes the default action specified in the policy statement.

The following figure shows an example of the route policy process.

Route policies can also match a specific route policy entry and continue to search for other entries within either the same route policy or the next route policy by specifying the *next-entry* or *next-policy* option in the entry's **action** command. Policies can be constructed to support multiple states to the evaluation and setting of various route attributes.

Figure 27: Route policy process example



sw0489

6.6.4 Damping

Damping initiates controls when routes flap. Route flapping can occur when an advertised route between nodes alternates (flaps) back and forth between two paths due to network problems which cause intermittent route failures. It is necessary to reduce the amount of routing state change updates propagated to limit processing requirements. Therefore, when a route flaps beyond a configured value (the suppress value), then that route is removed from the routing tables and routing protocols until the value falls below the reuse value.

A route can be suppressed according to the Figure of Merit (FoM) value. The FoM is a value that is added to a route each time it flaps. A new route begins with an FoM value of 0.


Damping is optional. If damping is configured, the following parameter values must be explicitly specified as there are no default values:

- **suppress**

- [half-life](#)
- [reuse](#)
- [max-suppress](#)

When a route's FoM value exceeds the suppress value, then the route is removed from the routing table. The route is considered to be stable when the FoM drops below the reuse value by means of the specified half life parameter. The route is returned to the routing tables. When routes have higher FoM and half life values, they are suppressed for longer periods of time. The following figure shows an example of a flapping route, the suppress threshold, the half life decay (time), and reuse threshold. The peaks represent route flaps, the slopes represent half life decay.

Figure 28: Damping example



6.7 Basic configurations

This section provides information to configure route policies and configuration examples of common tasks. The minimal route policy parameters that need to be configured are:

- policy statement with the following parameters specified:
 - at least one entry
 - entry action

Example: Route policy configuration output

```
A:ALA-B>config>router>policy-options# info
-----
. . .

    policy-statement "aggregate-customer-peer-only"
      entry 1
        from
          community "all-customer-announce"
        exit
        action accept
        exit
      exit
    default-action reject
    exit
```

```
exit
-----
A:ALA-B>config>router>policy-options#
A:ALA-B>config>router>policy-options#info
-----
    prefix-list "host"
        prefix 10.0.0.0/8 longer
    exit
    prefix-list "group"
        prefix 239.6.6.6/32 exact
    exit
policy-statement "block-igmp"
    description "Reject-Reports-From-Specific-Group-And-Host"
    entry 1
        from
            host-ip "host"
        exit
        action next-entry
    exit
    entry 2
        from
            group-address "group"
        exit
        action reject
    exit
    default-action accept
    exit
exit
policy-statement "permit-igmp"
    description "Accept-Reports-From-Specific-Group-And-Host"
    entry 1
        from
            host-ip "host3"
            group-address "group3"
        exit
        action accept
    exit
    exit
    default-action reject
    exit
-----
A:ALA-B>config>router>policy-options#
```

6.8 Configuring route policy components

This section describes the CLI syntax used to configure route policy components.

6.8.1 Beginning the policy statement

Use the following syntax to begin a policy statement configuration. In order for a policy statement to be complete an entry must be specified (see [Configuring an entry](#)).

```
config>router>policy-options
begin
policy-statement name
description text
```


Example: Error message

The following error message displays when the you try to modify a policy options command without entering **begin** first.

```
A:ALA-B>config>router>policy-options# policy-statement "allow all"
MINOR: CLI The policy-
options must be in edit mode by calling begin before any changes can be made.
```

Example: Command usage

The following example displays policy statement configuration command usage. These commands are configured in the **config>router** context.

```
config>router# policy-options
policy-options# begin
```

There are no default policy statement options. All parameters must be explicitly configured.

6.8.2 Creating a route policy

To enter the mode to create or edit route policies, you must enter the **begin** keyword at the **config>router>policy-options** prompt. Other editing commands include:

- The **commit** command saves changes made to route policies during a session.
- The **abort** command discards changes that have been made to route policies during a session.

Example

The following error message displays when the you try to modify a policy options command without entering **begin** first.

```
A:ALA-B>config>router>policy-options# policy-statement "allow all"
MINOR: CLI The policy-
options must be in edit mode by calling begin before any changes can

A:ALA-B>config>router>policy-options# info
#-----
# Policy
#-----

        policy-options
        begin
        policy-statement "allow all"
description "General Policy"
...
        exit
exit
-----
A:ALA-B>config>router>policy-options#
```

6.8.3 Configuring a default action

Specifying a default action is optional. The default action controls those packets not matching any policy statement entries. If no default action is specified for the policy, then the action associated with the protocol to which the routing policy was applied is performed.

A policy statement must include at least one entry (see [Configuring an entry](#)).

To enter the mode to create or edit route policies, you must enter the **begin** keyword at the **config>router>policy-options** prompt. Other editing commands include:

- The **commit** command saves changes made to route policies during a session.
- The **abort** command discards changes that have been made to route policies during a session.

6.8.4 Configuring an entry

An entry action must be specified. The other parameters in the **entry action** context are optional. See [Route policy command reference](#) for the commands and syntax.

Example

The following is a sample configuration output of entry parameters and includes the default action parameters which were displayed in [Configuring a default action](#).

```
A:ALA-B>config>router>policy-options# info
-----
    prefix-list "host"
      prefix 10.0.0.0/8 longer
    exit
    prefix-list "group"
      prefix 239.6.6.6/32 exact
    exit
    policy-statement "block-igmp"
      description "Reject-Reports-From-Specific-Group-And-Host"
      entry 1
        from
          host-ip "host"
        exit
        action next-entry
      exit
      entry 2
        from
          group-address "group"
        exit
        action reject
      exit
      default-action accept
    exit
  exit
-----
A:ALA-B>config>router>policy-options#
```

6.8.5 Configuring damping



Note:

- For each damping profile, all parameters must be configured.
- The *suppress* value must be greater than the *reuse* value (see [Figure 28: Damping example](#)).
- Damping can be enabled in the **config>router>bgp** context on the BGP global, group, and neighbor levels. If damping is enabled, but route policy does not specify a damping profile, the default damping profile will be used. This profile is always present and consists of the following parameters:

half-life:	15 minutes
max-suppress:	60 minutes
suppress:	3000
reuse:	750

Example: Damping configuration output

```
*A:cses-A13>config>router>policy-options# info
-----
      damping "damptest123"
        half-life 15
        max-suppress 60
        reuse 750
        suppress 1000
      exit
-----
*A:cses-A13>config>router>policy-options#
```

6.8.5.1 Configuring a prefix list

Example: Prefix list configuration output

```
A:ALA-B>config>router>policy-options# info
-----
      prefix-list "western"
        prefix 10.10.0.1/32 exact
        prefix 10.10.0.2/32 exact
        prefix 10.10.0.3/32 exact
        prefix 10.10.0.4/32 exact
      exit
-----
A:ALA-B>config>router>policy-options#
A:ALA-B>config>router>policy-options# info
-----
      prefix-list "host"
        prefix 10.0.0.0/8 longer
      exit
      prefix-list "group"
        prefix 239.6.6.6/32 exact
      exit
-----
```

```
A:ALA-B>config>router>policy-options#
```

6.9 Route policy configuration management tasks

This section describes the route policy configuration management tasks.

6.9.1 Editing policy statements and parameters

Route policy statements can be edited to modify, add, or delete parameters. To enter the mode to edit route policies, you must enter the **begin** keyword at the **config>router>policy-options** prompt. Other editing commands include:

- The **commit** command saves changes made to route policies during a session.
- The **abort** command discards changes that have been made to route policies during a session.

Example: Changed configuration output

```
A:ALA-B>config>router>policy-options>policy-statement# info
-----
      description "Level 1"
      entry 1
        to
          neighbor 10.10.10.104
        exit
        action accept
        exit
      exit
      entry 2
        from
          protocol ospf
        exit
        to
          protocol ospf
          neighbor 10.10.0.91
        exit
        action accept
        exit
      exit
      entry 4
        description "new entry"
        from
          protocol isis
          area 0.0.0.20
        exit
        action reject
      exit
      default-action accept
      metric set 10
      exit
-----
A:ALA-B>config>router>policy-options# info
-----
      prefix-list "host"
        prefix 10.0.0.0/8 longer
      exit
      prefix-list "group1"
        prefix 239.6.6.8/32 exact
```

```
exit
policy-statement "block-igmp"
  description "Reject-Reports-From-Specific-Group-And-Host"
  entry 1
    from
      host-ip "host"
    exit
    action next-entry
    exit
  exit
  entry 2
    from
      group-address "group1"
    exit
    action reject
  exit
  default-action accept
  exit
exit
-----
A:ALA-B>config>router>policy-options#
```

6.9.2 Deleting an entry

Use the following syntax to delete a policy statement entry.

```
config>router>policy-options
begin
commit
abort
policy-statement name
no entry entry-id
```

Example: Command usage to delete a policy statement entry

```
config>router>policy-options# begin
policy-options# policy-statement "1"
policy-options>policy-statement# no entry 4
policy-options>policy-statement# commit
```

6.9.3 Deleting a policy statement

Use the following syntax to delete a policy statement.

```
config>router>policy-options
begin
commit
abort
no policy-statement name
```

Example: Command usage to delete a policy statement

```
config>router>policy-options# begin
policy-options# no policy-statement 1
policy-options# commit
```

6.9.4 Use of route policies for IGMP filtering

Example

The following is a sample route policy configuration output for IGMP filtering. This policy needs to be configured with a SAP for filtering to take effect.

```
-----  
A:ALA-B>config>router>policy-options#info  
-----  
prefix-list "host"  
    prefix 10.0.0.0/8 longer  
exit  
prefix-list "group"  
    prefix 239.6.6.6/32 exact  
exit  
  
policy-statement "block-igmp"  
    description "Reject-Reports-From-Specific-Group-And-Host"  
    entry 1  
        from  
            host-ip "host"  
        exit  
        action next-entry  
        exit  
    exit  
    entry 2  
        from  
            group-address "group"  
        exit  
        action reject  
        exit  
    default-action accept  
    exit  
exit  
  
policy-statement "permit-igmp"  
    description "Accept-Reports-From-Specific-Group-And-Host"  
    entry 1  
        from  
            host-ip "host3"  
            group-address "group3"  
        exit  
        action accept  
        exit  
    exit  
    default-action reject  
exit  
-----  
A:ALA-B>config>router>policy-options#
```

6.10 Route policy command reference

6.10.1 Command hierarchies

- [Route policy configuration commands](#)
- [Show commands](#)

6.10.1.1 Route policy configuration commands

```

config
- [no] router [router-name]
- [no] triggered-policy
- [no] policy-options
  - abort
  - as-path name expression regular-expression
  - no as-path name
  - begin
  - commit
  - community name members comm-id [comm-id ... (up to 15 max)]
  - no community name [members comm-id]
  - [no] damping name
    - half-life minutes
    - no half-life
    - max-suppress minutes
    - no max-suppress
    - reuse integer
    - no reuse
    - suppress integer
    - no suppress
  - [no] policy-statement name
    - default-action {accept | next-entry | reject}
    - no default-action
      - aigp-metric metric
      - aigp-metric metric add
      - aigp-metric igp
      - no aigp-metric
      - as-path {add | replace} name
      - no as-path
      - as-path-prepend as-number [repeat]
      - no as-path-prepend
      - community {{add name [remove name]} | {remove name [add name]} |
{replace name}}
      - no community
      - damping {name | none}
      - no damping
      - local-preference local-preference
      - no local-preference
      - metric {add | subtract | set} metric
      - no metric
      - [no] next-hop-self
      - origin {igp | egp | incomplete}
      - no origin
      - preference preference
      - tag
      - type
    - description description-string

```

```

- no description
- [no] entry entry-id
  - action {accept | next-entry | next-policy | reject}
  - no action
    - aigp-metric metric
    - aigp-metric metric add
    - aigp-metric igp
    - no aigp-metric
    - as-path {add | replace} name
    - no as-path
    - as-path-prepend as-number [repeat]
    - no as-path-prepend
    - community {{add name [remove name]} | {remove name [add name]} |
{replace name}}
    - no community
    - damping {name | none}
    - no damping
    - local-preference local-preference
    - no local-preference
    - metric {add | subtract | set} metric
    - no metric
    - [no] next-hop-self
    - origin {igp | egp | incomplete}
    - no origin
    - [no] preference preference
    - [no] tag
    - [no] type
  - description description-string
- no description
- [no] from
  - [no] area
  - [no] as-path name
  - [no] as-path-group name
  - as-pathcommunity name
  - no as-pathcommunity
  - [no] external
  - family [ipv4] [vpn-ipv4][l2-vpn] [ms-pw] [route-target]
  - no family
  - group-address prefix-list-name
  - no group-address
  - [no] host-ip prefix-list-name
  - prefix-list name [name...(up to 5 max)]
  - no prefix-list
  - level {1 | 2}
  - no level
  - neighbor {ip-address | prefix-list name}
  - no neighbor
  - source-address ip-address
  - no source-address
  - [no] protocol protocol [all | {instance instance}]
  - [no] tag tag
  - no tag
  - type type
  - no type
- [no] to
  - level {1 | 2}
  - no level
  - neighbor {ip-address | prefix-list name}
  - no neighbor
  - [no] prefix-list name [name...(up to 5 max)]
  - protocol protocol [all | {instance instance}]

```



```
- no protocol
```

```
config
- [no] router
- [no] policy-options
- [no] prefix-list name
- prefix ip-prefix/prefix-length [exact | longer | through length | prefix-
length-range length1-length2]
- no prefix [ipv-prefix/prefix-length] [exact | longer | through length |
prefix-length-range length1-length2]
```

6.10.1.2 Show commands

```
show
- router
- policy [name | prefix-list name | admin]
```

6.10.2 Command descriptions

6.10.2.1 Route policy command reference

6.10.2.1.1 Generic commands

abort

Syntax

abort

Context

config>router>policy-options

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command discards changes made to a route policy.

begin

Syntax

begin

Context

config>router>policy-options

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enters the mode to create or edit route policies.

```
commit
```

Syntax

commit

Context

config>router>policy-options

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command saves changes made to a route policy.

```
description
```

Syntax

description *string*

no description

Context

config>router>policy-options>policy-statement

config>router>policy-options>policy-statement>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a text description, which is stored in the configuration file, to help identify the content of the entity.

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

Parameters

string

Specifies the description character string. Allowed values are any string up to 80 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

6.10.2.1.2 Route policy options

as-path

Syntax

as-path *name* **expression** *regular-expression*

no as-path *name*

Context

config>router>policy-options

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a route policy autonomous system (AS) path regular expression statement to use in route policy entries.

The **no** form of this command deletes the AS path regular expression statement.

Parameters

name

Specifies the AS path regular expression name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

regular-expression

Specifies the AS path regular expression. Allowed values are any string up to 256 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. A value of **null** specifies the AS path expressed as an empty regular expression string.

community

Syntax

community *name* **members** *comm-id* [*comm-id...*(up to 15 max)]
no community *name* [**members** *comm-id*]

Context

config>router>policy-options

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a route policy community list to use in route policy entries.
The **no** form of this command deletes the community list or the provided community ID.

Default

no community

Parameters

- name**
Specifies the community list name. Allowed values are any string up to 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.
- comm-id**
Specifies the community ID. Note that up to 15 community ID strings can be specified up to a total maximum of 72 characters.

Values	
72 chars maximum	
2byte-asnumber:comm-val reg-ex ext-comm well-known-comm	
ext-comm	type:{ip-address:comm-val reg-ex1®-ex2 ip-address®-ex2 2byte-asnumber:ext-comm-val 4byte-asnumber:comm-val}
2byte-asnumber	0..65535
comm-val	0..65535
reg-ex	72 chars maximum
type	target, origin

ip-address	a.b.c.d
ext-comm-val	0..4294967295
4byte-asnumber	0..4294967295
reg-ex1	63 chars max
reg-ex2	63 chars max
well-known-comm	null, no-export,no-export-subconfed, no-advertise

A community ID can be specified in different forms:

- as-num:comm.-value* — The *as-num* is the autonomous system number (ASN)

Values:	as-num:	1 to 65535
	comm-value:	0 to 65535

- type **{target | origin}** *as-num:comm.-value* — The keywords **target** or **origin** denote the community as an extended community of type route target or route origin respectively. The *as-num* and *comm-value* values allow the same preceding values for regular community values.
- reg-ex1 reg-ex2* — These values are a regular expression string. Allowed values are any string up to 63 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.
- well-known-comm* — Keywords **null**, **no-export**, **no-export-subconfed**, **no-advertise**.

policy-options

Syntax
[no] policy-options

Context
config>router

Platforms
Supported on all 7210 SAS platforms as described in this document

Description
Commands in this context configure route policies. Route policies are applied to the routing protocol used for IGMP group membership report filtering.
The **no** form of this command deletes the route policy configuration.

triggered-policy

Syntax

[no] **triggered-policy**

Context

config>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command triggers route policy re-evaluation.

By default, when a change is made to a policy in the **config router policy options** context and 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 router, the consequences could be dramatic. It is more effective to control changes on a peer by peer basis.

If the **triggered-policy** command is enabled, 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. That is, when a **triggered-policy** is enabled, any routine policy change or policy assignment change within the protocol will not take effect until the protocol is reset or a clear command is issued to re-evaluate route policies; for example, **clear router bgp neighbor x.x.x.x soft**. This keeps the peer up and the change made to a route policy is applied only to that peer, or group of peers.

6.10.2.1.3 Route policy damping commands

damping

Syntax

[no] **damping** *name*

Context

config>router>policy-options

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates the context to configure a route damping profile to use in route policy entries.

The **no** form of this command deletes the named route damping profile.

Parameters

name

Specifies the damping profile 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.

half-life

Syntax

half-life *minutes*

no half-life

Context

config>router>policy-options>damping

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the **half-life** parameter for the route damping profile.

The half life value is the time, expressed in minutes, required for a route to remain stable for the Figure of Merit (FoM) value to be reduced by one half; for example, if the half life value is 6 (minutes) and the route remains stable for 6 minutes, the new FoM value is 3 (minutes). After another 3 minutes pass and the route remains stable, the new FoM value is 1.5 (minutes).

When the FoM value falls below the [reuse](#) threshold, the route is again considered valid and can be reused or included in route advertisements. No half life value is specified. The half life value must be explicitly configured.

The **no** form of this command removes the half life parameter from the damping profile.

Parameters

minutes

Specifies the half life, in minutes, expressed as a decimal integer.

Values 1 to 45

max-suppress

Syntax

max-suppress *minutes*

no max-suppress

Context

config>router>policy-options>damping

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum suppression parameter for the route damping profile.

This value indicates the maximum time, expressed in minutes, that a route can remain suppressed.

The **no** form of this command removes the maximum suppression parameter from the damping profile.

Parameters

minutes

Specifies the maximum suppression time, in minutes, expressed as a decimal integer.

Values 1 to 720

reuse

Syntax

reuse *integer*

no reuse

Context

config>router>policy-options>damping

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the reuse parameter for the route damping profile.

When the FoM value falls below the **reuse** threshold, the route is again considered valid and can be reused or included in route advertisements.

The **no** form of this command removes the reuse parameter from the damping profile.

Parameters

integer

Specifies the reuse value, expressed as a decimal integer.

Values 1 to 20000

suppress

Syntax

suppress *integer*

no suppress

Context

config>router>policy-options>damping

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the suppression parameter for the route policy damping profile.

A route is suppressed when it has flapped frequently enough to increase the FoM value to exceed the **suppress** threshold limit. When the FoM value exceeds the **suppress** threshold limit, the route is removed from the route table or inclusion in advertisements.

The **no** form of this command removes the suppress parameter from the damping profile.

Parameters

integer

Specifies the suppress value, expressed as a decimal integer.

Values 1 to 20000

6.10.2.1.4 Route policy prefix commands

prefix-list

Syntax

[no] **prefix-list** *name*

Context

config>router>policy-options

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure a prefix list to use in route policy entries.

The **no** form of this command deletes the named prefix list.

Parameters

name

Specifies the prefix list name. Allowed values are any string up to 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.

prefix

Syntax

[no] **prefix** [*ipv-prefix/prefix-length*] [**exact** | **longer** | **through** *length* | **prefix-length-range** *length1-length2*]
no prefix [*ipv-prefix/prefix-length*] [**exact** | **longer** | **through** *length* | **prefix-length-range** *length1-length2*]

Context

config>router>policy-options>prefix-list

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a prefix entry in the route policy prefix list.
The **no** form of this command deletes the prefix entry from the prefix list.

Parameters

ip-prefix

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
<exact longer thro*>		: keyword
<length>		: [0 to 128] (prefix-length <= length)
<length1-length2>		: length1/length - [0 to 128] (prefix-length <= length1 <=length2)

exact

Specifies the prefix list entry only matches the route with the specified *ip-prefix* and prefix *mask* (length) values.

longer

Specifies that the prefix list entry matches any route that matches the specified *ip-prefix* and prefix *mask* length values greater than the specified *mask*.

through *length*

Specifies that the prefix list entry matches any route that matches the specified ip-prefix and has a prefix length between the specified *length* values inclusive.

Values 0 to 32

prefix-length-range *length1-length2*

Specifies a route must match the most significant bits and have a prefix length with the specified range. The range is inclusive of start and end values. The *length2* value is greater than the *length1* value.

Values 0 to 32

6.10.2.1.5 Route policy entry match commands

entry

Syntax

entry *entry-id*

no entry

Context

config>router>policy-options>policy-statement

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context edit route policy entries within the route policy statement.

Multiple entries can be created using unique entries. The 7210 SAS exits the filter when the first match is found and executes the action specified. For this reason, entries must be sequenced correctly from most to least explicit.

An entry does not require that matching criteria be defined (in which case, everything matches), but must have at least an action defined to be considered complete. Entries without an action are considered incomplete and will be rendered inactive.

The **no** form of this command removes the specified entry from the route policy statement.

Parameters

entry-id

Specifies the entry ID expressed as a decimal integer. An *entry-id* uniquely identifies match criteria and the corresponding action. Nokia recommends that multiple entries be specific *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 4294967295

from

Syntax

[no] from

Context

config>router>policy-options>policy-statement>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure policy match criteria based on a route source or the protocol from which the route is received.

If no condition is specified, all route sources are considered to match.

The **no** form of this command deletes the source match criteria for the route policy statement entry.

family

Syntax

family [ipv4] [vpn-ipv4]

no family

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies address families as matching conditions.

Parameters

ipv4

Specifies IPv4 routing information.

vpn-ipv4

Specifies IPv4 VPN routing information.

area

Syntax

area *area-id*

no area

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an OSPF area as a route policy match criterion.

This match criterion is only used in export policies.

All OSPF routes (internal and external) are matched using this criterion if the best path for the route is through the specified area.

The **no** form of this command removes the OSPF area match criterion.

Parameters

area-id

Specifies the OSPF area ID, expressed in dotted-decimal notation or as a 32-bit decimal integer.

Values	0.0.0.0 to 255.255.255.255 (dotted-decimal), 0 to 4294967295 (decimal)
---------------	--

aigp-metric

Syntax

aigp-metric *metric*

aigp-metric *metric* **add**

aigp-metric **igp**

no aigp-metric

Context

config>router>policy-options>policy-statement>default-action

config>router>policy-options>policy-statement>entry>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns a BGP AIGP metric to routes matching the entry. The effect of this command on a route matched and accepted by a route policy entry depends on how the policy is applied (BGP import policy versus BGP export policy), the type of route, and the specific form of this command.

In a BGP import policy, this command is used to:

- associate an AIGP metric with an iBGP route received with an empty AS path and no AIGP attribute
- associate an AIGP metric with an eBGP route received without an AIGP attribute that has an AS path containing only AS numbers belonging to the local AIGP administrative domain
- modify the received AIGP metric value prior to BGP path selection

In a BGP export policy, this command is used to:

- add the AIGP attribute and set the AIGP metric value in a BGP route originated by exporting a direct, static, or IGP route from the routing table
- remove the AIGP attribute from a route advertisement to a specific peer
- modify the AIGP metric value in a route advertisement to a specific peer

The **no** form of this command removes the AIGP attribute and any explicit AIGP metric value changes that were previously configured using this command.

Default

no aigp-metric

Parameters

add

Keyword to add the AIGP attribute.

igp

Keyword to set the AIGP metric value to the IGP metric value.

metric

Specifies the AIGP metric value.

Values 0 to 4294967295

as-path

Syntax

as-path *name*

no as-path

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an AS path regular expression statement as a match criterion for the route policy entry.

If no AS path criterion is specified, any AS path is considered to match.

AS path regular expression statements are configured at the global route policy level (**config>router>policy-options>as-path**).

The **no** form of this command removes the AS path regular expression statement as a match criterion.

Default

no as-path

Parameters

name

Specifies the AS path regular expression name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@," "start@variable@end"," @variable@end", or "start@variable@".

as-path-group

Syntax

as-path-group *name*

no as-path-group *name*

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a route policy AS path regular expression statement to use in route policy entries.

The **no** form of this command deletes the AS path regular expression statement.

Parameters

name

Specifies the AS path regular expression name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double

quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@," "start@variable@end"," @variable@end", or "start@variable@".

community

Syntax

community *name*

no community

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a community list as a match criterion for the route policy entry.

If no community list is specified, any community is considered a match.

The **no** form of this command removes the community list match criterion.

Default

no community

Parameters

name

Specifies the community list 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 *name* specified must already be defined.

external

Syntax

[no] **external**

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the external route matching criteria for the entry.

Default

no external

group-address

Syntax

group-address *prefix-list-name*

no group-address

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the multicast group address prefix list containing multicast group addresses that are embedded in the join or prune packet as a filter criterion. The prefix list must be configured before entering this command. Prefix lists are configured in the **config>router>policy-options>prefix-list** context.

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

Default

no group-address

Parameters

prefix-list-name

Specifies the prefix list name. Allowed values are any string up to 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 *prefix-list-name* is defined in the **config>router>policy-options>prefix-list** context.

host-ip

Syntax

host-ip *prefix-list-name*

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies a prefix list host IP address as a match criterion for the route policy statement entry.

Default

no host-ip

Parameters

prefix-list-name

Specifies the prefix list name. Allowed values are any string up to 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 *prefix-list-name* is defined in the **config>router>policy-options>prefix-list** context.

interface

Syntax

interface *interface-name*

no interface

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the router interface, specified either by name or address, as a filter criterion.

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

Default

no interface

Parameters

ip-int-name

Specifies the name of the interface as a match criterion for this entry. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

level

Syntax

level {1 | 2}

no level

Context

config>router>policy-options>policy-statement>entry>from

config>router>policy-options>policy-statement>entry>to

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the IS-IS route level as a match criterion for the entry.

Default

no level

Parameters

1 | 2

Keyword to match the IS-IS route learned from level 1 or level 2.

neighbor

Syntax

neighbor {*ip-address* | **prefix-list** *name*}

no neighbor

Context

config>router>policy-options>policy-statement>entry>to

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the neighbor address as found in the source address of the actual join and prune message as a filter criterion. If no neighbor is specified, any neighbor is considered a match.

The **no** form of the of the command removes the neighbor IP match criterion from the configuration.

Default

no neighbor

Parameters

ip-address

Specifies the neighbor IP address in dotted-decimal notation.

Values ipv4-address: a.b.c.d

prefix-list *name*

Specifies the prefix list 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 *name* specified must already be defined.

origin

Syntax

origin {igp | egp | incomplete | any}

no origin

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a BGP origin attribute as a match criterion for a route policy statement entry.

If no origin attribute is specified, any BGP origin attribute is considered a match.

The **no** form of this command removes the BGP origin attribute match criterion.

Default

no origin

Parameters

igp

Keyword to configure matching path information originating within the local AS.

egp

Keyword to configure matching path information originating in another AS.

incomplete

Keyword to configure matching path information learned by another method.

any

Keyword to ignore this criteria.

policy-statement

Syntax

[no] **policy-statement** *name*

Context

config>router>policy-options

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure a route policy statement.

Route policy statements enable appropriate processing of IGMP group membership reports received from hosts. The processing action taken is determined by the action associated with the entries configured in the policy statement.

The **policy-statement** is a logical grouping of match and action criteria.

The **no** form of this command deletes the policy statement.

Default

no policy-statement

Parameters

name

Specifies the route policy statement name. Allowed values are any string up to 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.

prefix-list

Syntax

prefix-list *name* [*name...*(up to 5 max)]

no prefix-list

Context

config>router>policy-options>policy-statement>entry>from

config>router>policy-options>policy-statement>entry>to

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a prefix list as a match criterion for a route policy statement entry.

If no prefix list is specified, any network prefix is considered a match.

The prefix lists specify the network prefix (this includes the prefix and length) a specific policy entry applies.

A maximum of five prefix names can be specified.

The **no** form of this command removes the prefix list match criterion.

Default

no prefix-list

Parameters

name

Specifies the prefix list name. Allowed values are any string up to 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.

protocol

Syntax

protocol *protocol* [**all** | {**instance** *instance*}]

no protocol

Context

config>router>policy-options>policy-statement>entry>from

config>router>policy-options>policy-statement>entry>to

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a routing protocol as a match criterion for a route policy statement entry. This command is used for both import and export policies depending how it is used.

If no protocol criterion is specified, any protocol is considered a match.

The **no** form of this command removes the protocol match criterion.

Default

no protocol

Parameters

protocol

Specifies the protocol name to match on.

Values direct, static, bgp, isis, ospf, aggregate, bgp-vpn, igmp, periodic

instance

Specifies the OSPF or IS-IS instance.

Values 1 to 31

all

OSPF- or IS-IS-only keyword.

source-address

Syntax

source-address *ip-address*

no source-address

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the source address that is embedded in the join or prune packet as a filter criterion.

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

This command specifies a multicast data source address as a match criterion for this entry.

Parameters

ip-address

Specifies the IP prefix for the IP match criterion in dotted-decimal notation.

ipv4-address - a.b.c.d

tag

Syntax

tag *tag*

no tag

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds an integer tag to the static route. These tags are then matched on to control route redistribution.

The **no** form of this command removes the tag field match criterion.

Default

no tag

Parameters

tag

Specifies to match a specific external LSA tag field.

Values **no-tag**, 1 to 4294967295

to

Syntax

[no] to

Context

config>router>policy-options>policy-statement>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure export policy match criteria based on a route destination or the protocol into which the route is being advertised.

If no condition is specified, all route destinations are considered to match.

The **to** command context only applies to export policies. If it is used for an import policy, match criteria is ignored.

The **no** form of this command deletes export match criteria for the route policy statement entry.

type

Syntax

type *type*

no type

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures an OSPF type metric as a match criterion in the route policy statement entry.

If no type is specified, any OSPF type is considered a match.

The **no** form of this command removes the OSPF type match criterion.

Parameters

1

Keyword to match OSPF routes with type 1 LSAs.

2

Keyword to match OSPF routes with type 2 LSAs.

6.10.2.1.6 Route policy action commands

action

Syntax

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

no action

Context

config>router>policy-options>policy-statement>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures actions to take for routes matching a route policy statement entry.

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

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

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

Default

no action

Parameters

accept

Specifies packets matching the entry match criteria will be accepted and processed appropriately.

next-entry

Specifies that the actions specified would be taken and policy evaluation would continue with the next policy entry (if any others are specified).

next-policy

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

reject

Specifies packets matching the entry match criteria will be rejected.

as-path

Syntax

as-path {add | replace} *name*

no as-path

Context

config>router>policy-options>policy-statement>default-action

config>router>policy-options>policy-statement>entry>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns a BGP AS path list to routes matching the route policy statement entry.

If no AS path list is specified, the AS path attribute is not changed.

The **no** form of this command disables the AS path list editing action from the route policy entry.

Default

no as-path

Parameters

add

Specifies that the AS path list is to be prepended to an existing AS list.

replace

Specifies AS path list replaces any existing as path attribute.

name

Specifies the AS path list name. Allowed values are any string up to 32 characters 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 *name* specified must already be defined.

as-path-prepend

Syntax

as-path-prepend *as-num* [*repeat*]

no as-path-prepend

Context

config>router>policy-options>policy-statement>default-action

config>router>policy-options>policy-statement>entry>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command prepends a BGP ASN once or multiple times to the AS path attribute of routes matching the route policy statement entry.

If an ASN is not configured, the AS path is not changed.

If the optional *number* is specified, the ASN is prepended as many times as indicated by the number.

The **no** form of this command disables the AS path prepend action from the route policy entry.

Default

no as-path-prepend

Parameters

as-num

Specifies the ASN to prepend expressed as a decimal integer.

Values 1 to 4294967295

repeat

Specifies the number of times to prepend the specified ASN expressed as a decimal integer.

Values 1 to 50

community

Syntax

community {{**add** *name* [**remove** *name*]} | {**remove** *name* [**add** *name*]} | {**replace** *name*}}

no community

Context

config>router>policy-options>policy-statement>default-action

config>router>policy-options>policy-statement>entry>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds or removes a BGP community list to or from routes matching the route policy statement entry.

If no community list is specified, the community path attribute is not changed.

The community list changes the community path attribute according to the **add** and **remove** keywords.

The **no** form of this command disables the action to edit the community path attribute for the route policy entry.

Default

no community

Parameters

add

Keyword to specify that the community list is added to any existing list of communities.

remove

Keyword to specify that the community list is removed from the existing list of communities.

replace

Keyword to specify that the community list replaces any existing community attribute.

name

Specifies the community list 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.

damping

Syntax

damping {*name* | **none**}
no damping

Context

config>router>policy-options>policy-statement>default-action
config>router>policy-options>policy-statement>entry>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a damping profile used for routes matching the route policy statement entry. If no damping criteria is specified, the default damping profile is used. The **no** form of this command removes the damping profile associated with the route policy entry.

Default

no damping

Parameters

name

Specifies the damping profile 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 *name* specified must already be defined.

none

Keyword to disable route damping for the route policy.

default-action

Syntax

default-action {**accept** | **next-entry** | **reject**}
no default-action

Context

config>router>policy-options>policy-statement

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure actions for routes packets that do not match any route policy statement entries when the **accept** parameter is specified.

The default action clause can be set to all available action states, including accept, reject, next-entry and next-policy. If the action states accept or reject, the policy evaluation terminates and the appropriate result is returned.

If a default action is defined and no matches occurred with the entries in the policy, the default action clause is used.

If a default action is defined and one or more matches occurred with the entries of the policy then the default action is not used.

The **no** form of this command deletes the **default-action** context for the policy statement.

Default

no default-action

Parameters

accept

Keyword to specify that route packets matching the entry match criteria will be accepted and propagated and processed appropriately.

next-entry

Keyword to specify that the actions specified will be made to the route attributes taken, and policy evaluation will continue with the next policy entry (if any others are specified).

reject

Keyword to specify that routes or packets matching the entry match criteria will be rejected.

local-preference

Syntax

local-preference *preference*

no local-preference

Context

config>router>policy-options>policy-statement>default-action

config>router>policy-options>policy-statement>entry

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns a BGP local preference to routes matching a route policy statement entry.

If no local preference is specified, the BGP configured local preference is used.

The **no** form of this command disables assigning a local preference in the route policy entry.

Default

no local-preference

Parameters

preference

Specifies the local preference expressed as a decimal integer.

Values 0 to 4294967295

metric

Syntax

metric {**add** | **subtract** | **set**} *metric*

no metric

Context

config>router>policy-options>policy-statement>default-action

config>router>policy-options>policy-statement>entry>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns a metric to routes matching the policy statement entry.

If no metric is specified, the configured metric is used. If neither is defined, no metric will be advertised.

The value assigned for the metric by the route policy is controlled by the required keywords.

The **no** form of this command disables assigning a metric in the route policy entry.

Default

no metric

Parameters

add

Keyword to add the specified integer to any existing metric. If the result of the addition results in a number greater than 4294967295, the value 4294967295 is used.

subtract

Keyword to subtract the specified integer from any existing metric. If the result of the subtraction results in a number less than 0, the value of 0 is used.

set

Keyword to replace any existing metric with the specified integer.

metric

Specifies the metric modifier expressed as a decimal integer.

Values 0 to 4294967295

next-hop-self

Syntax

[no] next-hop-self

Context

config>router>policy-options>policy-statement>default-action

config>router>policy-options>policy-statement>entry>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command advertises a next hop IP address belonging to this router even if a third-party next hop is available to routes matching the policy statement entry.

The **no** form of this command disables advertising the **next-hop-self** option for the route policy entry.

Default

no next-hop-self

origin

Syntax

origin {igp | egp | incomplete}

no origin

Context

config>router>policy-options>policy-statement>default-action

config>router>policy-options>policy-statement>entry>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command sets the BGP origin assigned to routes exported into BGP.

If the routes are exported into protocols other than BGP, this option is ignored.

The **no** form of this command disables setting the BGP origin for the route policy entry.

Default

no origin

Parameters

igp

Keyword to set the path information as originating within the local AS.

egp

Keyword to set the path information as originating in another AS.

incomplete

Keyword to set the path information as learned by some other means.

preference

Syntax

preference *preference*

no preference

Context

config>router>policy-options>policy-statement>default-action

config>router>policy-options>policy-statement>entry>action>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns a route preference to routes matching the route policy statement entry.

If no preference is specified, the default route table manager (RTM) preference for the protocol is used.

The **no** form of this command disables setting an RTM preference in the route policy entry.

Default

no preference

Parameters

preference

Specifies the route preference, expressed as a decimal integer.

Values 1 to 255 (0 represents unset - MIB only)

tag

Syntax

tag tag
no tag

Context

config>router>policy-options>policy-statement>default-action
config>router>policy-options>policy-statement>entry>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns an OSPF tag to routes matching the entry. The tag value is used to apply a tag to a route for either an OSPF or RIP route. A hexadecimal value of 4 octets can be entered.

For OSPF, all four octets can be used.

For RIP, only the two most significant octets are used if more than two octets are configured.

The **no** form of this command removes the tag.

Default

no tag

Parameters

tag

Specifies an OSPF or IS-IS tag assigned to routes matching the entry.

Values

Accepts decimal or hex formats:

OSPF and IS-IS:	[0x0..0xFFFFFFFF]H
RIP:	[0x0..0xFFFF]H

type

Syntax

type {*type*}
no type

Context

config>router>policy-options>policy-statement>default-action
config>router>policy-options>policy-statement>entry>action

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command assigns an OSPF type metric to routes matching the route policy statement entry and being exported into OSPF.

The **no** form of this command disables assigning an OSPF type in the route policy entry.

Default

no type

Parameters

type

Specifies the OSPF type metric.

Values	1 — Set as OSPF routes with type 1 LSAs
	2 — set as OSPF routes with type 2 LSAs

6.10.2.2 Show commands

policy

Syntax

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

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays configured policy statement information.

Parameters

- policy name**
Displays information similar to the info command for a specific policy statement. If a *name* is provided, the matching policy statement displays.
If no statement name is specified, a list of all policies statements and descriptions display.
- prefix-list name**
Displays the prefix lists configured in the route policy for the specified policy name.
- admin**
Keyword to display the entire policy option configuration, including any uncommitted configuration changes. This command is similar to the **info** command.

Output

The following outputs are examples of router policy information, and [Table 79: Output fields: router policy](#) describes the output fields.

- [Sample output - show router policy](#)
- [Sample output - show router policy admin](#)
- [Sample output - show router policy "From direct To RIP"](#)

Sample output - show router policy

The **show router policy** command displays all configured route policies.

```
A:ALA-1# show router policy
=====
Route Policies
=====
Policy          Description
-----
OSPF to OSPF    Policy Statement for 'OSPF to OSPF'
Direct And Aggregate  Policy Statement ABC
-----
Policies : 2
=====
A:ALA-1#
```

Sample output - show router policy admin

The **show router policy admin** command is similar to the **info** command which displays information about the route policies and parameters.

```
*A:7210-SAS>show>router# policy admin
  prefix-list "abc"
    prefix 10.1.1.0/24 longer
    prefix 10.1.1.1/32 exact
    prefix 10.1.0.0/16 prefix-length-range 16-24
  exit
  community "S00" members "origin:12345:1"
  community "sample" members "target:12345:10"
  as-path "null" "null"
  as-path "test" "1234"
  as-path "prevent loop" "null"
```

```
damping "re"
    reuse 100
exit
damping "max"
    max-suppress 20
exit
damping "sup"
    suppress 20000
exit
damping "half"
    half-life 10
exit
damping "test"
exit
policy-statement "abcd"
    description "Test for policy statements"
    entry 1
        from
            area 0.0.0.0
        exit
to
    protocol bgp
    exit
    action accept
    exit
exit
    entry 2
        from
            community "sample"
        exit
        to
            neighbor 10.2.2.2
        exit
        action accept
        exit
    exit
    entry 3
        from
            external
        exit
        to
            level 2
        exit
        action accept
        exit
    exit
    entry 4
        from
            family vpn-ipv4
        exit
        to
            protocol bgp-vpn
        exit
        action accept
        exit
    exit
entry 5
    from
        protocol bgp
    exit
    action accept
    next-hop 10.1.1.1
    exit
exit
```

```
        entry 6
            from
                protocol bgp
            exit
            action accept
                as-path add "null"
            exit
        exit
    entry 7
        from
            protocol bgp
        exit
        action accept
            as-path replace "sample"
        exit
    exit
    default-action accept
    exit
    exit
    policy-statement "test"
        entry 2
            from
            exit
            to
            exit
            action accept
            exit
        exit
        default-action accept
        exit
    exit
*A:7210-SAS>show>router#

*A:7210-2# show router policy admin
    prefix-list "host"
        prefix 10.0.0.0/8 longer
    exit
    prefix-list "group"
        prefix 239.6.6.6/32 exact
    exit
    policy-statement "block-igmp"
        description "Reject-Reports-From-Specific-Group-And-Host"
        entry 1
            from
                host-ip "host"
            exit
            action next-entry
            exit
        exit
        entry 2
            from
                group-address "group"
            exit
            action reject
            exit
        default-action accept
        exit
    exit
    policy-statement "permit-igmp"
        description "Accept-Reports-From-Specific-Group-And-Host"
        entry 1
            from
                host-ip "host1"
                group-address "group1"
```

```

        exit
        action accept
        exit
    exit
    default-action reject
exit

```

The **show router policy *name*** command displays information about a specific route policy.

```

*A:7210-2# show router policy permit-igmp
description "Accept-Reports-From-Specific-Group-And-Host"
entry 1
    from
        host-ip "host1"
        group-address "group1"
    exit
    action accept
    exit
exit
default-action reject
*A:7210-2#

```

The **show router policy prefix-list** command, lists the prefix-lists configured in the route policy.

```

*A:7210-2# show router policy prefix-list
=====
Prefix Lists
=====
Prefix List Name
-----
host
group
-----
Num Prefix Lists: 2
=====
*A:7210-2#

```

Sample output - show router policy "From direct To RIP"

The **show router policy *name*** command displays information about a specific route policy.

```

d*A:dut-c>config>router>policy-options>policy-statement# info detail
-----
        description "Policy From direct To rip"
        entry 2
            description "Entry 2 - From Prot. rip To rip"
            from

                no neighbor
                no prefix-list
                no as-path
                no as-path-group
                no community
                no type
                no area
                no level

```

```

no external
no host-ip
no group-address
no interface
no tag
no family
exit
to

no neighbor
no level
no prefix-list
exit

```

Table 79: Output fields: router policy

Label	Description
Policy	Displays a list of route policy names
Description	Displays the description of each route policy
Policies	The total number of policies configured
Damping	Displays the damping profile name
half-life	Displays the half-life parameter for the route damping profile
max-suppress	Displays the maximum suppression parameter configured for the route damping profile
Prefix List	Displays the prefix list name and IP address/mask and whether the prefix list entry only matches (exact) the route with the specified <i>ip-prefix</i> and prefix <i>mask</i> (length) values or values greater (longer) than the specified <i>mask</i>
AS Path Name	Displays a list of AS path names
AS Paths	Displays the total number of AS paths configured
Community Name	Displays a list of community names
Communities	Displays the total number of communities configured

7 Standards and protocol support



Note:

The information provided in this chapter is subject to change without notice and may not apply to all platforms.

Nokia assumes no responsibility for inaccuracies.

The following conventions are used in this section:

- T(A,N) indicates 7210 SAS-T in both Access-uplink mode and Network mode. Similarly, T(N) indicates 7210 SAS-T in network mode only.
- K5 indicates 7210 SAS-K 2F2T1C.
- K12 indicates 7210 SAS-K 2F4T6C.
- K30 indicates 7210 SAS-K 3SFP+ 8C.
- Sx/S-1/10GE indicates all variants of 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE platforms.
- Sx-10/100GE indicates only the variants of 7210 SAS-Sx 10/100GE. Sx-10/100GE QSFP28 indicates the 7210 SAS-Sx 10/100GE 64 SFP+ 4QSFP28 variant.
- Sx/S-1/10GE-VC indicates 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE platforms in standalone VC mode.
- R6 indicates 7210 SAS-R6.
- R12 indicates 7210 SAS-R12.
- D indicates 7210 SAS-D and 7210 SAS-D ETR. If a line item applies only to 7210 SAS-D ETR, then it is indicated as D-ETR.
- Dxp indicates 7210 SAS-Dxp-12p (2SFP+ 4F6T) 7210 SAS-Dxp-12p ETR (2SFP+ 4F6T), 7210 SAS-Dxp 16p (2SFP+ 4F10T), and 7210 SAS-Dxp-24p (2SFP+ 6F16T). If a line item applies only to a particular variant, the variant name will be called out explicitly against that item.
- This standards list is not applicable to platforms in the satellite mode of operation, as most of the features are supported on 7x50 SR platforms. For this reason, the host platforms standards compliance must be consulted for the satellite mode of operation.

7.1 BGP

draft-ietf-idr-add-paths-04, Advertisement of Multiple Paths in BGP is supported on T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-sidr-origin-validation-signaling-04, BGP Prefix Origin Validation State Extended Community is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1772, Application of the Border Gateway Protocol in the Internet is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1997, BGP Communities Attribute is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2385, Protection of BGP Sessions via the TCP MD5 Signature Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2439, BGP Route Flap Damping is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2545, Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2858, Multiprotocol Extensions for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2918, Route Refresh Capability for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3107, Carrying Label Information in BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3392, Capabilities Advertisement with BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4271, A Border Gateway Protocol 4 (BGP-4) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4360, BGP Extended Communities Attribute is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4456, BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only R6 and R12 supports RR server functionality. Rest of the platforms support only client function.

RFC 4659, BGP/MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4684, Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4760, Multiprotocol Extensions for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4798, Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4893, BGP Support for Four-octet AS Number Space is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5004, Avoid BGP Best Path Transitions from One External to Another is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5291, Outbound Route Filtering Capability for BGP-4 is supported on K12, K30, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5668, 4-Octet AS Specific BGP Extended Community is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6811, Prefix Origin Validation is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.2 Ethernet

IEEE 802.1AB, Station and Media Access Control Connectivity Discovery is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

TIA-1057, LLDP for Media endpoint devices is supported on Dxp, Sx/S-1/10GE, and Sx/S-1/10GE-VC

IEEE 802.1ad, Provider Bridges is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1ag, Connectivity Fault Management is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1ah, Provider Backbone Bridges is supported on T(N)

IEEE 802.1ax, Link Aggregation is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1D, MAC Bridges is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1p, Traffic Class Expediting is supported on D, E, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1Q, Virtual LANs is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1s, Multiple Spanning Trees is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1w, Rapid Reconfiguration of Spanning Tree is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1X, Port Based Network Access Control is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ab, 1000BASE-T is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ac, VLAN Tag is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ad, Link Aggregation is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ae, 10 Gb/s Ethernet is supported on Dxp, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ah, Ethernet in the First Mile is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ba, 40 Gb/s and 100 Gb/s Ethernet is supported on R6, R12, and Sx-10/100GE

IEEE 802.3i, Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3u, Fast Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3z, Gigabit Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3af, Power Over Ethernet (PoE) is supported on Dxp, T-ETR, Mxp-ETR, and Sx/S-1/10GE



Note:

Sx/S-1/10GE only on PoE variant and Sx-1/10GE fiber variant with two fixed copper ports. Only on Dxp-16p and Dxp-24p.

IEEE 802.3at, Power Over Ethernet (PoE+) is supported on Dxp, T-ETR, Mxp-ETR, and Sx/S-1/10GE



Note:

Sx/S-1/10GE only on PoE variant and Sx-1/10GE fiber variant with two fixed copper ports. Only on Dxp-16p and Dxp-24p.

IEEE 802.3bt, Power Over Ethernet (PoE++/HPoE) is supported on Dxp



Note:

Only on Dxp-16p and Dxp-24p.

ITU-T G.8032, Ethernet Ring Protection Switching is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ITU-T Y.1731, OAM functions and mechanisms for Ethernet based networks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.3 EVPN

draft-snr-bess-evpn-proxy-arp-nd-00, Proxy-ARP/ND function in EVPN networks is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12



Note:

Sx/S-1/10GE standalone mode only.

RFC 7432, BGP MPLS-Based Ethernet VPN is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12



Note:

Sx/S-1/10GE standalone mode only.

draft-ietf-bess-evpn-vpws-14, Virtual Private Wire Service support in Ethernet VPN is supported on Mxp

7.4 Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, Operational management of Loop Free Alternates is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

With Segment Routing.

RFC 5286, Basic Specification for IP Fast Reroute: Loop-Free Alternates is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

With Segment Routing.

RFC 7490, Remote Loop-Free Alternate (LFA) Fast Reroute (FRR) is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

With Segment Routing.

7.5 Internet Protocol (IP) — General

draft-grant-tacacs-02, The TACACS+ Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-rrp-unified-spec-02, Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6 is supported on Mxp

RFC 768, User Datagram Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 793, Transmission Control Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 854, TELNET Protocol Specifications is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 951, Bootstrap Protocol (BOOTP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1034, Domain Names - Concepts and Facilities is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1035, Domain Names - Implementation and Specification is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1350, The TFTP Protocol (revision 2) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1534, Interoperation between DHCP and BOOTP is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1542, Clarifications and Extensions for the Bootstrap Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2131, Dynamic Host Configuration Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2132, DHCP Options and BOOTP Vendor Extensions 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 support only password based authentication.

RFC 4253, The Secure Shell (SSH) Transport Layer Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4254, The Secure Shell (SSH) Connection Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4632, Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5880, Bidirectional Forwarding Detection (BFD) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 5881, Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 5883, Bidirectional Forwarding Detection (BFD) for Multihop Paths is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 6528, Defending against Sequence Number Attacks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7130, Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces is supported on T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

7.6 IP — Multicast

RFC 1112, Host Extensions for IP Multicasting is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

IGMP v1, v2, v3 is supported.

RFC 2236, Internet Group Management Protocol, Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3306, Unicast-Prefix-based IPv6 Multicast Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3376, Internet Group Management Protocol, Version 3 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3446, Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

MSDP supported only on Sx/S-1/10GE standalone.

RFC 3618, Multicast Source Discovery Protocol (MSDP) is supported on Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 4601, Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4604, Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

MLD not supported.

RFC 4607, Source-Specific Multicast for IP is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4608, Source-Specific Protocol Independent Multicast in 232/8 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4610, Anycast-RP Using Protocol Independent Multicast (PIM) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5059, Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5384, The Protocol Independent Multicast (PIM) Join Attribute Format is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6513, Multicast in MPLS/BGP IP VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6514, BGP Encodings and Procedures for Multicast in MPLS/IP VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6515, IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6625, Wildcards in Multicast VPN Auto-Discover Routes is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6826, Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 7246, Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 7385, IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

7.7 IP — Version 4

RFC 791, Internet Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 792, Internet Control Message Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 826, An Ethernet Address Resolution Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1519, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1812, Requirements for IPv4 Routers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1981, Path MTU Discovery for IP version 6 is supported on K12, K30, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2401, Security Architecture for Internet Protocol is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Supported only for OSPFv3 authentication. Not supported for services.

RFC 2460, Internet Protocol, Version 6 (IPv6) Specification is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

7.8 IP — Version 6

RFC 2464, Transmission of IPv6 Packets over Ethernet Networks is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 3021, Using 31-Bit Prefixes on IPv4 Point-to-Point Links is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3122, Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification is supported on K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3587, IPv6 Global Unicast Address Format is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4007, IPv6 Scoped Address Architecture is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4193, Unique Local IPv6 Unicast Addresses is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4291, Internet Protocol Version 6 (IPv6) Addressing Architecture is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4443, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4861, Neighbor Discovery for IP version 6 (IPv6) is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 4862, IPv6 Stateless Address Autoconfiguration (Router Only) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5095, Deprecation of Type 0 Routing Headers in IPv6 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5952, A Recommendation for IPv6 Address Text Representation is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, and T(A) for Management only.

RFC 6106, IPv6 Router Advertisement Options for DNS Configuration is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6164, Using 127-Bit IPv6 Prefixes on Inter-Router Links is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.9 IPsec

RFC 2401, Security Architecture for the Internet Protocol is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only for use with OSPFv3 authentication. Not supported for services.

RFC 2406, IP Encapsulating Security Payload (ESP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only for use with OSPFv3 authentication. Not supported for services.

7.10 IS-IS

draft-ietf-isis-mi-02, IS-IS Multi-Instance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-segment-routing-extensions-04, IS-IS Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-kaplan-isis-ext-eth-02, Extended Ethernet Frame Size Support is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ISO/IEC 10589:2002, Second Edition, Nov. 2002, Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3359, Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3719, Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3787, Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4971, Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5120, M-ISIS: Multi Topology (MT) Routing in IS-IS is supported on K12, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5130, A Policy Control Mechanism in IS-IS Using Administrative Tags is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5301, Dynamic Hostname Exchange Mechanism for IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5302, Domain-wide Prefix Distribution with Two-Level IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5303, Three-Way Handshake for IS-IS Point-to-Point Adjacencies is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5304, IS-IS Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5305, IS-IS Extensions for Traffic Engineering TE is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5306, Restart Signaling for IS-IS (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5308, Routing IPv6 with IS-IS is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5310, IS-IS Generic Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6232, Purge Originator Identification TLV for IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6233, IS-IS Registry Extension for Purges is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-mi-02, IS-IS Multi-Instance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

K12, K30 support only a single instance and can operate in multi-instance deployment as it supports the processing of TLVs for multi-instance support.

draft-ietf-isis-segment-routing-extensions-04, IS-IS Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

7.11 Management

draft-ietf-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 and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2021, Remote Network Monitoring Management Information Base Version 2 using SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2138, Remote Authentication Dial In User Service (RADIUS) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2206, RSVP Management Information Base using SMIv2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2213, Integrated Services Management Information Base using SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2571, An Architecture for Describing SNMP Management Frameworks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2573, SNMP Applications is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2575, View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP) is supported on D, E, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2578, Structure of Management Information Version 2 (SMIv2) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2579, Textual Conventions for SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2787, Definitions of Managed Objects for the Virtual Router Redundancy Protocol is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2819, Remote Network Monitoring Management Information Base is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2856, Textual Conventions for Additional High Capacity Data Types is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2863, The Interfaces Group MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2864, The Inverted Stack Table Extension to the Interfaces Group MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2933, Internet Group Management Protocol MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3014, Notification Log MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3164, The BSD syslog Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3165, Definitions of Managed Objects for the Delegation of Management Scripts is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3231, Definitions of Managed Objects for Scheduling Management Operations is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3273, Remote Network Monitoring Management Information Base for High Capacity Networks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3414, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3416, Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3417, Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3419, Textual Conventions for Transport Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3584, Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3593, Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals is supported on Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3635, Definitions of Managed Objects for the Ethernet-like Interface Types is supported on D, E, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3826, The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3877, Alarm Management Information Base (MIB) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4001, Textual Conventions for Internet Network Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4022, Management Information Base for the Transmission Control Protocol (TCP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4113, Management Information Base for the User Datagram Protocol (UDP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4220, Traffic Engineering Link Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4292, IP Forwarding Table MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4293, Management Information Base for the Internet Protocol (IP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5101, Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information is supported on Mxp, Sx/S-1/10GE, and R6

RFC 6241, Network Configuration Protocol (NETCONF) is supported on K5, K12, T(N), Mxp, Sx/S-1/10GE, R6, and R12

RFC 6242, Using the NETCONF Protocol over Secure Shell (SSH) is supported on K5, K12, T(N), Mxp, Sx/S-1/10GE, R6, and R12

7.12 MPLS — General

RFC 3031, Multiprotocol Label Switching Architecture is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3032, MPLS Label Stack Encoding is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3443, Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4182, Removing a Restriction on the use of MPLS Explicit NULL is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5332, MPLS Multicast Encapsulations is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12

7.13 MPLS — GMPLS

draft-ietf-ccamp-rsvp-te-srlg-collect-04, RSVP-TE Extensions for Collecting SRLG Information is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.14 MPLS — LDP

draft-ietf-mpls-ldp-ipv6-15, Updates to LDP for IPv6 is supported on Mxp

draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-pdutta-mpls-ldp-v2-00, LDP Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-pdutta-mpls-ldp-hello-reduce-04, Targeted LDP Hello Reduction is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3037, LDP Applicability is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5036, LDP Specification is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5443, LDP IGP Synchronization is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5561, LDP Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:

P2MP LSPs only.

7.15 MPLS — MPLS-TP

RFC 5586, MPLS Generic Associated Channel is supported on T(N), R6, and R12

RFC 5921, A Framework for MPLS in Transport Networks is supported on T(N), R6, and R12

RFC 5960, MPLS Transport Profile Data Plane Architecture is supported on T(N), R6, and R12

RFC 6370, MPLS Transport Profile (MPLS-TP) Identifiers is supported on T(N), R6, and R12

RFC 6378, MPLS Transport Profile (MPLS-TP) Linear Protection is supported on T(N), R6, and R12

RFC 6426, MPLS On-Demand Connectivity and Route Tracing is supported on T(N), R6, and R12

RFC 6428, Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile is supported on T(N), R6, and R12

RFC 6478, Pseudowire Status for Static Pseudowires is supported on T(N), R6, and R12

RFC 7213, MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing is supported on T(N), R6, and R12

7.16 MPLS — OAM

RFC 6424, Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6425, Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping is supported on T(N), Mxp, R6, and R12

7.17 MPLS — RSVP-TE

RFC 2702, Requirements for Traffic Engineering over MPLS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2747, RSVP Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2961, RSVP Refresh Overhead Reduction Extensions is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3097, RSVP Cryptographic Authentication -- Updated Message Type Value is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3209, RSVP-TE: Extensions to RSVP for LSP Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3477, Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4090, Fast Reroute Extensions to RSVP-TE for LSP Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4561, Definition of a Record Route Object (RRO) Node-Id Sub-Object is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4875, Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs) is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4950, ICMP Extensions for Multiprotocol Label Switching is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5817, Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.18 OSPF

draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1765, OSPF Database Overflow is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2328, OSPF Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3101, The OSPF Not-So-Stubby Area (NSSA) Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3509, Alternative Implementations of OSPF Area Border Routers is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3623, Graceful OSPF Restart Graceful OSPF Restart (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3630, Traffic Engineering (TE) Extensions to OSPF Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4222, Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4552, Authentication/Confidentiality for OSPFv3 is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4576, Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4577, OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4970, Extensions to OSPF for Advertising Optional Router Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5185, OSPF Multi-Area Adjacency is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5187, OSPFv3 Graceful Restart (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5243, OSPF Database Exchange Summary List Optimization is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5250, The OSPF Opaque LSA Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5340, OSPF for IPv6 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5838, Support of Address Families in OSPFv3 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6987, OSPF Stub Router Advertisement is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

7.19 Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, OAM Procedures for VPWS Interworking is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3916, Requirements for Pseudo- Wire Emulation Edge-to-Edge (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3985, Pseudo Wire Emulation Edge-to-Edge (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4385, Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4446, IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4447, Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4448, Encapsulation Methods for Transport of Ethernet over MPLS Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5659, An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6073, Segmented Pseudowire is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6310, Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6391, Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network is supported on K12, K30, Mxp, R6, and R12

RFC 6718, Pseudowire Redundancy is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6870, Pseudowire Preferential Forwarding Status bit is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7023, MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7267, Dynamic Placement of Multi-Segment Pseudowires is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.20 Quality of Service

RFC 2430, A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2598, An Expedited Forwarding PHB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3140, Per Hop Behavior Identification Codes is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3260, New Terminology and Clarifications for Diffserv is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.21 RIP

RFC 1058, Routing Information Protocol is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2082, RIP-2 MD5 Authentication is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2453, RIP Version 2 is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

7.22 Timing

GR-1244-CORE, Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

GR-253-CORE, SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

IEEE 1588-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, IEEE default profile is supported only includes the Dxp-12p ETR, Dxp-16p, Dxp-24p. Sx-10/100GE does not support IEEE default profile.

ITU-T G.781, Synchronization layer functions, issued 09/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.813, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8261, Timing and synchronization aspects in packet networks, issued 04/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8262, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8264, Distribution of timing information through packet networks, issued 10/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8265.1, Precision time protocol telecom profile for frequency synchronization, issued 10/2010 is supported on D-ETR, K5, K12, K30, T(A,N), Mxp, Sx-1/10GE, R6, and R12

ITU-T G.8275.1, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014 is supported on Dxp, K12, K30, T(A,N), Mxp, Sx-1/10GE, Sx-10/100GE, R6, and R12



Note:

For 7210 SAS-Sx 10/100GE, the support only includes the Sx 10/100GE QSFP28 variant. For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

IEC/IEEE 61850-9-3-2016, Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation is supported on Dxp-16p and Dxp-24p

IEEE C37.238-2017 - IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications is supported on Dxp-16p and Dxp-24p

RFC 5905, Network Time Protocol Version 4: Protocol and Algorithms Specification is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

7.23 VPLS

RFC 4761, Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

On 7210 platforms, only BGP-AD is supported with TLDP signalling for PW. No BGP signalling is supported for PW establishment.

RFC 4762, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5501, Requirements for Multicast Support in Virtual Private LAN Services is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6074, Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)