



7210 Service Access System

Release 26.3.R1

7210 SAS-Mxp, R6, R12, S, Sx, T Routing Protocols Guide

3HE 27846 AAAA TQZZA 01

Edition: 01

March 2026

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.

© 2026 Nokia.

Table of contents

List of tables.....	15
List of figures.....	20
1 Getting started.....	22
1.1 About this guide.....	22
1.1.1 Document structure and content.....	23
1.2 7210 SAS modes of operation.....	23
1.3 7210 SAS port modes.....	25
1.4 7210 SAS router configuration process.....	27
1.5 Conventions.....	28
1.5.1 Precautionary and information messages.....	28
1.5.2 Options or substeps in procedures and sequential workflows.....	28
2 Multicast.....	30
2.1 Overview of multicast.....	30
2.1.1 Multicast models (SSM).....	30
2.1.1.1 SSM.....	31
2.2 Multicast features.....	31
2.2.1 IGMP.....	31
2.2.1.1 IGMP versions and interoperability requirements.....	32
2.2.1.2 IGMP version transition.....	32
2.2.1.3 SSM groups.....	32
2.2.2 PIM-SM.....	33
2.2.2.1 PIM-SM functions.....	33
2.2.2.2 Encapsulating data packets in the register tunnel.....	35
2.2.2.3 PIM bootstrap router mechanism.....	35
2.2.2.4 PIM-SM routing policies.....	36
2.2.2.5 Reverse Path Forwarding checks.....	37
2.2.2.6 Distributing PIM joins over multiple ECMP paths.....	39
2.2.3 MSDP.....	42
2.2.3.1 Anycast RPs for MSDP.....	42
2.2.3.2 MSDP procedure.....	43
2.2.3.3 MSDP peer groups.....	44

2.2.3.4	MSDP mesh groups.....	44
2.2.3.5	MSDP routing policies.....	44
2.2.3.6	Draft-Rosen multicast in VPNs.....	45
2.2.4	Dynamic multicast signaling over P2MP LDP.....	45
2.2.5	Multicast debugging tools.....	45
2.2.5.1	Mtrace.....	46
2.2.5.2	Mrinfo.....	47
2.2.6	Configuration guidelines for 7210 SAS.....	47
2.3	Configuring multicast parameters with CLI.....	48
2.3.1	Multicast configuration overview.....	48
2.3.2	Basic configuration.....	49
2.3.3	Common configuration tasks.....	52
2.3.4	Configuring IGMP parameters.....	53
2.3.4.1	Enabling IGMP.....	53
2.3.4.2	Configuring an IGMP interface.....	53
2.3.4.3	Configuring static parameters.....	54
2.3.4.4	Configuring SSM translation.....	55
2.3.5	Configuring PIM parameters.....	56
2.3.5.1	Enabling PIM.....	56
2.3.5.2	Configuring PIM interface parameters.....	57
2.3.5.3	Importing PIM join or register policies.....	61
2.3.6	Configuring MSDP parameters.....	62
2.3.7	Disabling IGMP or PIM.....	63
2.3.8	Disabling MSDP.....	64
2.4	Multicast command reference.....	65
2.4.1	Command hierarchies.....	65
2.4.1.1	Configuration commands.....	65
2.4.1.2	IGMP commands.....	65
2.4.1.3	PIM commands.....	66
2.4.1.4	MSDP commands.....	67
2.4.1.5	Operational commands.....	68
2.4.1.6	Show commands.....	68
2.4.1.7	Clear commands.....	69
2.4.1.8	Debug commands.....	70
2.4.2	Command descriptions.....	71
2.4.2.1	Configuration commands.....	71

2.4.2.2	Show commands.....	128
2.4.2.3	Clear commands.....	177
2.4.2.4	Debug commands.....	186
3	RIP.....	203
3.1	RIP overview.....	203
3.1.1	RIP features.....	204
3.1.1.1	RIP version types.....	204
3.1.1.2	RIPv2 authentication.....	204
3.1.1.3	Metrics.....	204
3.1.1.4	Timers.....	205
3.1.1.5	Import and export policies.....	205
3.1.1.6	RIP packet format.....	205
3.1.2	Hierarchical levels.....	207
3.2	RIP configuration process overview.....	208
3.3	Configuration notes.....	208
3.3.1	General.....	208
3.4	Configuring RIP with CLI.....	208
3.5	RIP configuration overview.....	209
3.5.1	Preconfiguration requirements.....	209
3.5.2	RIP hierarchy.....	209
3.6	Basic RIP configuration.....	209
3.7	Common configuration tasks.....	210
3.7.1	Configuring interfaces.....	210
3.7.2	Configuring a route policy.....	211
3.7.3	Configuring RIP parameters.....	213
3.7.4	Configuring global-level parameters.....	214
3.7.5	Configuring group-level parameters.....	215
3.7.6	Configuring neighbor-level parameters.....	215
3.8	RIP configuration management tasks.....	216
3.8.1	Modifying RIP parameters.....	216
3.8.2	Deleting a group.....	217
3.8.3	Deleting a neighbor.....	218
3.9	RIP command reference.....	218
3.9.1	Command hierarchies.....	218
3.9.1.1	Configuration commands.....	218

3.9.1.2	Show RIP commands.....	220
3.9.1.3	Clear RIP commands.....	220
3.9.1.4	Debug RIP commands.....	221
3.9.2	Command descriptions.....	221
3.9.2.1	RIP configuration commands.....	221
3.9.2.2	Show commands.....	237
3.9.2.3	Clear commands.....	251
3.9.2.4	Debug RIP commands.....	253
4	OSPF.....	257
4.1	Configuring OSPF.....	257
4.1.1	OSPF areas.....	258
4.1.1.1	Backbone area.....	258
4.1.1.2	Stub area.....	259
4.1.1.3	Not-So-Stubby Area.....	259
4.1.2	OSPFv3 authentication.....	262
4.1.3	Virtual links.....	263
4.1.4	Neighbors and adjacencies.....	263
4.1.5	Link-state advertisements.....	264
4.1.6	Metrics.....	264
4.1.7	Authentication.....	264
4.1.8	Multiple OSPF instances.....	264
4.1.8.1	Route export policies for OSPF.....	265
4.1.8.2	Preventing route redistribution loops.....	265
4.1.9	IP subnets.....	265
4.1.10	Preconfiguration recommendations.....	266
4.2	IP Fast-Reroute (IP FRR) for OSPF and IS-IS prefixes.....	266
4.2.1	LFA configuration.....	266
4.2.1.1	Reducing the scope of the LFA calculation by SPF.....	267
4.2.2	ECMP considerations.....	267
4.2.3	OSPF and IS-IS support for Loop-Free Alternate calculation.....	267
4.3	Loop-Free Alternate Shortest Path First (LFA SPF) policies.....	270
4.3.1	Configuration of route next-hop policy template.....	271
4.3.1.1	Configuring affinity or admin group constraint in route next-hop policy.....	271
4.3.1.2	Configuring SRLG group constraint in route next-hop policy.....	272
4.3.1.3	Interaction of IP and MPLS admin group and SRLG.....	273

4.3.1.4	Configuring protection type and next-hop type preference in route next-hop policy template.....	273
4.3.2	Application of route next-hop policy template to an interface.....	274
4.3.3	Excluding prefixes from LFA SPF.....	274
4.3.4	Modification to LFA next-hop selection algorithm.....	275
4.4	Segment routing in Shortest Path Forwarding.....	276
4.4.1	LFA protection using segment routing backup node SID.....	276
4.4.1.1	Detailed operation of LFA protection using backup node SID.....	277
4.4.1.2	Duplicate SID handling.....	279
4.4.1.3	OSPF control plane extensions.....	280
4.5	OSPF configuration process overview.....	282
4.6	Configuration notes.....	282
4.6.1	General.....	282
4.6.1.1	OSPF defaults.....	283
4.7	Configuring OSPF with CLI.....	283
4.8	OSPF configuration guidelines.....	283
4.9	Basic OSPF configuration.....	283
4.9.1	Configuring the router ID.....	284
4.10	Configuring OSPF components.....	285
4.10.1	Configuring OSPF parameters.....	285
4.10.2	Configuring an OSPF area.....	285
4.10.3	Configuring a stub area.....	286
4.10.4	Configuring a Not-So-Stubby Area.....	286
4.10.5	Configuring a virtual link.....	287
4.10.6	Configuring an interface.....	288
4.10.7	Configuring authentication.....	289
4.10.8	Assigning a designated router.....	292
4.10.9	Configuring route summaries.....	293
4.10.10	Configuring route preferences.....	294
4.11	OSPF configuration management tasks.....	296
4.11.1	Modifying a router ID.....	296
4.11.2	Deleting a router ID.....	296
4.11.3	Modifying OSPF parameters.....	297
4.12	OSPF command reference.....	298
4.12.1	Command hierarchies.....	298
4.12.1.1	Configuration commands for OSPF.....	298

4.12.1.2	Configuration commands for OSPF3.....	300
4.12.1.3	Show commands.....	302
4.12.1.4	Clear commands.....	302
4.12.1.5	Debug commands.....	302
4.12.2	Command descriptions.....	303
4.12.2.1	Configuration commands.....	303
4.12.2.2	Show commands.....	355
4.12.2.3	Clear commands.....	414
4.12.2.4	OSPF debug commands.....	416
5	IS-IS.....	426
5.1	Configuring IS-IS.....	426
5.1.1	Routing.....	427
5.1.2	IS-IS frequently used terms.....	428
5.1.3	ISO network addressing.....	429
5.1.3.1	IS-IS PDU configuration.....	430
5.1.3.2	IS-IS operations.....	431
5.1.4	IS-IS route summarization.....	431
5.1.5	IS-IS multi-topology for IPv6.....	431
5.1.6	IS-IS administrative tags.....	431
5.1.6.1	Setting route tags.....	432
5.1.6.2	Using route tags.....	432
5.1.7	Segment routing in Shortest Path Forwarding.....	432
5.1.7.1	Segment routing operational procedures.....	434
5.1.7.2	Segment routing tunnel management.....	440
5.1.7.3	Remote LFA with segment routing.....	441
5.1.7.4	Datapath support.....	445
5.1.7.5	Control protocol changes.....	446
5.1.7.6	BGP label route resolution using segment routing tunnels.....	450
5.1.7.7	Service packet forwarding with segment routing.....	451
5.1.7.8	Mirror services.....	452
5.1.8	IGP-LDP synchronization.....	452
5.1.9	IS-IS import policy on the 7210 SAS-Mxp.....	453
5.2	IS-IS configuration process overview.....	453
5.3	Configuration notes.....	454
5.4	Configuring IS-IS with CLI.....	454

5.5	IS-IS configuration overview.....	454
5.5.1	Router levels.....	454
5.5.2	Area address attributes.....	454
5.5.3	Interface level capability.....	455
5.5.4	Route leaking.....	455
5.6	Basic IS-IS configuration.....	456
5.7	Common configuration tasks.....	457
5.8	Configuring IS-IS components.....	457
5.8.1	Enabling IS-IS.....	457
5.8.2	Modifying router-level parameters.....	458
5.8.3	Configuring ISO area addresses.....	459
5.8.4	Configuring global IS-IS parameters.....	459
5.8.5	Configuring interface parameters.....	460
5.8.5.1	Example: configuring a Level 1 area.....	461
5.8.5.2	Example: modifying a router level capability.....	463
5.9	IS-IS configuration management tasks.....	463
5.9.1	Disabling IS-IS.....	463
5.9.2	Removing IS-IS.....	464
5.9.3	Modifying global IS-IS parameters.....	464
5.9.4	Modifying IS-IS interface parameters.....	465
5.9.5	Configuring leaking.....	466
5.9.6	Redistributing external IS-IS routers.....	468
5.10	IS-IS command reference.....	469
5.10.1	Command hierarchies.....	469
5.10.1.1	Configuration commands.....	469
5.10.1.2	Global commands.....	469
5.10.1.3	Interface commands.....	471
5.10.1.4	Show commands.....	472
5.10.1.5	Clear commands.....	472
5.10.1.6	Debug commands.....	472
5.10.2	Command descriptions.....	473
5.10.2.1	IS-IS configuration commands.....	473
5.10.2.2	Show commands.....	524
5.10.2.3	Clear commands.....	552
5.10.2.4	Debug commands.....	555

6	BGP.....	562
6.1	BGP overview.....	562
6.2	BGP communication.....	562
6.2.1	Message types.....	563
6.3	Group configuration and peers.....	564
6.4	Hierarchical levels.....	565
6.5	Route reflection.....	565
6.6	Fast external failover.....	568
6.7	Sending of BGP communities.....	569
6.8	ECMP and BGP route tunnels.....	569
6.9	Next-hop resolution of BGP labeled routes to tunnels.....	569
6.9.1	VPN-IPv4 and VPN-IPv6 route resolution.....	570
6.10	Route selection criteria.....	571
6.11	Enabling best external.....	572
6.11.1	BGP decision process with best external.....	572
6.11.2	Advertisement rules with best external.....	572
6.11.3	Displaying best-external routes.....	573
6.12	BGP path attributes.....	573
6.12.1	NEXT_HOP attribute.....	574
6.12.1.1	Next-hop indirection.....	574
6.13	BGP Routing Information Base.....	575
6.13.1	LOC-RIB features.....	575
6.13.2	BGP fast reroute.....	575
6.13.2.1	Calculating backup paths.....	576
6.13.2.2	Failure detection and switchover to the backup path.....	576
6.13.3	BGP fast reroute in a VPRN.....	577
6.13.3.1	BGP fast reroute in a VPRN configuration.....	577
6.13.4	RIB-OUT features.....	577
6.13.4.1	BGP export policies.....	577
6.13.4.2	Outbound Route Filtering.....	579
6.13.4.3	RT constrained route distribution.....	579
6.13.4.4	Minimum Route Advertisement Interval.....	581
6.13.4.5	Advertise-inactive.....	581
6.13.4.6	Split-horizon.....	582
6.14	Add-paths.....	582

6.14.1	Receiving multiple paths per prefix from a BGP peer.....	582
6.14.2	Path selection with add-paths.....	584
6.14.3	BGP decision process with ADD-PATH.....	584
6.14.4	Advertising multiple paths using ADD-PATH.....	585
6.14.5	Limiting the number of paths per prefix.....	586
6.15	AIGP metric.....	586
6.16	Command interactions and dependencies.....	587
6.16.1	Changing the ASN.....	587
6.16.2	BGP advertisement.....	587
6.16.3	Changing the local ASN.....	587
6.16.4	Changing the router ID at the configuration level.....	588
6.16.5	Hold time and keep alive timer dependencies.....	588
6.16.6	Import and export route policies.....	588
6.16.7	Route damping and route policies.....	589
6.16.8	AS Override.....	589
6.17	Configuration guidelines for BGP.....	589
6.18	BGP configuration process overview.....	589
6.19	Configuration notes.....	590
6.19.1	General.....	590
6.19.1.1	BGP defaults.....	590
6.19.1.2	BGP MIB notes.....	591
6.20	Configuring BGP with CLI.....	592
6.20.1	BGP configuration overview.....	592
6.20.1.1	Preconfiguration requirements.....	592
6.20.1.2	BGP hierarchy.....	592
6.20.1.3	Internal and external BGP configurations.....	593
6.21	Basic BGP configuration.....	593
6.22	Common configuration tasks.....	595
6.22.1	Configuring a basic autonomous system.....	595
6.22.2	Creating an autonomous system.....	595
6.22.3	Configuring a router ID.....	596
6.22.4	BGP components.....	597
6.22.5	Configuring BGP.....	597
6.22.6	Configuring group attributes.....	598
6.22.7	Configuring neighbor attributes.....	598
6.22.8	Configuring AIGP.....	599

6.22.9	BGP configuration management tasks.....	601
6.22.9.1	Modifying an ASN.....	601
6.22.9.2	Modifying the BGP router ID.....	601
6.22.9.3	Modifying the router-level router ID.....	602
6.22.9.4	Deleting a neighbor.....	603
6.22.9.5	Deleting groups.....	604
6.22.9.6	Editing BGP parameters.....	604
6.23	BGP command reference.....	604
6.23.1	Command hierarchies.....	605
6.23.1.1	Configuration commands.....	605
6.23.1.2	Global BGP commands.....	605
6.23.1.3	Group BGP commands.....	607
6.23.1.4	Neighbor BGP commands.....	608
6.23.1.5	Other BGP-related commands.....	609
6.23.1.6	Show commands.....	610
6.23.1.7	Clear commands.....	610
6.23.1.8	Debug commands.....	610
6.23.2	Command descriptions.....	611
6.23.2.1	Configuration commands.....	611
6.23.2.2	Other BGP-related commands.....	664
6.23.2.3	Show commands.....	665
6.23.2.4	Clear commands.....	715
6.23.2.5	Debug commands.....	718
7	Route policies.....	725
7.1	Configuring route policies.....	725
7.1.1	Policy statements.....	725
7.1.1.1	Default action behavior.....	726
7.1.1.2	Denied IP prefixes.....	726
7.1.1.3	Controlling route flapping.....	726
7.2	Regular expressions.....	727
7.2.1	BGP and OSPF route policy support.....	731
7.2.1.1	BGP route policies.....	732
7.2.1.2	Re-advertised route policies.....	733
7.2.2	When to use route policies.....	733
7.3	Route policy configuration process overview.....	733

7.4	Configuration notes.....	734
7.4.1	General.....	734
7.5	Configuring route policies with CLI.....	734
7.6	Route policy configuration overview.....	734
7.6.1	When to create routing policies.....	735
7.6.2	Default route policy actions.....	735
7.6.3	Policy evaluation.....	736
7.6.4	Damping.....	738
7.7	Basic configurations.....	739
7.8	Configuring route policy components.....	740
7.8.1	Beginning the policy statement.....	740
7.8.2	Creating a route policy.....	740
7.8.3	Configuring a default action.....	741
7.8.4	Configuring an entry.....	741
7.8.5	Configuring damping.....	742
7.8.5.1	Configuring a prefix list.....	743
7.9	Route policy configuration management tasks.....	743
7.9.1	Editing policy statements and parameters.....	743
7.9.2	Deleting an entry.....	744
7.9.3	Deleting a policy statement.....	744
7.10	Use of route policies for IGMP filtering.....	745
7.11	Route policy command reference.....	746
7.11.1	Command hierarchies.....	746
7.11.1.1	Route policy configuration commands.....	746
7.11.1.2	Show commands.....	748
7.11.2	Command descriptions.....	748
7.11.2.1	Generic commands.....	748
7.11.2.2	Route policy options.....	750
7.11.2.3	Route policy damping commands.....	753
7.11.2.4	Route policy prefix commands.....	756
7.11.2.5	Route policy entry match commands.....	758
7.11.2.6	Route policy action commands.....	771
7.11.2.7	Show commands.....	783
8	Standards and protocol support.....	788
8.1	BGP.....	788

8.2	Certificate Management.....	790
8.3	Ethernet.....	790
8.4	EVPN.....	791
8.5	Fast Reroute.....	792
8.6	Internet Protocol (IP) — General.....	792
8.7	IP — Multicast.....	795
8.8	IP — Version 4.....	796
8.9	IP — Version 6.....	797
8.10	IPsec.....	798
8.11	IS-IS.....	799
8.12	Management.....	800
8.13	MPLS — General.....	803
8.14	MPLS — GMPLS.....	803
8.15	MPLS — LDP.....	804
8.16	MPLS — MPLS-TP.....	804
8.17	MPLS — OAM.....	805
8.18	MPLS — RSVP-TE.....	805
8.19	OSPF.....	805
8.20	Pseudowire.....	807
8.21	Quality of Service.....	807
8.22	RIP.....	808
8.23	Timing.....	808
8.24	VPLS.....	809

List of tables

Table 1: Supported modes of operation and configuration methods.....	24
Table 2: Supported port modes by mode of operation.....	26
Table 3: 7210 SAS platforms supporting port modes.....	27
Table 4: Configuration process.....	27
Table 5: Join filter policy match conditions.....	36
Table 6: Register filter policy match conditions.....	36
Table 7: Output fields: mrinfo.....	126
Table 8: Output fields: mtrace.....	128
Table 9: Output fields: IGMP group.....	130
Table 10: Output fields: IGMP SSM-translate.....	131
Table 11: Output fields: IGMP interface.....	132
Table 12: Output fields: IGMP static.....	135
Table 13: Output fields: IGMP statistics.....	136
Table 14: Output fields: IGMP status.....	138
Table 15: Output fields: PIM anycast.....	139
Table 16: Output fields: PIM CRP.....	140
Table 17: Output fields: PIM group.....	143
Table 18: Output fields: PIM interface.....	146
Table 19: Output fields: PIM neighbor.....	148
Table 20: Output fields: PIM RP.....	149
Table 21: Output fields: PIM RP-hash.....	150

Table 22: Output fields: S-PMSI.....	155
Table 23: Output fields: PIM statistics.....	157
Table 24: Output fields: PIM status.....	160
Table 25: Output fields: MSDP group.....	167
Table 26: Output fields: MSDPpeer.....	169
Table 27: Output fields: MSDP source.....	170
Table 28: Output fields: MSDP source-active.....	172
Table 29: Output fields: MSDP source-active rejected.....	174
Table 30: Output fields: MSDP statistics.....	175
Table 31: Output fields: MSDP status.....	176
Table 32: RIP support on interfaces and platforms.....	203
Table 33: Route preference defaults by route type.....	233
Table 34: Output fields: RIP database.....	238
Table 35: Output fields: group.....	240
Table 36: Output fields: RIP group detail.....	241
Table 37: Output fields: neighbor standard.....	243
Table 38: Output fields: neighbor detail.....	245
Table 39: Output fields: peer.....	247
Table 40: Output fields: statistics.....	249
Table 41: Handling of duplicate SIDs.....	279
Table 42: OSPF control plane extension fields.....	280
Table 43: OSPF control plane extension flags.....	281
Table 44: Route preference defaults by route type.....	294

Table 45: Route preference defaults by route type.....	313
Table 46: Route preference defaults by route type.....	320
Table 47: Output fields: OSPF area.....	358
Table 48: Output fields: OSPF database.....	365
Table 49: Output fields: OSPF interface.....	371
Table 50: Output fields: OSPF interface detail.....	378
Table 51: Output fields: OSPF neighbor.....	383
Table 52: Output fields: OSPF/OSPF3 neighbor detail.....	386
Table 53: Output fields: OSPF opaque-database.....	390
Table 54: Output fields: prefix SIDs.....	394
Table 55: Output fields: OSPF range.....	395
Table 56: Output fields: OSPF sham-link.....	397
Table 57: Output fields: OSPF sham-link neighbor.....	398
Table 58: Output fields: OSPF SPF.....	399
Table 59: Output fields: OSPF statistics.....	401
Table 60: Output fields: OSPF status.....	406
Table 61: Output fields: OSPF virtual-link.....	409
Table 62: Output fields: OSPF virtual-neighbor.....	413
Table 63: Datapath support.....	445
Table 64: Potential adjacency capabilities.....	455
Table 65: Default preferences.....	483
Table 66: Potential adjacency capabilities.....	496
Table 67: Route preference defaults by route type.....	511

Table 68: Output fields: IS-IS adjacency.....	527
Table 69: Output fields: IS-IS database.....	530
Table 70: Output fields: IS-IS hostname.....	533
Table 71: Output fields: IS-IS interface.....	535
Table 72: Output fields: prefix SIDs.....	538
Table 73: Output fields: IS-IS routes.....	541
Table 74: Output fields: IS-IS SPF.....	543
Table 75: Output fields: IS-IS SPF log.....	544
Table 76: Output fields: IS-IS statistics.....	546
Table 77: Output fields: IS-IS status.....	548
Table 78: Output fields: IS-IS summary address.....	550
Table 79: Output-fields: IS-IS topology.....	552
Table 80: BGP fast reroute scenarios (base router context).....	576
Table 81: BGP fast reroute scenarios (VPRN context).....	577
Table 82: 7210 SAS and IETF MIB variations.....	591
Table 83: MIB variable with SNMP.....	591
Table 84: Output fields: BGP PIC.....	668
Table 85: Output fields: BGP damping.....	674
Table 86: Output fields: BGP group.....	678
Table 87: Output fields: BGP neighbor.....	688
Table 88: Output fields: BGP neighbor received routes.....	692
Table 89: Output fields: show neighbor add-path.....	694
Table 90: Output fields: BGP next-hop.....	699

Table 91: Output fields: BGP paths.....	701
Table 92: Output fields: BGP routes.....	710
Table 93: Output fields: BGP summary.....	713
Table 94: Regular expression operators.....	728
Table 95: AS path and community regular expression examples.....	729
Table 96: Default route policy actions.....	735
Table 97: Output fields: route policy.....	787

List of figures

Figure 1: Anycast RP for PIM-SM implementation example.....	38
Figure 2: RIP packet format.....	206
Figure 3: RIPv1 format.....	207
Figure 4: RIPv2 format.....	207
Figure 5: RIP configuration and implementation flow.....	208
Figure 6: Backbone area.....	259
Figure 7: PEs connected to an MPLS-VPN super backbone.....	260
Figure 8: Sham links.....	261
Figure 9: Example topology with primary and LFA routes.....	268
Figure 10: Example topology with broadcast interfaces.....	269
Figure 11: Label stack for remote LFA in ring topology.....	277
Figure 12: Backup ABR node SID.....	278
Figure 13: OSPF configuration and implementation flow.....	282
Figure 14: Checking corresponding bit.....	308
Figure 15: IS-IS routing domain.....	427
Figure 16: Using area addresses to form adjacencies.....	430
Figure 17: Programming multiple tunnels to the same destination.....	435
Figure 18: Handling of the same prefix and SID in different IS-IS instances.....	438
Figure 19: Example topology remote LFA algorithm.....	442
Figure 20: Remote LFA next-hop in segment routing.....	444
Figure 21: IS-IS configuration and implementation flow.....	453

Figure 22: Configuring a Level 1 area.....	461
Figure 23: Configuring a Level 1/2 area.....	463
Figure 24: BGP configuration.....	564
Figure 25: Fully meshed BGP configuration.....	566
Figure 26: BGP configuration with route reflectors.....	567
Figure 27: BGP Update message with path identifier for VPN-IPv4 NLRI.....	583
Figure 28: BGP Update message with path identifier for IPv4 NLRI.....	583
Figure 29: BGP configuration and implementation flow.....	590
Figure 30: BGP route policy diagram.....	732
Figure 31: OSPF route policy diagram.....	732
Figure 32: Route policy configuration and implementation flow.....	734
Figure 33: Route policy process example.....	737
Figure 34: Next policy logic example.....	738
Figure 35: Damping example.....	739

1 Getting started

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

1.1 About this guide



Note:

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

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

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

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



Note:

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

- **network mode of operation**

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

- **standalone mode of operation**

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

- **standalone-VC mode of operation**

7210 SAS-Sx/S 1/10GE

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

1.1.1 Document structure and content

This guide uses the following structure to describe routing protocols and route policies content.



Note:

This guide generically covers Release 26.x.Rx content and may include some content that will be released in later maintenance loads. See the *7210 SAS Software Release Notes 26.x.Rx*, part number 3HE 27852 000x TQZZA, for information about features supported in each load of the Release 26.x.Rx software.

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

1.2 7210 SAS modes of operation

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



Note:

Not all CLI commands are supported on all 7210 SAS platforms in all modes of operation. Users can only configure CLI commands supported by the current operating mode of the router. See the *7210 SAS Software Release Notes 26.x.Rx*, part number 3HE 27852 000x TQZZA, and to the appropriate 7210 SAS software user guide for information about features and capabilities supported by a 7210 SAS platform when operating in a specific mode.

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

- **access-uplink**

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

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

- **network**

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

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

- **satellite**

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

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

- **standalone**

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

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

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

- **standalone-VC**

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

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

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

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

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

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

Table 1: Supported modes of operation and configuration methods

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

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

1.3 7210 SAS port modes

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

- **access port mode**

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

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

- **access-uplink port mode**

- ¹ By default, the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C boot up in the [network](#) mode of operation. These platforms also allow the use of [access-uplink port mode](#) (without explicit BOF configuration), which provides the option to use Layer 2 uplinks instead of IP/MPLS uplinks to the network core, similar to the 7210 SAS-K 2F1C2T router.
- ² See section [7210 SAS port modes](#) for information about port mode configuration
- ³ Implicitly operates in [network](#) mode when [standalone](#) mode of operation is configured
- ⁴ Supports MPLS uplinks only and implicitly operates in network mode

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

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

- **network port mode**

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

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

- **hybrid port mode**

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

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



Note:

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

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

Table 2: Supported port modes by mode of operation

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

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

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

Table 3: 7210 SAS platforms supporting port modes

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

1.4 7210 SAS router configuration process

This section provides process flow information to configure IP routing protocols. The following table lists the tasks necessary to configure multicast, RIP, OSPF, IS-IS, BGP, and route policies.

Table 4: Configuration process

Area	Task	Chapter
Protocol configuration	Configure routing protocols:	

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.

Area	Task	Chapter
	• Multicast	Multicast
	• RIP	RIP
	• 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

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).



Note:

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

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 is 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 describes the multicast requirements when a Nokia router is deployed as part of the user core 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. The following 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 needed 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 then treats the query as if there is no EXCLUDE definition present.

2.2.1.3 SSM groups

IGMPv3 allows 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 allows 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 performs a (*,G) join as normal, but can combine this with a prune for each of the sources the receiver does not want 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 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 eventually times out.

A multicast data sender starts sending data destined for a multicast group. The sender 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 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 switches to native forwarding. To do this, when the RP receives a register-encapsulated data packet from source S on group G, it initiates 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 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 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 the S DR to prevent the DR unnecessarily encapsulating the packets. At the end of phase 2, traffic flows 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 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 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 correct 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 is detected and the mapping tables are modified so that the unreachable RP is no longer used and the new tables are rapidly distributed throughout the domain.

2.2.2.4 PIM-SM routing policies

Multicast traffic can be restricted from specific 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 the:
Interface	RTR interface by name
Neighbor	The neighbors source address in the IP header
Group Address	Multicast Group address in the join/prune message
Source Address	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:

- notifies the RP that a source has active data for the group
- 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 ignores 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

Match condition	Matches
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 because of 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:

- S1 sends a multicast packet.
- 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.
- 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.
- RP1 is configured with RP2 and RP3 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.
- RP1 sends a copy of the register message from S1 DR to both RP2 and RP3. RP1 uses its own IP address as the source address for the PIM register message.
- 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.
- 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.
- 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 is sent when the next register is received. If RP2 decides not to wait, the register-stop message is sent now.
- 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.
- RP3 receives the register message from RP1 and de-encapsulates it, but, because no receivers are 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 previous steps, but it is RP3 that sends a copy of the register originated by S2 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 (if enabled in the **config>service>vprn** context) 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 output indicating ECMP hashing is disabled

```
*B:BB# show router 100 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: Distribution output of PIM joins over multiple ECMP paths

```

*A:BA# show router 100 pim group
=====
PIM Groups ipv4
=====
Group Address          Type      Spt Bit Inc Intf      No.0ifs
  Source Address          RP
-----
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 MSDP

Multicast Source Discovery Protocol (MSDP) defines a protocol to exchange information about multicast sources across multiple PIM-SM domains. MSDP-speaking routers in a PIM-SM domain have an MSDP peering relationship with the MSDP-speaking peer in another domain. The peering relationship is made up of a TCP connection in which control information is exchanged. Each domain has one or more connections to this virtual topology.

When a PIM-SM rendezvous point (RP) learns about a new multicast source within its own domain from a standard PIM register mechanism, it encapsulates the first data packet in an MSDP source-active (SA) message and sends it to all MSDP peers.

The SA message is flooded (after a reverse path forwarding (RPF) check) by each peer to its MSDP peers until the SA message reaches every MSDP router in the interconnected networks. If the receiving MSDP peer is an RP and the RP has a (*.G) entry (receiver) for the group, the RP creates a state for the source and joins to the shortest path tree for the source. The encapsulated data is de-encapsulated and forwarded down the shared tree of that RP. When the packet is received by the last hop router of the receiver, the last hop router may also join the shortest path tree to the source.

The MSDP speaker periodically sends SA messages that include all sources.

2.2.3.1 Anycast RPs for MSDP

MSDP provides a mechanism that allows RPs to share information about active sources. When RPs in remote domains learn about the active sources, they pass on that information to the local receivers and multicast data is forwarded between the domains. MSDP allows each domain to maintain an independent RP that does not rely on other domains but enables RPs to forward traffic between domains. PIM-SM is used to forward the traffic between the multicast domains.

Using PIM-SM, multicast sources and receivers register themselves with their local RP and this registration is performed by the closest multicast router to which the source and receiver are connected. The RP maintains information about the sources and receivers for a specific group. RPs in other domains do not have any knowledge about sources located in other domains.

MSDP is required to provide inter-domain multicast services using Any Source Multicast (ASM). Anycast RP for MSDP enables fast convergence when should an MSDP/PIM PR router fail by allowing receivers and sources to rendezvous at the closest RP.

2.2.3.2 MSDP procedure

When an RP in a PIM-SM domain first learns of a new sender, for example, by PIM register messages, it constructs an SA message and sends it to its MSDP peers. The SA message contains the following fields:

- source address of the data source
- group address the data source sends to
- IP address of the RP



Note:

An RP that is not a designated router on a shared network does not originate SA messages for directly-connected sources on that shared network. It only originates SA messages in response to a register message received from the designated router.

Each MSDP peer receives and forwards the message away from the RP address in a peer-RPF flooding fashion. The notion of peer-RPF flooding is with respect to forwarding SA messages. The Multicast RPF Routing Information Base (MRIB) is examined to determine which peer toward the originating RP of the SA message is selected. Such a peer is called an RPF peer.

If the MSDP peer receives the SA message from a non-RPF peer toward the originating RP, it drops the message. Otherwise, it forwards the message to all its MSDP peers (except the peer from which it received the SA message).

When an MSDP peer that is also an RP for its own domain receives a new SA message, it determines if there are any group members in the domain interested in any group described by an (S,G) entry within the SA message. That is, the RP checks for a (*,G) entry with a non-empty outgoing interface list. This implies that a system in the domain is interested in the group. In this case, the RP triggers an (S,G) join event toward the data source as if a join/prune message was received and addressed to the RP. This sets up a branch of the source-tree to this domain. Subsequent data packets arrive at the RP by this tree branch and are forwarded down the shared tree inside the domain. If leaf routers choose to join the source-tree they have the option to do so according to existing PIM-SM conventions. If an RP in a domain receives a PIM join message for a new group G, the RP must trigger an (S,G) join event for each active (S,G) for that group in its SA cache.

This procedure is called flood-and-join because if an RP in the domain is not interested in the group, the SA message can be ignored when there are no receivers or members interested in that domain; otherwise, the RP joins a distribution tree.

2.2.3.2.1 MSDP peering scenarios

The *draft-ietf-mboned-msdp-deploy-nn.txt*, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios* describes how protocols work together to provide intra- and inter-domain ASM service.

The following inter-domain peering scenarios are supported for ASM services:

- peering between PIM border routers (single-hop peering)
- peering between non-border routers (multi-hop peering)
- MSDP peering without BGP
- MSDP peering between mesh groups
- MSDP peering at a multicast exchange

The following intra-domain peering scenarios are supported for ASM services:

- peering between routers configured for both MSDP and MBGP
- MSDP peer is not BGP peer (meaning, no BGP peer)

2.2.3.3 MSDP peer groups

MSDP peer groups are typically created when multiple peers have a set of common operational parameters. Group parameters not specifically configured are inherited from the global level.

2.2.3.4 MSDP mesh groups

MSDP mesh groups are used to reduce SA flooding primarily in intra-domain configurations. When several speakers in an MSDP domain are fully meshed, they can be configured as a mesh group. The originator of the SA message forwards the message to all members of the mesh group. Because of this, forwarding the SA between non-originating members of the mesh group is not necessary.

2.2.3.5 MSDP routing policies

MSDP routing policies allow for filtering of inbound and outbound SA messages. Policies can be configured at the following levels:

- **global level**
This level applies to all peers.
- **group level**
This level applies to all peers in a peer group.
- **neighbor level**
This level applies only to specified peer.

The most specific level is used. If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If no policy is applied, source active messages are passed.

Match conditions include:

- **neighbor**
This condition matches on a neighbor address is the source address in the IP header of the source active message.
- **route filter**
This condition matches on a multicast group address embedded in the source active message.

- **source address filter**

This condition matches on a multicast source address embedded in the source active message.

2.2.3.6 Draft-Rosen multicast in VPNs

RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*, describes a method of providing a VPN service, which secures connections to the network, allowing more efficient service to remote users without compromising the security of firewalls. The Rosen draft specifies the protocols and procedures that must be implemented for a service provider to provide a unicast VPN. The draft extends that specification by describing the protocols and procedures that a service provider must implement to support multicast traffic in a VPN.

See the "Virtual Private Routed Network Service" section of the *7210 SAS-Mxp, S, Sx, T Services Guide* for more information.

2.2.4 Dynamic multicast signaling over P2MP LDP

This feature provides a flexible multicast signaling solution to connect native IP multicast source and receivers that are running PIM multicast protocol via an intermediate MPLS (P2MP LDP LSP) network. It allows each native IP multicast flow to be connected via an intermediate P2MP LSP by dynamically mapping each PIM multicast flow to a P2MP LDP LSP.

It is not required to manually configure a mapping of (S,G) to a P2MP LSP on the edge node of MPLS network. A signaling method is defined that allows dynamic mapping of PIM signaling to P2MP LDP tree setup on the leaf node of P2MP LSP and also P2MP LDP signaling to be handed back to PIM on root node of P2MP LSP. Because of dynamic mapping of multicast IP flow to P2MP LSP, provisioning and maintenance overhead is eliminated as multicast distribution services are added and removed from the network.

P2MP LDP LSP signaling is initiated from node that receives PIM JOIN from a downstream node. The **p2mp-ldp-tree-join** command must be configured on PIM outgoing interface that received PIM JOIN to enable handover of multicast tree signaling from PIM to P2MP LDP LSP.

Leaf node of P2MP LDP LSP selects the upstream-hop as the root node of LDP FEC based on route table lookup. On the root node of P2MP LDP LSP, multicast tree signaling is handed back to PIM and propagated upstream as native-IP PIM JOIN.

Only PIM-SSM is supported with this feature. A single instance of P2MP LDP LSP is supported between the root and leaf nodes per multicast flow, that is, no stitching of dynamic trees.

If multiple criteria exist to set up a multicast flow, the following priorities are:

1. Multicast (statically provisioned) over P2MP LSP (RSVP-TE or LDP)
2. Dynamic multicast signaling over P2MP LDP
3. PIM native-IP multicast

2.2.5 Multicast debugging tools

This section describes multicast debugging tools for the 7210 SAS.

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

2.2.5.1 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 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.

2.2.5.1.1 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 wanted interface for the path from the source. In such cases, the wanted interface should be specified explicitly as the receiver.

2.2.5.1.2 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.

For each attempt, if no response is received within the timeout, a "***" is printed. After the specified number of attempts have failed, mtrace tries to query the next hop router with a DVMRP_ASK_NEIGHBORS2 request (as used by the mrinfo program) to determine the router type.

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.

2.2.5.2 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.

2.2.6 Configuration guidelines for 7210 SAS

The following are the configuration guidelines for 7210 SAS:

- On 7210 SAS devices, on ingress of a port, multicast traffic can be processed in the context of either **igmp-snooping** (Layer 2 Ethernet multicast) or **I3-multicast** (Layer 3 multicast), but not both. It is not possible to configure SAPs on the port such that one SAP is a receiver for multicast

traffic that is processed by IGMP snooping, and the other is the receiver for multicast traffic that is processed by IP multicast in the context of Layer 3 service or RVPLS. An option per port (using the **configure>port>ethernet>multicast-ingress {l2-mc | ip-mc}** command) is available to enable one or the other. By default, backward compatibility is enabled for IGMP snooping. To allow processing of received multicast traffic as IP multicast in the context of Layer 3 service or RVPLS, the user must explicitly change the default by using the **configure>port>ethernet>multicast-ingress** command. For more information about the command, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide*.

- On 7210 SAS devices in network mode, on egress of a port, users have an option to configure Layer 2 or Layer 3 multicast replication. That is, with RVPLS IGMPv3 snooping-based multicast, a port on which receivers are present can be configured to do either Layer 2 multicast replication—where IP TTL is not decremented and the source MAC address is not replaced with the 7210 SAS chassis MAC or IP interface MAC address—or Layer 3 multicast replication—where IP TTL is decremented and the source MAC address is replaced. Users have an option to modify this with the **configure>port>ethernet>multicast-egress** command. For more information about the command, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide*.
- If a VPLS SAP is configured on the port that has IP multicast enabled, then multicast traffic received on the SAP is dropped. Unicast, Broadcast and unknown-unicast packets received on the SAP are forwarded appropriately. This behavior is true only for VPLS SAPs and does not apply to VPLS SDPs, Epipe SAPs, and Epipe SDPs.
- 7210 SAS platforms can be used as RPs.
- Static RP configuration and automatic RP discovery using PIM BSR messages is supported. 7210 SAS-R6 or 7210 SAS-R12 platforms can be configured as candidate RP (or candidate BSR).
- 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 share tree. However, the traffic rate threshold cannot be configured to trigger the switch over.
- There is no hardware support for LAG hashing of replicated multicast traffic on 7210 SAS-R6 and 7210 SAS-R12 for IES access interfaces configured on a LAG. On these platforms, to egress on IES access interfaces, software assigns multicast traffic to the member ports of the LAG based on system-defined hash logic. LAG hashing is supported by hardware for multicast traffic on 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and on network port interfaces for 7210 SAS-R6 and 7210 SAS-R12.
- RFP check 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, PIM, and MSDP.

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 wants 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)

- For MSDP:
 - enable MSDP (required)
 - configure peer
 - configure local address

Example: Enabled IGMP, PIM, and MSDP configuration output

```
A:LAX>config>router>igmp# info
-----
interface "lax-vls"
exit
interface "p1-ix"
exit
-----
A:7210SAS>config>router>igmp# info detail
-----
        interface "lax-vls"
            no import
            version 3
            no shutdown
        exit
        interface "p1-ix"
            no import
            version 3
            no shutdown
        exit
        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 "p1-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#
*A:Dut-B>config>router>msdp# info
-----
peer 10.20.1.6
    local-address 10.20.1.1
exit
group "msdpgroup"
    peer 10.20.1.4
        local-address 10.20.1.1
    exit
```

```
        peer 10.20.1.5
          local-address 10.20.1.1
        exit
      exit
-----
*A:Dut-B>config>router>msdp# info detail
-----
      no active-source-limit
      no receive-msdp-msg-rate
      data-encapsulation
      no export
      no import
      no local-address
      no rpf-table
      no sa-timeout
      no shutdown
      peer 10.20.1.6
        no active-source-limit
        no receive-msdp-msg-rate
        no authentication-key
        no default-peer
        no export
        no import
        local-address 10.20.1.1
        no shutdown
      exit
      group "msdpgroup"
        no active-source-limit
        no receive-msdp-msg-rate
        no export
        no import
        no local-address
        mode standard
        no shutdown
        peer 10.20.1.4
          no active-source-limit
          no receive-msdp-msg-rate
          no authentication-key
          no default-peer
          no export
          no import
          local-address 10.20.1.1
          no shutdown
        exit
      peer 10.20.1.5
        no active-source-limit
        no receive-msdp-msg-rate
        no authentication-key
        no default-peer
        no export
        no import
        local-address 10.20.1.1
        no shutdown
      exit
    exit
-----
```

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: Detailed output for an enabled IGMP

```
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
```

Use the following syntax to configure IGMP interfaces.

```
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

The following shows the command usage 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
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

The following shows the command usage 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 syntax to configure IGMP parameters.

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

Example

The following shows the command usage to configure IGMP parameters.

```
config>router# igmp
config>router>igmp# ssm-translate
config>router>igmp>ssm# grp-range 239.255.0.1 239.2.2.2
config>router>igmp>ssm>grp-range# source 10.1.1.1
```

Example: SSM translation 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
  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 may 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

Example: Command usage to configure PIM interface parameters

```
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: PIM configuration output

```
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

```

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#

```

```

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
        shutdown
bsr-candidate
    exit
    rp-candidate
        shutdown
    exit
exit
-----
A:SJC>config>router>pim#

```

Example

```

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#
```

```
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

```
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#

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

```
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#
```

```
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]
```

Example

The following 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 Configuring MSDP parameters

Use the following commands to configure basic MSDP parameters.

```
config>router# msdp
    peer ip-address
        active-source-limit number
        authentication-key [authentication-key | hash-key] [hash | hash2 | custom]
        default-peer
        export policy-name [policy-name]
        import policy-name [policy-name]
        local-address ip-address
        receive-msdp-msg-rate number interval seconds [threshold threshold]
        no shutdown
    no shutdown
```

Use the following CLI syntax to configure MSDP parameters.

```
config>router>msdp# peer 10.20.1.1
config>router>msdp>peer# local-address 10.20.1.6
config>router>msdp>peer# no shutdown
config>router>msdp>peer# exit
config>router>msdp# no shutdown
config>router>msdp#
```

Example: Displaying the MSDP configuration

```
ALA-48>config>router>msdp# info
-----
    peer 10.20.1.1
        local-address 10.20.1.6
    exit
-----
ALA-48>config>router>msdp#
```

2.3.7 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 239.255.0.1 239.2.2.2  
            source 10.1.1.1  
        exit  
    exit  
    interface "lax-sjc"  
        static  
            group 239.1.1.1  
            starg  
        exit  
    exit  
    interface "lax-vls"  
        static  
            group 239.255.0.2  
            source 172.22.184.197  
        exit  
    exit  
    interface "pl-ix"  
    exit  
exit  
#-----  
echo "PIM Configuration"  
#-----  
    pim  
        shutdown  
        import join-policy "foo"  
        interface "system"
```

```

        exit
        interface "lax-sjc"
        exit
        interface "lax-vls"
        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
        exit
        rp-candidate
            shutdown
        bsr-candidate
            exit
        exit
    exit
#-----
....
#-----
A:LAX>config>router#

```

2.3.8 Disabling MSDP

Use the following CLI syntax to disable MSDP.

```

config>router#
    msdp
        shutdown

```

Example: Command usage to disable multicast

```

config>router#
config>router>msdp# shutdown
config>router>msdp# exit

```

Example: Configuration output

```

A:LAX>config>router# info
-----
...
#-----
echo "MSDP Configuration"
#-----
    msdp
        shutdown
        peer 10.20.1.1
            local-address 10.20.1.6
        exit
        group "test"
            active-source-limit 50000
            receive-msdp-msg-rate 100 interval 300 threshold 5000
            export "LDP-export"

```



```
import "LDP-import"  
local-address 10.10.10.103  
mode mesh-group  
peer 10.10.10.104  
exit  
exit  
exit  
#-----
```

2.4 Multicast command reference

2.4.1 Command hierarchies

- [Configuration commands](#)
- [IGMP commands](#)
- [PIM commands](#)
- [MSDP 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 [1..1000]  
- no max-sources  
- [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]
- 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
- 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 MSDP commands

```

config
- router
  - [no] msdp
    - [no] active-source-limit number
    - [no] data-encapsulation
    - export policy-name [policy-name...(up to 5 max)]
    - no export
    - [no] group group-name
      - active-source-limit number
      - no active-source-limit
      - export policy-name [policy-name...(up to 5 max)]
      - no export
      - import policy-name [policy-name...(up to 5 max)]
      - no import
      - local-address ip-address
      - no local-address
      - mode {mesh-group | standard}
    - [no] peer peer-address
      - active-source-limit number
      - no active-source-limit
      - authentication-key {authentication-key | hash-key} [hash | hash2]
      - no authentication-key
      - [no] default-peer
      - export policy-name [policy-name...(up to 5 max)]
      - no export
      - import policy-name [policy-name...(up to 5 max)]
      - no import
      - local-address ip-address

```

```

- no local-address
- receive-msdp-msg-rate number interval seconds [threshold threshold]
- no receive-msdp-msg-rate
- [no] shutdown
- receive-msdp-msg-rate number interval seconds [threshold threshold]
- no receive-msdp-msg-rate
- [no] shutdown
- import policy-name [policy-name...(up to 5 max)]
- no import
- local-address ip-address
- no local-address
- [no] peer ip-address
- active-source-limit number
- no active-source-limit
- authentication-key {authentication-key | hash-key} [hash | hash2]
- no authentication-key
- [no] default-peer
- export policy-name [policy-name]
- no export
- import policy-name [policy-name]
- no import
- local-address ip-address
- no local-address
- receive-msdp-msg-rate number interval seconds [threshold threshold]
- no receive-msdp-msg-rate
- [no] shutdown
- receive-msdp-msg-rate number interval seconds [threshold threshold]
- no receive-msdp-msg-rate
- rpf-table {rtable-m | rtable-u | both}
- no rpf-table
- sa-timeout seconds
- no sa-timeout
- [no] shutdown
- [no] source prefix/mask
- active-source-limit number
- no active-source-limit

```

2.4.1.5 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]

```

2.4.1.6 Show commands

```

show
- router
- igmp
- group [grp-ip-address]
- group summary
- 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
- s-pmsi [group-ip group-ip]source-ip source- ip] [detail]
- s-pmsi [mdSrcAddr [mdGrpAddr]] [group-ip group-ip] [source-ip source- ip]
[detail]
- s-pmsi ext-tunnel-id ext-tunnel-id [tunnel-id tunnel-id [group-ip group-ip]
[source-ip source- ip] [detail]
- s-pmsi root-addr root-addr [lsp-idlsp-id] [group-ip group-ip] [source-ip source-
ip] [detail]
- statistics [ip-int-name | int-ip-address | mpls-if-name] [family]
- status [detail] [family]
- tunnel-interface [ip-int-name | mt-int-name | int-ip-address] [group [group-ip-
address] source ip-address] [type {starstarrp | starg | sg}] [detail] [family]
show
- router
- msdp
- group [group-name] [detail]
- peer [ip-address] [group group-name] [detail]
- source [ip-address/mask] [type {configured | dynamic | both}] [detail]
- source-active [{group ip-address | local | originator ip-address | peer ip-
address | source ip-address | group ip-address source ip-address}] [detail]
- source-active-rejected [peer-group name] [group ip-address] [source ip-address]
[originator ip-address] [peer ip-address]
- statistics [peer ip-address]
- status
show
- router {router-instance}
- mvpn
- mvpn-list
show
- router
- tunnel-table [ip-address [/mask]] [protocol | sdp sdp-id]
- tunnel-table [summary]

```

2.4.1.7 Clear commands

```

clear
- router
- igmp
- database [interface ip-int-name|ip-address] group grp-ip-address [source src-ip-
address]
- database [group grp-ip-address [source src-ip-address]]
- statistics [interface ip-int-name | ip-address]
- 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]
- s-pmsi [mdSrcAddr] [mdGrpAddr] [vprnSrcAddr vprnGrpAddr]

```

```

- statistics [{[interface ip-int-name | int-ip-address]} {[group grp-ip-address
[source ip-address]]}][family]]
- msdp
- cache [peer ip-address] [group ip-address] [source ip-address] [originrp ip-
address]
- statistics [peer ip-address]
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]

```

2.4.1.8 Debug commands

```

debug
- router
- igmp
- [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
- dynmlldp [detail]
- no dynmlldp
- 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
- mvpn-rtcache [group grp-ip-address] [peer ip-address]
- no mvpn-rtcache
- 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
- s-pmsi [{vpnSrcAddr [vpnGrpAddr]} [mdSrcAddr]] [detail]
- no s-pmsi
- tunnel-interface [ldp-p2mp p2mp-id] [sender ip-address] [detail]
- no tunnel-interface [ldp-p2mp p2mp-id] [sender ip-address]
- tunnel-interface [rsvp-p2mp lsp-name] [sender ip-address] [detail]
- no tunnel-interface [rsvp-p2mp lsp-name] [sender ip-address]

debug
- router
  - [no] msdp
  - packet [pkt-type] [peer ip-address]
  - pim [grp-address]
  - no pim
  - rtm [rp-address]
  - no rtm
  - sa-db [group grpAddr] [source srcAddr] [rp rpAddr]
  - no sa-db
```

2.4.2 Command descriptions

- [Configuration commands](#)
- [Show commands](#)
- [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.

The **shutdown** command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

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

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

Default

```
no shutdown:    config>router>igmp
                config>router>igmp>interface ip-int-name
                config>router>pim
```

Special Cases

IGMP Protocol Handling

On all 7210 SAS platforms, IGMP is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure router igmp** command instantiates the protocol in the **no shutdown** state, and resources are allocated to enable the node to process the protocol.

To deallocate resources, you must issue the **configure router igmp shutdown** and **configure router no igmp** commands to allow the node to boot up correctly after the reboot. It is not sufficient to issue only a **configure router igmp shutdown** command.

The resources for IGMP are allocated when the IGMP context is enabled either in the base routing instance or the VPRN service instance. Resources are deallocated when the configuration of the last IGMP context under either base routing instances or VPRN service is removed or shut down.

PIM Protocol Handling

On all 7210 SAS platforms, PIM is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure router pim** command instantiates the protocol in the **no shutdown** state, and resources are allocated to enable the node to process the protocol.

To deallocate resources, you must issue the **configure router pim shutdown** and **configure router no pim** commands to allow the node to boot up correctly after the reboot. It is not sufficient to issue only a **configure router pim shutdown** command.

The resources for PIM are allocated when the PIM context is enabled either in the base routing instance or the VPRN service instance. Resources are deallocated when the

configuration of the last PIM context under either base routing instances or VPRN service is removed or shut down.

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 are processed.

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

Default

no mc-maximum-routes

Parameters

number

Specifies the maximum number of routes 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 threshold

Specifies the percentage at which a warning log message and SNMP trap should be 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, the IGMP protocol 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 itself 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

This command configures an IGMP interface. The interface is a local identifier of the network interface on which reception of the specified multicast address is 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 **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 is 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 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>policy-options** 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 configures the maximum number of groups for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups are not allowed. This command is applicable for IPv4 and IPv6.

Default

max-group 0

Parameters

value

Specifies the maximum number of groups for this interface.

Values 1 to 1024

max-sources

Syntax

```
max-sources [value]  
no max-sources
```

Context

```
config>router>igmp>group-interface
```

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 for this interface.

Values 1 to 1000

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.

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 the source IPv4 address (S) for the static IGMP group being configured. Multicast traffic to the group (G) is forwarded out the interface on which this static group is configured if the source address in the IPv4 header of the multicast packets matches S.

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

The **source** command, in combination with the **group** command, 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.

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 be enabled only if no existing source addresses for this group are specified.

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

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. Routers that run different versions of IGMP 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

Values

>= 1000

query-interval

Syntax

- query-interval seconds
- no query-interval

Context

config>router>igmp

Platforms

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

Description

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

Default

query-interval 125

Parameters

- seconds

Specifies the 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

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 number of seconds 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 that are translated to SSM (S,G) entries. If the static entry needs to be created, it must be translated from an IGMPv1 IGMPv2 request to a Source Specific Multicast (SSM) join. An SSM translate source can only be added if the starg command is not enabled. An error message is generated if you try to configure the **source** command with 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 configures group ranges that are translated to SSM (S,G) entries.

Parameters

start

Specifies an IP address for the start of the group range.

end

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

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 configures 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.

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). By default, this command is not enabled.

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 configures a logical IP routing interface.

Interface names are case-sensitive and must be unique within the group of IP interfaces defined for the **config router interface** and **config service ies interface** commands. 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. An interface name cannot be in the form of an IP address. Interface names can be any string of up to 32 alphanumeric characters. 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 refreshes of PIM Assert messages on an interface.

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

Default

no assert-period

Parameters

assert-period

Specifies the period for 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.

For information about the protocols and platforms that support BFD, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide*.

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

Default

no bfd-enable

enable-mdt-spt

Syntax

[no] enable-mdt-spt

Context

config>router>pim

Platforms

7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command enables SPT switchover for default MDT. When enabled, the PIM instance resets all MDTs and reinitiates setup.

The **no** form of this command disables SPT switchover for default MDT. When 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... (up to 5 max)]]

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 prevent 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.

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 enabled, each multicast stream that needs to be forwarded over an ECMP link is re-evaluated for the total multicast bandwidth utilization. Re-evaluation occurs on the specified ECMP interface.

The **no** form of this command disables the 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 both IPv4 and IPv6.

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 never times out 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 advantage is that it eliminates duplicate traffic forwarding to the LAN. It also improves performance by removing the required interaction between the control and data planes.



Note:

This command is still fully interoperable with the implementations described in 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)*. However, there may be conformance tests that 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

improved-assert

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 the 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 in an 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

multicast-senders auto

Parameters

auto

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

always

Keyword that treats all traffic received on the interface as coming from a directly-attached multicast source.

never

Keyword to specify 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 or point-to-point interfaces, it means that 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 configures the priority value to elect the designated router (DR). The DR election priority is a 32-bit unsigned number, and the numerically larger priority is always preferred.

The **no** form of this command reverts to the default 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 **sticky-dr** operation on this interface. When 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 this interface, it continues to act as the DR for the LAN even after the old DR comes back up.

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

Default

no sticky-dr

Parameters

priority dr-priority

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

no three-way-hello

Context

config>router>pim>interface

Platforms

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

Description

This command enables three-way hello. By default, the **three-way-hello** command is disabled on all interfaces and the standard two-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 the capability of the router to enable 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 not enabled, 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 configured RP. 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 Any valid loopback address configured on the node.

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 *ip-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.



Caution:

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 Any valid IP address within the scope described previously.

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 the PIM configuration. The policies control the flow of bootstrap messages from the RP. 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...*(up to 5 max)]

no bootstrap-import

Context

config>router>pim>rp

Platforms

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

Description

This command to applies import policies to the PIM configuration. The policies control the flow of bootstrap messages to the RP. 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.

Default

bsr-candidate shutdown

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 is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values a.b.c.d

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 configures the length of a mask that is to be combined with the group address before the hash function is called. All groups with the same hash result 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 value of 0 means the router is not eligible to be the bootstrap router. A value of 1 means 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 candidate RP parameters.

Routers use a set of available rendezvous points distributed in bootstrap messages to get the correct group-to-RP mapping. A set of routers within a domain are also configured as candidate RPs (C-RPs); typically, these are 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 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 (224.0.0.0 to 239.255.255.255).

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 (255.255.0.0).

Values 4 to 32

netmask

Specifies the subnet mask in dotted decimal notation (0.0.0.0 to 255.255.255.255).

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 an RP. This value is used to elect RP for a group range.

Default

priority 192

Parameters

priority
Specifies the priority to become an RP. A value of 0 is considered 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 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 configures the RP address used by the router for the range of multicast groups configured by the range command.

Parameters

ip-address

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

Values a.b.c.d

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 to be used by the router as the RP. The **config>router>pim>rp>static>address a.b.c.d** 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 group-prefix from the configuration.

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 (255.255.0.0).

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 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) 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, a **non-dr-attract-traffic** flag can be used in the PIM context to have the router ignore the DR state and attract traffic when it is not the DR. When using this flag, the router may not send the stream down to the DSLAM while it is not the DR.

When enabled, the designated router state is ignored. When disabled, the designated router value is honored.

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 RPF interface toward the source/rendezvous point. However, the user can specify use of the unicast route table (rtable-u).

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

Default

rpf-table rtable-u

Parameters

rtable-u

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

spt-switchover-threshold

Syntax

spt-switchover-threshold {*grp-ip-prefix/ip-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 switches 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 7210 SAS, this command enables or disables switchover to the SPT. To disable switchover to SPT, a threshold value of **infinity** must be configured (that is, to continue using the shared tree for ever, configure the IP multicast prefix with this command and set the threshold to **infinity**). To use the SPT, do not configure the IP multicast address prefix using this command, and the default behavior applies 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-ip-address

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

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

ip-prefix-length

Specifies the mask associated with the IP prefix expressed as a mask length or in dotted decimal notation; for example /16 for a sixteen-bit mask. The mask can also be entered in dotted decimal notation (255.255.0.0).

Values 4 to 32

spt-threshold

Specifies the 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 threshold. When the **infinity** keyword is specified, no switchover occurs at any time, regardless of the traffic level detected

Values 1 to 4294967294 | infinity (threshold in kbps)

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 enabled, this command treats packets as if the SPT bit is set, regardless of whether it is set.

Default

ssm-assert-compatible-mode disable

ssm-default-range-disable

Syntax

ssm-default-range-disable ipv4

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 enables and disables 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

Parameters

ipv4

Keyword to specify IPv4 as the SSM default range.

ipv6

Keyword to specify IPv6 as the SSM default range.

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 an SSM group instance.

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 address ranges of the multicast groups for this router. When there are parameters present, the command configures the SSM group ranges for IPv6 addresses and netmasks.

Parameters

ip-prefix/mask

Specifies the IP prefix for the range used by the ABR to advertise that summarizes the area into another area.

Values

- ipv4-prefix: a.b.c.d
- ipv4-prefix-le: 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-le: 0 to 128

Values

Specifies the to 32 (mask length), 0.0.0.0 to 255.255.255.255 (dotted decimal)

netmask

Specifies the subnet mask in dotted decimal notation.

Values

0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

2.4.2.1.5 MSDP commands

```
msdp
```

Syntax

[no] msdp

Context

config>router

Platforms

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

Description

This command enables a Multicast Source Discovery Protocol (MSDP) instance. When an MSDP instance is created, the protocol is enabled. To start or suspend execution of the MSDP without affecting the configuration, use the **[no] shutdown** command.

For the MSDP to function, at least one peer must be configured.

When MSDP is configured and started, an appropriate event message is generated.

When the **no** form of this command is executed, all sessions are terminated and an appropriate event message is generated.

When all peering sessions are terminated, an event message per peer is not generated.

The **no** form of this command deletes the MSDP instance, removing all associated configuration parameters.

Default

no msdp

active-source-limit

Syntax

active-source-limit *number*

no active-source-limit

Context

config>router>msdp

config>router>msdp>group

config>router>msdp>group>peer

config>router>msdp>peer

config>router>msdp>source

Platforms

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

Description

This command configures the maximum number of active source messages that are accepted by MSDP, effectively controlling the number of active sources that can be stored on the system.

The **no** form of this command removes the user-configured limit on the number of source-active (SA) records.

Default

no active-source-limit

Parameters

number

Specifies the number of active sources that can be maintained by MSDP stored on the system.

Values 0 to 1000000

data-encapsulation

Syntax

[no] data-encapsulation

Context

config>router>msdp

Platforms

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

Description

This command configures a rendezvous point (RP) using MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP SA messages.

The **no** form of this command disables the encapsulation of multicast data in SA register messages. The system sends only multicast (S,G) sender information to the remote MSDP peer.

Default

data-encapsulation

export

Syntax

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

no export

Context

config>router>msdp

config>router>msdp>peer

config>router>msdp>group

config>router>msdp>group>peer

Platforms

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

Description

This command specifies the policies to export the SA state from the SA list into MSDP.

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. The user can specify a maximum of five policy names.

The **no** form of this command applies no export policies and all SA entries are announced.

Default

no export

Parameters

policy-name

Specifies the export policy name, up to 32 characters. Up to five policy-name arguments can be specified.

If you configure an export policy at the global level, each individual peer inherits the global policy. If you configure an export policy at the group level, each individual peer in a group inherits the group policy. If you configure an export policy at the peer level, the policy only applies to the peer where it is configured.

group

Syntax

[no] group *group-name*

Context

config>router>msdp

Platforms

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

Description

This command creates or modifies an MSDP group. To configure multiple MSDP groups, include multiple group statements.

By default, the group options are inherited from the global MSDP options. To override these global options, group-specific options within the group statement can be configured.

If the specified group name is already configured, this command only provides the context to configure the options pertaining to this group.

If the group name provided is not already configured, the group name must be created and the commands in the **config>router>msdp>group** context, which configure parameters for the group, become available.

For a group to be of use, at least one peer must be configured.

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

Default

no group

Parameters

group-name

Specifies a MSDP group name, up to 32 characters.

import

Syntax

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

no import

Context

config>router>msdp

config>router>msdp>peer

config>router>msdp>group

config>router>msdp>group>peer

Platforms

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

Description

This command specifies the policies to import the SA state from MSDP into the SA list.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple import commands are issued, the last command entered overrides the previous command. A maximum of five policy names can be specified.

If an import policy is configured at the global level, each individual peer inherits the global policy.

If an import policy is configured at the group level, each individual peer in a group inherits the group policy.

If an import policy is configured at the peer level, the policy only applies to the peer where it is configured.

The **no** form of this command applies no import policies and all source active messages are allowed.

Default

no import

Parameters

policy-name

Specifies the import policy name, up to 32 characters. Up to five policy-name arguments can be specified.

local-address

Syntax

local-address *ip-address*

no local-address

Context

config>router>msdp

config>router>msdp>peer

config>router>msdp>group

config>router>msdp>group>peer

Platforms

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

Description

This command configures the local end of an MSDP session. For MSDP to function, at least one peer must be configured. When configuring a peer, you must include this **local-address** command to configure the local end of the MSDP session. This address must be present on the node and is used to validate incoming connections to the peer and to establish connections to the remote peer.

If the user enters this command, the specified address is validated and used as the local address for MSDP peers from that point. If a subsequent **local-address** command is entered, it replaces the existing configuration and the existing sessions are terminated.

Similarly, when the **no** form of this command is entered, the existing **local-address** is removed from the configuration and the existing sessions are terminated.

When a session is terminated, all information pertaining to and learned from that peer is removed.

When a new peering session is created or a peering session is lost, an event message should be generated.

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

Default

no local-address

Parameters

ip-address

Specifies an existing address on the node.

Values a.b.c.d

peer

Syntax

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

Context

config>router>msdp

config>router>msdp>group

Platforms

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

Description

This command configures peer parameters. MSDP must have at least one peer configured. A peer is defined by configuring a local address that can be used by the node to set up a peering session and the address of a remote MSDP router. It is the address of this remote peer that is configured in this command and it identifies the remote MSDP router address.

After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. Multiple peering sessions may be required, in which case multiple peer statements should be included in the configurations.

By default, the options applied to a peer are inherited from the global or group level. To override these inherited options, include peer-specific options within the peer statement.

If the peer address provided is already a configured peer, this command only provides the context to configure the parameters pertaining to this peer.

If the peer address provided is not already a configured peer, the peer instance must be created and the commands in the **config>router>msdp>peer** or **config>router>msdp>group>peer** contexts, which configure parameters for the peer instance, become available.

The peer address provided is validated and, if valid, is used as the remote address for an MSDP peering session.

At least one peer must be configured for MSDP to function.

The **no** form of this command removes the existing peering address from the configuration and the existing session is terminated. When a session is terminated, all SA information pertaining to and learned from that peer is removed. When a new peering session is created or a peering session is lost, an event message should be generated.

Parameters

ip-address

Specifies the peer IP address. The address configured in this statement must identify the remote MSDP router with which the peering session must be established.

authentication-key

Syntax

authentication-key {*authentication-key* | *hash-key*} [**hash** | **hash2**]

no authentication-key

Context

config>router>msdp>peer

config>router>msdp>group>peer

Platforms

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

Description

This command configures a Message Digest 5 (MD5) authentication key to be used with a specific MSDP peering session. The authentication key must be configured per peer; therefore, no global or group configuration is possible.

The **no** form of this command configures acceptance of all MSDP messages and disables the MD5 signature option authentication key.

Default

no authentication-key

Parameters

authentication-key

Specifies the authentication key. The key can be any combination of printable, 7-bit ASCII characters, up to 255 characters in the **config>router>msdp>peer** context, or up to 127 characters in the **config>router>msdp>group>peer** context. If the string contains special characters (#, \$, spaces, and so on), enclose the entire string in quotation marks (" ").

hash-key

Specifies a hash key. The key can be any combination of ASCII characters up to 451 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, although, for security purposes, the actual unencrypted key value is not provided.

hash

Keyword to specify that the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2

Keyword to specify that the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable

cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

default-peer

Syntax

[no] default-peer

Context

config>router>msdp>peer

config>router>msdp>group>peer

Platforms

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

Description

This command enables the use of a default peer. Using the default peer mechanism, a peer can be selected as the default MSDP peer. As a result, all SA messages from the peer are accepted without the usual peer-reverse path forwarding (RPF) check.

The MSDP peer-RPF check is different from the normal multicast RPF checks. The peer-RPF check is used to stop SA messages from looping. A router validates SA messages originated from other routers in a deterministic fashion.

A set of rules is applied to validate received SA messages, and the first rule that applies determines the peer-RPF neighbor. All SA messages from other routers are rejected. The rules applied to SA messages originating at Router S received at Router R from Router N are as follows:

- If Router N and Router S are the same, the message is originated by a direct peer-RPF neighbor and is accepted.
- If Router N is a configured peer or a member of the Router R mesh group, its SA messages are accepted.
- If Router N is the Border Gateway Protocol (BGP) next hop of the active multicast RPF route toward Router S, Router N is the peer-RPF neighbor and its SA messages are accepted.
- If Router N is an external BGP peer of Router R and the last autonomous system (AS) number in the BGP AS-path to Router S is the same as Router N's AS number, Router N is the peer-RPF neighbor, and its SA messages are accepted.
- If Router N uses the same next hop as the next hop to Router S, Router N is the peer-RPF neighbor, and its SA messages are accepted.
- If Router N does not fit any of the preceding rules, Router N is not a peer-RPF neighbor, and its SA messages are rejected.

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

Default

no default-peer

receive-msdp-msg-rate

Syntax

receive-msdp-rate *number interval seconds* [**threshold** *threshold*]

no receive-msdp-rate

Context

config>router>msdp

config>router>msdp>peer

config>router>msdp>group

config>router>msdp>group>peer

Platforms

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

Description

This command limits the number of MSDP messages that are read from the TCP session. An MSDP/ RP router may receive a large number of MSDP protocol message packets in a particular SA message.

After the number of MSDP packets (including SA messages) defined in the threshold have been processed, the rate of all other MSDP packets is rate limited by no longer accepting messages from the TCP session until the time (seconds) has elapsed.

The **no** form of this command sets no limit on the number of MSDP and SA limit messages that are accepted.

Default

no receive-msdp-msg-rate

Parameters

number

Specifies the number of MSDP messages (including SA messages) that are read from the TCP session per the specified number of seconds.

Values 10 to 10000

Default 0

seconds

Specifies the time that, together with the *number* parameter, defines the number of MSDP messages (including SA messages) that are read from the TCP session within the configured number of seconds.

Values 1 to 600

Default 0

threshold

Specifies the number of MSDP messages that can be processed before the MSDP message rate limiting function is activated; this is particularly useful during system startup and initialization.

Values 1 to 1000000

Default 0

shutdown

Syntax

[no] shutdown

Context

config>router>msdp

config>router>msdp>peer

config>router>msdp>group

Platforms

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

Description

This command administratively enables or disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command and must be shut down before they may be deleted.

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

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

Default

no shutdown

mode

Syntax

mode {mesh-group | standard}

Context

config>router>msdp>group

Platforms

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

Description

This command configures groups of peers in a full mesh topology to limit excessive flooding of SA messages to neighboring peers.

MSDP peers can be configured and grouped in a full-mesh topology that prevents excessive flooding of SA messages to neighboring peers.

In a meshed configuration, all members of the group must have a peer connection with every other mesh group member. If this rule is not followed, unpredictable results may occur.

Default

mode standard

Parameters

mesh-group

Keyword to specify that SA messages received from a mesh group member are always accepted but are not flooded to other members of the same mesh group. These SA messages are only flooded to non-mesh group peers or members of other mesh groups.

standard

Keyword to specify a non-meshed mode.

rpf-table

Syntax

rpf-table {rtable-m | rtable-u | both}

no rpf-table

Context

config>router>msdp

Platforms

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

Description

This command configures the sequence of route tables used to find an 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 or RP. However, the user can specify one of the following options:

- use the unicast route table only
- use the multicast route table only
- use both route tables

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

Default

rpf-table rtable-u

Parameters

rtable-m

Specifies that only the multicast route table is used by the multicast protocol (MSDP) for IPv4 RPF checks. This route table contains routes submitted by static routes, IS-IS, and OSPF.

rtable-u

Specifies that only the unicast route table is used by the multicast protocol (MSDP) for IPv4 RPF checks. This route table contains routes submitted by all the unicast routing protocols.

both

Specifies that the multicast route table is always looked up first and, if there is a route, use it. If MSDP does not find a route in the first lookup, it tries to find it in the unicast route table. The multicast route table (rtable-m) is checked before the unicast route table (rtable-u).

sa-timeout

Syntax

sa-timeout *seconds*

no sa-timeout

Context

config>router>msdp

Platforms

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

Description

This command configures the value for the SA entries in the cache. If these entries are not refreshed within the timeout value, they are removed from the cache. Normally, entries are refreshed at least once a minute. But under a high load with many MSDP peers, the refresh cycle could be incomplete. A higher timeout value (more than 90 seconds) could be useful to prevent instabilities in the MSDP cache.

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

Default

sa-timeout 90

Parameters

seconds

Specifies the time, in seconds, to wait for a response from the peer before declaring the peer unavailable.

Values 90 to 600

source

Syntax

[no] **source** *ip-prefix/mask*

Context

config>router>msdp

Platforms

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

Description

This command limits the number of active source messages that the router accepts from sources in the specified address range.

If the specified prefix and mask are already configured, this command only provides the context to configure the parameters pertaining to this active source message filter.

If the prefix and mask provided are not already configured, the source node instance must be created and the commands in the **config>router>msdp>source** context, which configure parameters for the source node instance, become available.

The source active **msdp** messages are not rate limited based on the source address range.

The **no** form of this message removes the SA rate limiter for this source address range.

Parameters

ip-prefix

Specifies the IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.

Values ip-prefix/mask: ip-prefix a.b.c.d (host bits must be 0)

mask

Specifies the subnet mask for the range expressed as a decimal integer mask length or in dotted decimal notation.

Values 0 to 32 (mask length), 0.0.0.0 to 255.255.255.255 (dotted decimal)

2.4.2.1.6 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 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 to determine whether bidirectional adjacencies exist.

Parameters

ip-address

Specifies the IP address of the multicast capable target router should be entered.

dns-name

Specifies the DNS name, if the DNS name resolution is configured.

Values *ip-address*: ipv4 unicast address (a.b.c.d)
 dns-name: [max 63 chars]

router *router-instance*

Specifies the router instance that this command applies to.

Default *router-name* — "Base" | "management", Default — Base

service-name

Specifies the service instance that this command applies to, up to 64 characters.

Output

The following output is an example of mrinfo, 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 -> 239.200.200.5 [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]
? 10.1.1.1 -> ? 10.1.1.2 [1/0/pim]
? 10.1.1.1 -> ? 10.1.1.9 [1/0/pim]
? 10.1.1.1 -> ? 0.0.0.0 [1/0/pim/leaf]
? 10.1.1.1 -> ? 10.1.1.7 [1/0/pim]
? 10.1.2.1 -> ? 10.1.2.7 [1/0/pim]
```

Table 7: Output fields: mrinfo

Label	Description
General flags	
version	Indicates software version on queried router
prune	Indicates the router understands pruning
genid	Indicates the router sends generation IDs
mtrace	Indicates the router handles mtrace requests
Neighbors flags	
1	Metric
0	Threshold (multicast time-to-live)
pim	PIM enabled on the interface
down	Operational status of the interface
disabled	Administrative status of the interface
leaf	No downstream neighbors on the 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 if DNS name resolution is configured.

Values ip-address: ipv4 unicast address (a.b.c.d)
 dns-name: [max 63 chars]

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

hop *hop*

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

Values 1 to 255

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

router *router-instance*

Specifies the router instance that this command applies to.

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

service-name *service-name*

Specifies the service instance that this command applies to, up to 64 characters.

wait-time *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 multicast path tracing 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
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 is 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]
```


group summary

Context

show>router>igmp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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.

Output

The following output is an example of IGMP multicast 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
=====
(*,224.0.0.1)
  Fwd List   : 239.112.1.2           Up Time : 0d 00:00:21
(10.11.0.1,224.0.0.1)
  Fwd List   : 239.112.1.1           Up Time : 0d 00:00:30
  Blk List   : 239.112.1.2           Up Time : 0d 00:00:21
(10.11.0.2,224.0.0.1)
  Fwd List   : 239.112.1.1           Up Time : 0d 00:00:30
(*,224.0.0.2)
  Fwd List   : 239.112.1.2           Up Time : 0d 00:00:21
(10.11.0.1,224.0.0.2)
  Blk List   : 239.112.1.2           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
=====
(*,224.0.0.1)                 1         0
(10.11.0.1,224.0.0.1)         1         1
(10.11.0.2,224.0.0.1)         1         0
```

```
(*,224.0.0.2)          1          0
(10.11.0.1,225.0.0.2)  0          1
-----
(*,G)/(S,G) Entries : 5
=====
*B:Dut-C#

A:NYC# show router igmp group 224.24.24.24
=====
IGMP Groups
=====
(*,239.24.24.24)          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 that4 are statically configured
Fwd List	Displays the list of interfaces in the forward list
Blk List	Displays the list of interfaces in the blocked list

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.

Output

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

Sample output

```
=====
```

IGMP SSM Translate Entries		
Group Range	Source	Interface
<239.1.1.1 - 239.1.1.2>	10.1.1.1	
<239.1.1.1 - 239.1.1.5>	10.1.1.2	ies-abc

Table 10: 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

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 only the specified IP interface name.

ip-address

Displays the information associated with only the specified IP address.

group *grp-address*

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

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 11: Output fields: IGMP interface](#) describes the fields.

Sample output

```
show router igmp interface output

*A:ALA-BA# show router 100 igmp interface
=====
IGMP Interfaces
=====
Interface Adm Oper Querier Cfg/Opr Num Policy
Version Groups
-----
IGMP_to_CE Up Up 10.1.1.1 1/1 3 none
-----
Interfaces : 1
=====

*A:ALA-BA#

*A:7210SAS>show>router# interface
=====
Interface Table (Router: Base)
=====
Interface-Name      Adm      Opr (v4/v6)  Mode      Port/SapId
IP-Address          PfxState
-----
system              Up        Up/Down      Network    system
10.3.3.3/32         n/a
toD_1               Up        Up/Down      Network    1/1/6
10.1.3.3/24         n/a
toE_1               Up        Up/Down      Network    1/1/5
10.1.2.3/24         n/a
toIxia_1            Up        Up/Down      IES        1/1/16:1
10.2.1.3/24         n/a
-----
Interfaces : 4
=====

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

Table 11: 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
Querier	Displays the address of the IGMP querier on the IP subnet to which the interface is attached

Label	Description
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 is zero.
Cfg/Opr Version Admin/Oper version	Cfg — 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 — operational version of IGMP running on this interface. For example, if the cfg value is 3 but all the routers in the local subnet of this interface use IGMP version v1 or v2, the operational version is v1 or v2
Num Groups	Displays the number of multicast groups that 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 was 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 reports 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 <i>source-list</i> 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 <i>source-list</i> parameter.
V1 Host Timer	Displays the time remaining until the local router assumes there are no longer any IGMP version 1 members on the IP subnet attached to this interface. Upon hearing an IGMPv1 Membership Report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores an IGMPv2 Leave messages for this group that it receives on this interface.
V2 Host Timer	Displays the time remaining until the local router assumes there are no longer IGMP version 2 members on the IP subnet attached to this interface. Upon hearing an IGMPv2 Membership Report, this value is reset to the group membership timer. While

Label	Description
	this time remaining is non-zero, the local router ignores 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 is set to "dynamic." For statically configured groups, the value is set to "static."
Compat Mode	Used for routers to be compatible with older version routers. IGMPv3 hosts MUST operate in version 1 and version 2 compatibility modes. IGMPv3 hosts MUST keep state per local interface about the compatibility mode of each attached network. A host compatibility mode is determined by the Host Compatibility Mode variable, which can be in one of the following 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 Older Version Querier Present timers for the interface.

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 only the specified IP interface name.

ip-addr

Displays the information associated with only the specified IP address.

Output

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

Sample output

```
*A:ALA-BA# show router 100 igmp static
```

=====		
IGMP Static Group Source		
=====		
Source	Group	Interface

10.11.11.11	239.136.22.3	IGMP_to_CE
*	239.1.1.1	IGMP_to_CE
10.22.22.22	239.255.255.255	IGMP_to_CE

Static (*,G)/(S,G) Entries : 3		
=====		
*A:ALA-BA#		

Table 12: Output fields: IGMP static

Label	Description
Source	Displays entries that represent a source address from which receivers are interested or 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*]

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 only the specified IP interface name.

ip-address

Displays the information associated with only the specified IP address.

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	Queries — number of IGMP general queries transmitted or received on this interface Report — total number of IGMP V1, V2, or V3 reports transmitted or received on this interface Leaves — total number of IGMP leaves transmitted on this interface
Received	Displays the total number of IGMP packets received on this interface

Label	Description
Transmitted	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 that do 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
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 retries a query

2.4.2.2.2 Router PIM commands

anycast

Syntax

anycast [family] [detail]

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

family

Indicates an IPv4 address.

detail

Displays detailed information.

Output

The following output is an example of a PIM anycast configuration 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 [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.

Output

The following output is an example of PIM CRP information, 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  224.0.0.0/4        192        150      0d 00:02:19
239.22.187.239  224.0.0.0/4        192        150      0d 00:02:19
239.22.187.240  224.0.0.0/4        192        150      0d 00:02:09
-----
Candidate RPs : 3
=====
A:7210SAS#
```

Table 16: Output fields: PIM CRP

Label	Description
RP Address	Displays the CRP address
Group Address	Displays the range of multicast group addresses for which the CRP is the Candidate RP
Priority	Displays the CRP priority for becoming an RP. This value is used to elect RP for a group range. A value of 0 is considered the highest priority.
Holdtime	Displays the hold time of the CRP. 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.

Label	Description
Expiry	Displays the minimum time remaining before the CRP is 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 are displayed.

type starg

Specifies that only (*,G) entries are displayed.

type sg

Specifies that only (S,G) entries are displayed.

detail

Displays detailed group information.

family

Displays either IPv4 or IPv6 information.

Output

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

Sample output

```
*A:7210SAS# show router 1 pim group

=====
PIM Groups ipv4
=====
Group Address          Type      Spt Bit Inc Intf      No.Oifs
Source Address         RP
-----
239.1.1.1              (S,G)          mpls-if-74734  1
10.1.1.2
239.1.1.2              (S,G)          mpls-if-74734  1
10.1.1.2
239.1.1.3              (S,G)          mpls-if-74734  1
10.1.1.2
239.1.1.4              (S,G)          mpls-if-74734  1
10.1.1.2
239.1.1.5              (S,G)          mpls-if-74734  1
10.1.1.2
239.1.1.6              (S,G)          mpls-if-74734  1
10.1.1.2
239.1.1.7              (S,G)          mpls-if-74734  1
10.1.1.2
239.1.1.8              (S,G)          mpls-if-74734  1
10.1.1.2
239.1.1.9              (S,G)          mpls-if-74734  1
10.1.1.2

*A:7210SAS#

*A:7210SAS#show router 1 pim group detail

=====
PIM Source Group ipv4
=====
Group Address      : 239.1.1.1
Source Address     : 10.1.1.2
RP Address         : 0
Advt Router        : 10.17.17.17
Flags              :
Type               : (S,G)
MRIB Next Hop      : 10.17.17.17
MRIB Src Flags     : remote
Up Time            : 5d 21:12:47
Keepalive Timer    : Not Running
Resolved By        : rtable-u

Up JP State        : Joined
Up JP Rpt          : Not Joined StarG
Up JP Expiry       : 0d 00:00:52
Up JP Rpt Override : 0d 00:00:00

Register State     : No Info
Reg From Anycast RP: No

Rpf Neighbor       : 10.17.17.17
Incoming Intf      : mpls-if-74734
Outgoing Intf List : vprn1_1

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 The source address is 0 if the type is configured as starg . The source address is the address of the RP if the type is configured as starstarrp .
RP Address	Displays the RP address
Type	Specifies the type of entry, (*,*, rp)/(*,G) or (S,G)
Spt Bit	Specifies whether to forward on (*,*, rp)/(*,G) or on (S,G) state. This 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.
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 contains 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 is 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

Label	Description
	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 a "notPruned" 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 the router does not see a join (S,G, rpt) message, on expiry of the override timer, it sends its 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	Displays 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 occurs based on the register stop timer and the response received from the RP.
Register Stop Exp	Displays the time remaining before the register state may transition to a different state
Register from Anycast RP	Displays whether 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 only with the specified IP interface name.
- ip-address**
Displays the interface information associated with only 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 is zero.
If the type is **starstarrp**, the value of this object is the address of the RP.
- type**
Specifies the type of this entry.

Values **starstarrp, starg, sg**
- detail**
Displays detailed interface information.
- family**
Displays IPv4 or IPv6 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
-----
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, consider 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** parameter is not defined in this example, the **show** command output would display two adjacencies.

Parameters

- neighbor ip-int-name**
Displays the interface information associated only with the specified IP interface name.
- neighbor ip-address**
Displays the interface information associated only 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 either IPv4 or IPv6 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 neighbors address
Up Time	Displays the time since this PIM neighbor (last) became a neighbor of the local router
Expiry Time	Displays the minimum time remaining before this PIM neighbor is aged out 0 — Means that this neighbor is never be aged out. This happens when the PIM neighbor sends a Hello message with a hold-time 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 is 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, old Hello information about that neighbor is discarded and replaced 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 RP set information built by the router.

Parameters

family

Displays either IPv4 or IPv6 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 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 PIM RP hashing 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        229.200.200.4   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

Label	Description
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

s-pmsi

Syntax

s-pmsi [**group-ip** *group-ip*] [**source-ip** *source-ip*] [**detail**]

s-pmsi [*mdSrcAddr* [*mdGrpAddr*]] [**group-ip** *group-ip*] [**source-ip** *source-ip*] [**detail**]

s-pmsi **ext-tunnel-id** *ext-tunnel-id* [**tunnel-id** *tunnel-id*] [**group-ip** *group-ip*] [**source-ip** *source-ip*] [**detail**]

s-pmsi **root-addr** *root-addr* [**lsp-id** *lsp-id*] [**group-ip** *group-ip*] [**source-ip** *source-ip*] [**detail**]

Context

show>router>pim

Platforms

7210 SAS-T (network mode only), 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Mxp, and 7210 SAS-Sx/S
1/10GE (in standalone mode)

Description

Displays the list of selective provider multicast service interfaces that are currently active.

Parameters

group-ip

Specifies the group IP address.

source-ip

Specifies the source IP address.

mdSrcAddr

Displays PIM S-PMSI information associated with the specified source address of the multicast sender.

mdGrpAddr

Displays PIM S-PMSI information associated with the specified group address of the multicast sender.

ext-tunnel-id

Displays PIM S-PMSI information associated with the specified external tunnel ID.

tunnel-id

Displays PIM S-PMSI information associated with the specified tunnel ID.

- root-addr**
Displays PIM S-PMSI information associated with the specified root address.
- lsp-id**
Displays PIM S-PMSI information associated with the specified LSP.
- detail**
Displays detailed output.

Output

The following outputs are examples of S-PMSI information, and [Table 22: Output fields: S-PMSI](#) describes the output fields.

- [Sample output PIM selective provider tunnel](#)
- [Sample output PIM selective provider tunnel detail](#)
- [Sample output RX tracking for RSVP S-PMSI tunnel](#)
- [Sample output RX tracking for RSVP S-PMSI tunnel detail](#)
- [Sample output TX tracking for RSVP S-PMSI tunnel detail](#)

Sample output PIM selective provider tunnel

```
*B:node-6# show router 100 pim s-psmi
=====
PIM Selective provider tunnels
=====
MD Src Address      MD Grp Address      MT Index      Num VPN SGs
-----
239.200.200.7       239.0.89.72         24603         1
239.200.200.7       239.0.89.73         24604         1
239.200.200.7       239.0.89.74         24605         1
239.200.200.7       239.0.89.75         24606         1
239.200.200.7       239.0.89.76         24607         1
239.200.200.7       239.0.89.77         24608         1
239.200.200.7       239.0.89.78         24609         1
239.200.200.7       239.0.89.79         24610         1
239.200.200.7       239.0.89.80         24611         1
239.200.200.7       239.0.89.81         24612         1
239.200.200.7       239.0.89.82         24613         1
239.200.200.7       239.0.89.83         24614         1
239.200.200.7       239.0.89.84         24615         1
239.200.200.7       239.0.89.85         24616         1
239.200.200.7       239.0.89.86         24617         1
239.200.200.7       239.0.89.87         24618         1
...
=====
*B:node-6#
```

Sample output PIM selective provider tunnel detail

```
*B:node-6# show router 100 pim s-psmi
=====
PIM Selective provider tunnels
=====
MD Src Address      MD Grp Address      MT Index      Num VPN SGs
-----
239.200.200.7       239.0.89.72         24603         1
239.200.200.7       239.0.89.73         24604         1
```



```

239.200.200.7      239.0.89.74      24605      1
239.200.200.7      239.0.89.75      24606      1
239.200.200.7      239.0.89.76      24607      1
239.200.200.7      239.0.89.77      24608      1
239.200.200.7      239.0.89.78      24609      1
239.200.200.7      239.0.89.79      24610      1
239.200.200.7      239.0.89.80      24611      1
239.200.200.7      239.0.89.81      24612      1
239.200.200.7      239.0.89.82      24613      1
239.200.200.7      239.0.89.83      24614      1
239.200.200.7      239.0.89.84      24615      1
239.200.200.7      239.0.89.85      24616      1
239.200.200.7      239.0.89.86      24617      1
239.200.200.7      239.0.89.87      24618      1
...
=====

```

```
*B:node-6#
```

Sample output RX tracking for RSVP S-PMSI tunnel

```

*A:Dut-C# show router 1 pim s-pmsi
=====
PIM RSVP Spmsi tunnels
=====
P2mp ID   Tunnel ID   Ext Tunnel Adrs      SPMSI Index   Num VPN   State
          SGs
-----
0          0          10.20.1.4            1030144       1         RX Tracking
0          0          10.20.1.4            1030144       1         RX Tracking
=====
PIM RSVP Spmsi Interfaces : 2
=====
*A:Dut-C# show router 21 pim s-pmsi
=====
PIM LDP Spmsi tunnels
=====
Lsp ID    Root Addr      SPMSI Index   Num VPN   State
          SGs
-----
0          10.20.1.4      1030144       1         RX Tracking
0          10.20.1.4      1030144       1         RX Tracking
=====
PIM LDP Spmsi Interfaces : 2
=====
*A:Dut-C#

```

Sample output RX tracking for RSVP S-PMSI tunnel detail

```

*A:Dut-C# show router 1 pim s-pmsi detail
=====
PIM RSVP Spmsi tunnels
=====
P2MP ID      : 0          Tunnel ID      : 0
Ext Tunnel Adrs : 10.20.1.4   Spmsi IfIndex  : 1030144
Number of VPN SGs : 1          Uptime        : 0d 00:02:48
VPN Group Address : 239.100.0.0
VPN Source Address : 10.1.101.2
State         : RX Tracking   Mdt Threshold  : 0
=====
PIM RSVP Spmsi tunnels
=====
P2MP ID      : 0          Tunnel ID      : 0
Ext Tunnel Adrs : 10.20.1.4   Spmsi IfIndex  : 1030144

```

```

Number of VPN SGs : 1          Uptime : 0d 00:02:47
VPN Group Address : ff0e:db8:225:100::
VPN Source Address : 2001:db8:1:101::2
State : RX Tracking          Mdt Threshold : 0
=====
PIM RSVP Spmsi Interfaces : 2
=====
*A:Dut-C#
*A:Dut-C# show router 21 pim s-pmsi detail
=====
PIM LDP Spmsi tunnels
=====
LSP ID : 0
Root Addr : 10.20.1.4          Spmsi IfIndex : 1030144
Number of VPN SGs : 1          Uptime : 0d 00:03:35
VPN Group Address : 239.100.0.0
VPN Source Address : 10.1.101.2
State : RX Tracking          Mdt Threshold : 0
=====
PIM LDP Spmsi tunnels
=====
LSP ID : 0
Root Addr : 10.20.1.4          Spmsi IfIndex : 1030144
Number of VPN SGs : 1          Uptime : 0d 00:03:34
VPN Group Address : ff0e:db8:225:100::
VPN Source Address : 2001:db8:1:101::2
State : RX Tracking          Mdt Threshold : 0
=====
PIM LDP Spmsi Interfaces : 2
=====
*A:Dut-C#

```

Sample output TX tracking for RSVP S-PMSI tunnel detail

```

*A:Dut-C# show router 1 pim s-pmsi detail
=====
PIM RSVP Spmsi tunnels
=====
P2MP ID : 1          Tunnel ID : 61442
Ext Tunnel Addr : 10.20.1.4      Spmsi IfIndex : 74230
Number of VPN SGs : 1          Uptime : 0d 00:05:11
VPN Group Address : 239.100.0.0
VPN Source Address : 10.1.101.2
State : TX Join Pending      Mdt Threshold : 1
Join Timer : N/A            Holddown Timer : 0d 00:00:47
Receiver Count : 4
=====
PIM RSVP Spmsi tunnels
=====
P2MP ID : 1          Tunnel ID : 61443
Ext Tunnel Addr : 10.20.1.4      Spmsi IfIndex : 74231
Number of VPN SGs : 1          Uptime : 0d 00:05:10
VPN Group Address : ff0e:225:100::
VPN Source Address : 2001:db8:1:101::2
State : TX Join Pending      Mdt Threshold : 1
Join Timer : N/A            Holddown Timer : 0d 00:00:50
Receiver Count : 4
=====
PIM RSVP Spmsi Interfaces : 2
=====
*A:Dut-C#
*A:Dut-D# show router 21 pim s-pmsi detail
=====

```

```

PIM LDP Spmsi tunnels
=====
LSP ID           : 8194
Root Addr        : 10.20.1.4          Spmsi IfIndex      : 74228
Number of VPN SGs : 1                Uptime             : 0d 00:05:56
VPN Group Address : 239.100.0.0
VPN Source Address : 10.1.101.2
State            : TX Join Pending    Mdt Threshold      : 1
Join Timer       : N/A                Holddown Timer     : 0d 00:00:02
Receiver Count   : 4
=====
PIM LDP Spmsi tunnels
=====
LSP ID           : 8195
Root Addr        : 10.20.1.4          Spmsi IfIndex      : 74229
Number of VPN SGs : 1                Uptime             : 0d 00:05:55
VPN Group Address : ff0e:db8:225:100::
VPN Source Address : 2001:db8:1:101::2
State            : TX Join Pending    Mdt Threshold      : 1
Join Timer       : N/A                Holddown Timer     : 0d 00:00:05
Receiver Count   : 4
=====
PIM LDP Spmsi Interfaces : 2
=====
*A:Dut-D#

```

Table 22: Output fields: S-PMSI

Label	Description
MD Grp Address	Displays the IP multicast group address for which this entry contains information
MD Src Address	Displays the source address of the multicast sender A value of 0 (zero) indicates the type is configured as starg .
MT Index MT IfIndex	Displays the index number
Num VP SGs	Displays the number of VPN (S,G)s
Uptime	Displays the length of time that the S-PMSI has been up
Egress Fwding Rate	Displays the egress forwarding rate for the S-PMSI
VPN Group Address	Displays the VPN group address for the S-PMSI
VPN Source Address	Displays the VPN source address for the S-PMSI
Expiry Timer	Displays the minimum time remaining before this S_PMSI is aged out A value of 0 (zero) means that this S-PMSI is never aged out, which occurs when the PIM neighbor sends a Hello message with hold time set to 0xffff.

statistics

Syntax

statistics [*ip-int-name* | *int-ip-address* | *mpls-if-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**
Only displays the interface information associated with the specified IP interface name.
- int-ip-address**
Only displays the interface information associated with the specified IP address.
- mpls-if-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 either IPv4 or IPv6 information.

Output

The following output is an example of PIM statistics information, and [Table 23: 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 23: Output fields: PIM statistics

Label	Description
PIM Statistics	Displays the PIM statistics for a particular interface
Message Type	Displays the type of message Hello — displays the number of PIM Hello messages received or transmitted on this interface Join Prune — displays the number of PIM join prune messages received or transmitted on this interface Asserts — displays the number of PIM assert messages received or transmitted on this interface Register — displays the number of register messages received or transmitted on this interface Null Register — displays the number of PIM null register messages received or transmitted on this interface Register Stop — displays the number of PIM register stop messages received or transmitted on this interface BSM — displays the number of PIM Bootstrap messages (BSM) received or transmitted on this interface Candidate RP Adv — displays the number of candidate RP advertisements Total Packets — displays the total number of packets transmitted and received on this interface
Received	Displays the number of messages received on this interface

Label	Description
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 that were discarded because of bad checksum
Rx Bad Encoding	Displays the number of PIM messages with bad encodings received on this interface
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 that 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 group contained in the message
Register Policy Drops	Displays the number of times the register policy match resulted in dropping PIM register message
Bootstrap Import Policy Drops	Displays the number of Bootstrap messages received on this interface but were dropped because of Bootstrap import policy

Label	Description
Bootstrap Export Policy Drops	Displays the number of Bootstrap messages that were not transmitted on this interface because of Bootstrap export policy
Source Group Statistics	Displays the 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 PIM status. The Oper Status reflects the combined operational status of IPv4 or IPv6 PIM protocol status. If both are down, Oper Status is reflected as down. If IPv4 or IPv6 reflects up, the Oper Status reflects 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 either IPv4 or IPv6 information.

Output

The following output is an example of PIM status information, and [Table 24: 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 towards 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 24: 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

Label	Description
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 towards	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 C-RP mechanism
Address	Displays the local RP address
Priority	Displays the CRP priority for becoming an RP. A value of 0 is the highest priority
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 whether hash-based multicast balancing of traffic over ECMP links is enabled

tunnel-interface

Syntax

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

Context

show>router>pim

Platforms

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

Description

This command displays PIM tunnel interface information.

Parameters

- ip-int-name**
Specifies the IP interface name. A string up to 32 characters.
- mt-int-name**
Specifies the Multicast Tunnel (MT) interface for a VPRN.
- int-ip-address**
Specifies the interface IPv4 or IPv6 address.
- group-ip-address**
Specifies the IP multicast group address, or 0.
- ip-address**
Specifies the source or RP IPv4 or IPv6 address.
- type**
Specifies the type of entry.

Values starstarrp, starg, sg
- detail**
Displays detailed interface information.
- family**
Specifies the IPv4 or IPv6 address family.

Output

The following output is an example of PIM tunnel interface information.

Sample output

```
*A:Dut-C# show router pim tunnel-interface
=====
PIM Interfaces ipv4
=====
Interface                Originator Address  Adm  Opr  Transport Type
-----
mpls-if-73728             N/A                Up   Up   Tx-IPMSI
mpls-if-73729             N/A                Up   Up   Tx-IPMSI
mpls-if-73730             N/A                Up   Up   Tx-IPMSI
mpls-if-73731             N/A                Up   Up   Tx-IPMSI
mpls-if-73732             N/A                Up   Up   Tx-IPMSI
-----
Interfaces : 5
```

```
=====
```

mvpn

Syntax

mvpn

Context

show>router

Platforms

7210 SAS-T (network mode only), 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Mxp, and 7210 SAS-Sx/S
1/10GE (in standalone mode)

Description

This command displays multicast VPN information for the specified router instance.

Output

The following output is an example of PIM MVPN information.

Sample output

```
*A:7210SAS# show router 10 mvpn

=====
MVPN 10 configuration data
=====
signaling : Bgp auto-discovery : Default
UMH Selection : Highest-IP intersite-shared : Enabled
vrf-import : N/A
vrf-export : N/A
vrf-target : unicast
C-Mcast Import RT : target:10.16.16.16:3

ipmsi : ldp
i-pmsi P2MP AdmSt : Up

spmsi : ldp
s-pmsi P2MP AdmSt : Up
max-p2mp-spmsi : 251
data-delay-interval: 3 seconds
enable-asm-mdt : N/A
data-threshold : 224.0.0.0/4 --> 1 kbps

=====
*A:7210SAS#
```

mvpn-list

Syntax

mvpn-list

Context

show>router

Platforms

7210 SAS-T (network mode only), 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Mxp, and 7210 SAS-Sx/S
1/10GE (in standalone mode)

Description

This command displays multicast VPN list information for the specified router instance.

Output

The following output is an example of PIM MVPN list information.

Sample output

```
A:7210SAS# show router mvpn-list

=====
MVPN List
=====
VprnID Sig A-D iPmsi/sPmsi GroupAddr/Lsp-Template (S,G)/(*,G)
-----
10 Bgp Default Mldp/Mldp N/A 512/0
20 Bgp Default Mldp/Mldp N/A 512/0
30 Bgp Default None/None N/A 0/0
-----
Total PIM I-PMSI tunnels : 0
Total RSVP I-PMSI tunnels : 0
Total MLDP I-PMSI tunnels : 2
Total PIM TX S-PMSI tunnels : 0
Total RSVP TX S-PMSI tunnels : 0
Total MLDP TX S-PMSI tunnels : 502
Total PIM RX S-PMSI tunnels : 0
Total RSVP RX S-PMSI tunnels : 0
Total MLDP RX S-PMSI tunnels : 0
Total (S,G) : 1024
Total (*,G) : 0
Total Mvpns : 3
Sig = Signal Pim-a = pim-asm Pim-s = pim-ssm A-D = Auto-Discovery
=====
*A:7210SAS#
```

tunnel-table

Syntax

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

Context

```
show>router
```

Platforms

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

Description

This command displays tunnel table information.

Parameters

- protocol**

Specifies the protocol.

Values bgp, ldp, rsvp, sdp
- sdp-id**

Specifies the SDP ID.

Values 1 to 17407

Output

The following output is an example of PIM tunnel table information.

Sample output

```
*A:Dut-C# show router pim tunnel-interface
=====
PIM Interfaces ipv4
=====
Interface                Originator Address  Adm  Opr  Transport Type
-----
mpls-if-73728             N/A                Up   Up   Tx-IPMSI
mpls-if-73729             N/A                Up   Up   Tx-IPMSI
mpls-if-73730             N/A                Up   Up   Tx-IPMSI
mpls-if-73731             N/A                Up   Up   Tx-IPMSI
mpls-if-73732             N/A                Up   Up   Tx-IPMSI
-----
Interfaces : 5
=====
```

2.4.2.2.3 MSDP commands

group

Syntax

group [group-name] [detail]

Context

show>router>msdp

Platforms

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

Description

This command displays information about MSDP groups.

Parameters

- group-name**
Displays information about the specified group name, up to 32 characters. If no *group-name* is specified, information about all group names displays.
- detail**
Displays detailed MSDP group information.

Output

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

Sample output

```
*A:ALA-48>show>router>msdp# group
=====
MSDP Groups
=====
Group Name           Mode      Act Srcs  Local Address
-----
main                 Mesh-group None None
loop1                Mesh-group None None
loop2                Mesh-group None None
loop3                Mesh-group None None
loop4                Mesh-group None None
loop5                Mesh-group None None
-----
Groups : 6
=====
*A:ALA-48>show>router>msdp#

*A:ALA-48>show>router>msdp# group test
=====
```

```

MSDP Groups
=====
Group Name           Mode      Act Srcs  Local Address
-----
test                 Mesh-group 50000    10.10.10.103
-----
Groups : 1
=====
*A:ALA-48>show>router>msdp#

*A:ALA-48>show>router>msdp# group test detail
=====
MSDP Groups
=====
Group Name           : test
-----
Local Address        : 10.10.10.103
Admin State          : Up
Receive Msg Rate     : None
Receive Msg Time     : None
Mode                 : Mesh-group
SA Limit             : 50000
Export Policy        : None Specified / Inherited
Import Policy        : None Specified / Inherited
-----
Groups : 1
=====
*A:ALA-48>show>router>msdp#

```

Table 25: Output fields: MSDP group

Label	Description
Group Name	Displays the MSDP group name.
Mode	Displays the groups of peers in a full mesh topology to limit excessive flooding of SA messages to neighboring peers.
Act Srcs	Displays the configured maximum number of active source messages that are accepted by MSDP.
Local Address	Displays the local end of an MSDP session.
Admin State	Displays the administrative state.
Receive Msg Rate	Displays the rate at which the messages are read from the TCP session.
Receive Msg Time	Displays the time of MSDP messages that are read from the TCP session within the configured number of seconds.
Receive Msg Thd	Displays the configured threshold number of MSDP messages that can be processed before the MSDP message rate limiting function.
SA Limit	Displays the SA limit.

peer

Syntax

peer [*ip-address*] [**group** *group-name*] [**detail**]

Context

show>router>msdp

Platforms

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

Description

This command displays information about MSDP peers.

Parameters

- ip-address**
Displays information about the specified IP address. If no IP address is specified, information about all MSDP IP addresses displays.
- group-name**
Displays information about the specified group name. If no *group-name* is specified, information about all MSDP peers displays.
- detail**
Displays detailed MSDP peer information.

Output

The following output is an example of MSDP peer information, and [Table 26: Output fields: MSDPpeer](#) describes the output fields.

Sample output

```
A:ALA-48# show router msdp peer
=====
MSDP Peers
=====
Peer           Local Address   State           Last State Change   SA Learnt
-----
10.20.1.1      10.20.1.6       Established      08/30/2002 03:22:13 1008
-----
Peers : 1
=====
A:ALA-48#

A:ALA-48# show router msdp peer detail
=====
MSDP Peers
-----
Peer Address      : 10.20.1.1
-----
```



```

Group Name      : None
Local Address   : 10.20.1.6
Last State Change : 08/30/2002 03:22:13 Last Act Src Limit : N/A
Peer Admin State : Up                      Default Peer      : No
Peer Connect Retry : 0                      State             : Established
SA accepted      : 1008                     SA received         : 709
State timer expires: 18                     Peer time out       : 62
Active Source Limit: None                   Receive Msg Rate     : 0
Receive Msg Time  : 0                       Receive Msg Thd      : 0
Auth Status       : Disabled                 Auth Key             : None
Export Policy     : None Specified / Inherited
Import Policy     : None Specified / Inherited
-----
Peers : 1
=====
A:ALA-48#

```

Table 26: Output fields: MSDPpeer

Label	Description
Peer	Displays the IP address of the peer.
Local Address	Displays the local IP address.
State	Displays the current state of the peer.
Last State Change	Displays the date and time of the peer's last state change.
SA Learn	Displays the number of SA messages learned through a peer.

source

Syntax

source [*ip-address/mask*] [**type** {**configured** | **dynamic** | **both**}] [**detail**]

Context

show>router>msdp

Platforms

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

Description

This command displays the discovery method for this multicast source.

Parameters

configured

Displays user-created sources.

dynamic

Displays dynamically created sources.

both

Displays both user-configured and dynamically created sources.

detail

Displays detailed MSDP source information.

Output

The following output is an example of MSDP source information, and [Table 27: Output fields: MSDP source](#) describes the output fields.

Sample output

```
*A:Dut-G>config>router>msdp>source# /show router msdp source
=====
MSDP Sources
=====
Source           Type           SA Limit   Num Excd   Last Exceeded
-----
100.112.1.2/32   Configured    10         0          N/A
-----
Sources : 1
=====
```

Table 27: Output fields: MSDP source

Label	Description
Source	Displays the source address.
Type	Displays the source type.
SA Limit	Displays the limit of SA messages allowed.
Num Excd	Displays the number of SA messages that have exceeded the limit.
Last Exceeded	Displays the date and time that the limit was last exceeded.
Sources	Displays the number of sources.

source-active

Syntax

source-active [{**group** *ip-address* | **local** | **originator** *ip-address* | **peer** *ip-address* | **source** *ip-address* | **group** *ip-address* **source** *ip-address*}] [**detail**]

Context

show>router>msdp

Platforms

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

Description

This command displays SA messages accepted by MSDP.

Parameters

- group ip-address**
Displays information about the specified group IP address.
- local**
Displays information about local SA messages.
- originator ip-address**
Displays information about the specified originator IP address.
- peer ip-address**
Displays information about the specified peer IP address.
- source ip-address**
Displays information about the specified source IP address.
- group ip-address**
Displays information about the specified group IP address.
- detail**
Displays detailed MSDP SA information.

Output

The following output is an example of MSDP SA information, and [Table 28: Output fields: MSDP source-active](#) describes the output fields.

Sample output

```
A:ALA-48# show router msdp source-active
=====
MSDP Source Active Info
=====
Grp Address      Src Address      Origin RP        Peer Address State Timer
-----
239.100.0.0      10.112.1.2      10.20.1.1        10.20.1.1 69
239.100.0.1      10.112.1.2      10.20.1.1        10.20.1.1 69
239.100.0.2      10.112.1.2      10.20.1.1        10.20.1.1 69
239.100.0.3      10.112.1.2      10.20.1.1        10.20.1.1 69
239.100.0.4      10.112.1.2      10.20.1.1        10.20.1.1 69
239.100.0.5      10.112.1.2      10.20.1.1        10.20.1.1 69
239.100.0.6      10.112.1.2      10.20.1.1        10.20.1.1 69
239.100.0.7      10.112.1.2      10.20.1.1        10.20.1.1 69
239.100.0.8      10.112.1.2      10.20.1.1        10.20.1.1 69
239.100.0.9      10.112.1.2      10.20.1.1        10.20.1.1 69
-----
MSDP Source Active : 10
=====
A:ALA-48#
```

```
A:ALA-48# show router msdp source-active detail
=====
MSDP Source Active
=====
Group Address      : 239.100.0.0      Source Address      : 10.112.1.2
Origin RP          : 10.20.1.1      Peer Address        : 10.20.1.1
State Timer        : 64              Up Time             : 3d 01:44:25
Group Address      : 239.100.0.1      Source Address      : 10.112.1.2
Origin RP          : 10.20.1.1      Peer Address        : 10.20.1.1
State Timer        : 64              Up Time             : 48d 18:22:29
Group Address      : 239.100.0.2      Source Address      : 10.112.1.2
Origin RP          : 10.20.1.1      Peer Address        : 10.20.1.1
State Timer        : 64              Up Time             : 48d 18:22:29
Group Address      : 239.100.0.3      Source Address      : 10.112.1.2
Origin RP          : 10.20.1.1      Peer Address        : 10.20.1.1
State Timer        : 64              Up Time             : 48d 18:22:29
Group Address      : 239.100.0.4      Source Address      : 10.112.1.2
Origin RP          : 10.20.1.1      Peer Address        : 10.20.1.1
State Timer        : 64              Up Time             : 48d 18:22:29
Group Address      : 239.100.0.5      Source Address      : 10.112.1.2
Origin RP          : 10.20.1.1      Peer Address        : 10.20.1.1
Origin RP          : 10.20.1.1      Peer Address        : 10.20.1.1
State Timer        : 64              Up Time             : 48d 18:22:29
Group Address      : 239.100.0.7      Source Address      : 10.112.1.2
Origin RP          : 10.20.1.1      Peer Address        : 10.20.1.1
State Timer        : 64              Up Time             : 48d 18:22:29
Group Address      : 239.100.0.8      Source Address      : 10.112.1.2
Origin RP          : 10.20.1.1      Peer Address        : 10.20.1.1
State Timer        : 64              Up Time             : 48d 18:22:29
Group Address      : 239.100.0.9      Source Address      : 10.112.1.2
Origin RP          : 10.20.1.1      Peer Address        : 10.20.1.1
State Timer        : 64              Up Time             : 48d 18:22:29
-----
MSDP Source Active : 10
=====
A:ALA-48#
```

Table 28: Output fields: MSDP source-active

Label	Description
Grp Address	Displays the IP address of the group.
Src Address	Displays the IP address of the source.
Origin RP	Displays the origination RP address.
Peer Address	Displays the address of the peer.
State Timer	Displays the time-out value. If the value reaches zero, the SA is removed.

source-active-rejected

Syntax

source-active-rejected [**peer-group** *name*] [**group** *ip-address*] [**source** *ip-address*] [**originator** *ip-address*] [**peer** *ip-address*]

Context

show>router>msdp

Platforms

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

Description

This command displays SA messages rejected by MSDP.

Parameters

- peer-group *name***
Displays information about rejected SA messages for the specified peer group.
- group *ip-address***
Displays information about the specified group IP address.
- source *ip-address***
Displays information about the source address of the SA entry that is rejected.
- originator *ip-address***
Displays information about the specified originator IP address.
- peer *ip-address***
Displays information about the peer from which this rejected SA entry was last received.

Output

The following output is an example of MSDP rejected SA information, and [Table 29: Output fields: MSDP source-active rejected](#) describes the output fields.

Sample output

```
*A:ALA-48# show router msdp source-active-rejected
=====
MSDP Source Active Rejected Info
=====
```

Grp Address	Src Address	Origin RP	Peer Address	Reject Reason
239.100.0.1	10.0.0.1	10.20.0.1	239.0.0.1	Import Policy
239.100.0.2	10.0.0.2	10.20.0.2	239.0.0.2	Export Policy
239.100.0.3	10.0.0.3	10.20.0.3	239.0.0.3	RPF Failure
239.100.0.4	10.0.0.4	10.20.0.4	239.0.0.4	Limit Exceeded
239.100.0.5	10.0.0.5	10.20.0.5	239.0.0.5	Limit Exceeded
239.100.0.6	10.0.0.6	10.20.0.6	239.0.0.6	Limit Exceeded
239.100.0.7	10.0.0.7	10.20.0.7	239.0.0.7	Limit Exceeded

```
-----
```

```
SA Rejected Entries : 7
=====
*A:ALA-48#
```

Table 29: Output fields: MSDP source-active rejected

Label	Description
Grp Address	Displays the IP address of the group.
Src Address	Displays the IP address of the source.
Origin RP	Displays the origination rendezvous point (RP) address.
Peer Address	Displays the address of the peer.
Reject Reason	Displays the reason why this source active entry is rejected.

statistics

Syntax

statistics [**peer** *ip-address*]

Context

show>router>msdp

Platforms

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

Description

This command displays statistics information related to an MSDP peer.

Parameters

ip-address

Displays information about the specified peer IP address.

Output

The following output is an example of MSDP statistics information, and [Table 30: Output fields: MSDP statistics](#) describes the output fields.

Sample output

```
A:ALA-48# show router msdp statistics
=====
MSDP Statistics
=====
Glo ActSrc Lim Excd: 0
-----
Peer Address      : 10.20.1.1
```

```

-----
Last State Change : 0d 11:33:16      Last message Peer : 0d 00:00:17
RPF Failures      : 0                Remote Closes     : 0
SA Msgs Sent      : 0                SA Msgs Recvd     : 709
SA req. Msgs Sent : 0                SA req. Msgs Recvd : 0
SA res. Msgs Sent : 0                SA res. Msgs Recvd : 0
KeepAlive Msgs Sent: 694            KeepAlive Msgs Recd: 694
Unknown Msgs Sent : 0                Error Msgs Recvd  : 0
-----
Peers : 1
=====
A:ALA-48#

```

Table 30: Output fields: MSDP statistics

Label	Description
Last State Change	Displays the date and time the peer state changed.
RPF Failures	Displays the number of RPF failures.
SA Msgs Sent	Displays the number of SA messages sent.
SA req. Msgs Sent	Displays the number of SA request messages sent.
SA res. Msgs Sent	Displays the number of SA response messages sent.
KeepAlive Msgs Sent	Displays the number of keepalive messages sent.
Unknown Msgs Sent	Displays the number of unknown messages received.
Last message Peer	Displays the time the last message was received from the peer.
Remote Closes	Displays the number of times the remote peer close.
SA Msgs Recvd	Displays the number of SA messages received.
SA req. Msgs Recvd	Displays the number of SA request messages received.
SA res. Msgs Recvd	Displays the number of SA response messages received.
KeepAlive Msgs Recd	Displays the number of keepalive messages received.
Error Msgs Recvd	Displays the number of unknown messages received.

status

Syntax

status

Context

show>router>msdp

Platforms

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

Description

This command displays MSDP status information.

Output

The following output is an example of MSDP status information, and [Table 31: Output fields: MSDP status](#) describes the output fields.

Sample output

```
A:ALA-48# show router msdp status
=====
MSDP Status
=====
Admin State                : Up
Local Address              : None
Global Statistics
Active Src Limit           : None
Act Src Lim Excd           : 0
Num. Peers                 : 1
Num. Peers Estab           : 1
Num. Source Active         : 10
Policies                   : None
Data Encapsulation         : Enabled
Receive Msg Rate
Rate                       : 0
Time                       : 0
Threshold                  : 0
Last Msdp Enabled          : 08/30/2002 03:21:43
=====
A:ALA-48#
```

Table 31: Output fields: MSDP status

Label	Description
Admin State	Displays the administrative state.
Local Address	Displays the local IP address.
Active Src Limit	Displays the active source limit.
Act Src Lim Excd	Displays the active source limit which has been exceeded.
Num. Peers	Displays the number of peers.
Num. Peers Estab	Displays the number of peers established.
Num. Source Active	Displays the number of active sources.
Policies	Displays the policy to export the SA state from the SA list into MSDP.

Label	Description
Data Encapsulation	Displays the RP using MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP SA messages.
Rate	Displays the receive message rate.
Time	Displays the receive message time.
Threshold	Displays the number of MSDP messages that can be processed before the MSDP message rate limiting function is activated.
RPF Table	Displays the name of the reverse path forwarding table.
Last msdp Enabled	Displays the time the last MSDP was triggered.

2.4.2.3 Clear commands

database

Syntax

database [**interface** *ip-int-name* | *ip-address*] **group** *group-ip-address* [**source** *src-ip-address*]
database grp-interface *interface-name* [**fwd-service** *service-id*]
database [**interface** *ip-int-name* | *ip-address*] **group** *grp-ip-address* **source** *src-ip-address*
database host [*ip-address*]
database interface *ip-int-name* | *ip-address* [**group** *grp-ip-address*] [**source** *src-ip-address*]

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 *group-ip-address*

Clears the multicast group address (ipv4 or ipv6) or zero in the specified address group.

source *ip-address*

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

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 [**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 group or source cannot be specified at the same time.

Parameters

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, the (S,G) statistics are reset to zero.

source *ip-address*

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

family

Clears IPv4 information.

version

Syntax

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

ip-int-name

Clears IGMP or PIM statistics on the specified interface.

ip-address

Clears IGMP or PIM statistics 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.

s-pmsi

Syntax

s-pmsi [*mdSrcAddr*] [*mdGrpAddr*] [*vprnSrcAddr vprnGrpAddr*]

Context

clear>router>pim

Platforms

7210 SAS-T (network mode only), 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Mxp, and 7210 SAS-Sx/S 1/10GE (in standalone mode)

Description

This command clears the PIM selective provider multicast service interface cache.

Parameters

mdSrcAddr

Clears the specified source address used for Multicast Distribution Tree (MDT).

mdGrpAddr

Clears the specified group address used for Multicast Distribution Tree (MDT).

vprnSrcAddr

Clears the specified source address of the multicast sender.

vprnGrpAddr

Clears the specified multicast group address.

msdp

Syntax

msdp

Context

clear>router

Platforms

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

Description

Commands in this context clear and reset MSDP entities and statistics.

cache

Syntax

cache [**peer** *ip-address*] [**group** *ip-address*] [**source** *ip-address*] [**originrp** *ip-address*]

Context

clear>router>msdp

Platforms

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

Description

This command clears the MSDP cache.

Parameters

peer *ip-address*

Clears the cache of the IP address of the peer to which MSDP SA requests for groups matching this entry's group range were sent.

group *ip-address*

Clears the group IP address of the SA entry.

source *ip-address*

Clears the source IP address of the SA entry.

originrp *ip-address*

Clears the origin RP address type of the SA entry.

statistics

Syntax

statistics [**peer** *ip-address*]

Context

clear>router>msdp

Platforms

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

Description

This command clears IP address statistics of the peer to which MSDP SA requests for groups matching this entry's group range were sent.

Parameters

ip-address
Clears the statistics of the specified IP address.

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 from 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

Specifies 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

Specifies 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 from 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

Specifies 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

Specifies 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

2.4.2.4 Debug commands

2.4.2.4.1 Debug IGMP commands

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 specified interface name or IP address.

Parameters

ip-int-name

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

ip-address

Displays the information associated with only 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 debugging.

Output

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 and 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 only the specified IP interface name.

ip-address

Displays the information associated with only the specified IP address.

2.4.2.4.2 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 and 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 and 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 on 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 and 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 on 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 and 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 PIM bootstrap mechanism.

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

Parameters

detail

Debugs detailed information on the PIM bootstrap 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 and 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 and disables debugging for PIM database.

Parameters

group *grp-ip-address*

Debugs information associated with the specified database.

Values multicast group address (ipv4 or ipv6) or zero

source *ip-address*

Debugs information associated with the specified database.

Values source address (ipv4 or ipv6)

detail

Debugs detailed IP database information.

dynmldp

Syntax

dynmldp [detail]

no dynmldp

Context

debug>router>pim

Platforms

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

Description

This command enables and disables debugging for dynamic MLDP.

Parameters

detail

Debugs detailed dynamic MLDP 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 and disables debugging for 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 and disables debugging for 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.

mrib

Syntax

mrib [**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 and 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 and disables debugging for PIM messaging.

Parameters

detail

Debugs detailed messaging information.

mvpn-rtcache

Syntax

mvpn-rtcache [**group** *grp-ip-address*] [**peer** *ip-address*]

no mvpn-rtcache

Context

debug>router>pim

Platforms

7210 SAS-T (network mode only), 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Mxp, and 7210 SAS-Sx/S
1/10GE (in standalone mode)

Description

This command enables and disables debugging for the PIM MVPN route cache.

Parameters

grp-ip-address

Debugs information associated with the specified group.

Values multicast group address (ipv4) or zero

peer-ip-address

Debugs information associated with the specified peer.

Values peer address (ipv4)

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 and 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 and 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 and disables debugging for 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 the disables debugging for PIM RTM.

Parameters

detail

Debugs detailed RTM information.

s-pmsi

Syntax

s-pmsi [{*vpnSrcAddr* [*vpnGrpAddr*]} [*mdSrcAddr*]] [**detail**]

no s-pmsi

Context

debug>router>pim

Platforms

7210 SAS-T (network mode only), 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Mxp, and 7210 SAS-Sx/S
1/10GE (in standalone mode)

Description

This command enables debugging for PIM selective provider multicast service interface.

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

Parameters

vpnSrcAddr

Specifies the VPN source address.

vpnGrpAddr

Specifies the VPN group address

mdSrcAddr

Specifies the source address of the multicast domain.

detail

Displays detailed information for selective PMSI.

tunnel-interface

Syntax

```
tunnel-interface [rsvp-p2mp lsp-name] [sender ip-address] [detail]  
tunnel-interface [ldp-p2mp p2mp-id] [sender ip-address] [detail]  
no tunnel-interface [rsvp-p2mp lsp-name] [sender ip-address]  
no tunnel-interface [ldp-p2mp p2mp-id] [sender ip-address]
```

Context

debug>router>pim

Platforms

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

Description

This command enables and disables the debugging for PIM tunnel interfaces.

Parameters

lsp-name

Specifies the LSP for RSVP P2MP.

ip-address

Specifies the IP address of the sender.

p2mp-id

Specifies the P2MP ID for LDP P2MP.

detail

Displays detailed information for PIM tunnel interfaces.

2.4.2.4.3 Debug MSDP commands

msdp

Syntax

```
[no] msdp
```

Context

debug>router

Platforms

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

Description

This command enables debugging for MSDP.

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

packet

Syntax

packet [*pkt-type*] [*peer ip-address*]

Context

debug>router>msdp

Platforms

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

Description

This command enables debugging for MSDP packets.

Parameters

pkt-type

Debugs information associated with the specified packet type.

Values keep-alive, source-active, sa-request, sa-response

ip-address

Debugs information associated with the specified peer IP address.

pim

Syntax

pim [*grp-address*]

no pim

Context

debug>router>msdp

Platforms

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

Description

This command enables debugging for MSDP PIM.

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

Parameters

grp-address

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

```
rtm
```

Syntax

```
rtm [rp-address]
```

```
no rtm
```

Context

```
debug>router>msdp
```

Platforms

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

Description

This command enables debugging for the MSDP route table manager (RTM).

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

Parameters

rp-address

Debugs the IP multicast address for which this entry contains information.

```
sa-db
```

Syntax

```
sa-db [group grpAddr] [source srcAddr] [rp rpAddr]
```

```
no sadb
```

Context

```
debug>router>msdp
```

Platforms

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

Description

This command enables debugging for MSDP SA requests.

The **no** form of this command disables the MSDP SA database debugging.

Parameters

grpAddr

Debugs the IP address of the group.

srcAddr

Debugs the source IP address.

rpAddr

Debugs the specified RP address.

3 RIP



Note:

Routing Information Protocol (RIP) is only supported on the 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone).

This chapter provides information about configuring RIP parameters.

3.1 RIP overview

The RIP is an interior gateway protocol (IGP) that uses a distance-vector algorithm to determine the best route to a destination, using hop count as the metric. In order for the protocol to provide complete information About routing, every router in the domain must participate in the protocol.

RIP is a routing protocol based on a distance vector (Bellman-Ford) algorithm, which advertises network reachability by advertising prefix/mask and the metric (also known as hop count or cost). RIP selects the route with the lowest metric as the best route. RIP differs from link-state database protocols, such as OSPF and IS-IS, in that RIP advertises reachability information directly and link-state-database-based protocols advertise topology information. Each node is responsible for calculating the reachability information from the topology.

The router software supports RIPv1 and RIPv2. RIPv1, specified in RFC 1058, was written and implemented before the introduction of CIDR. It assumes the netmask information for non-local routes, based on the class the route belongs to:

- class A - 8-bit mask
- class B - 16-bit mask
- class C - 24-bit mask

RIPv2 was written after CIDR was developed and transmits netmask information with every route. Because of the support for CIDR routes and other enhancements in RIPv2 such as triggered updates, multicast advertisements, and authentication, most production networks use RIPv2. However, there are some older systems (hosts and routers) that only support RIPv1, especially when RIP is used simply to advertise default routing information.

The following table lists the interfaces and platforms RIP is supported on.

Table 32: RIP support on interfaces and platforms

Interface	Platforms supported
IES IP interface IPv4	7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone)
R-VPLS with IES IP interface IPv4	7210 SAS-Sx/S 1/10GE (standalone)
Network port IP interface IPv4	7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone)
VPRN IP interface	Not supported

Interface	Platforms supported
RVPLS with VPRN IP interface	Not supported

3.1.1 RIP features

RIP, a UDP-based protocol, updates its neighbors, and the neighbors update their neighbors, and so on. Each host that uses RIP has a routing process that sends and receives datagrams on UDP port number 520.

Each RIP router advertises all RIP routes periodically via RIP updates. Each update can contain a maximum of 25 route advertisements. This limit is imposed by RIP specifications. RIP can sometimes be configured to send as many as 255 routes per update. The formats of the RIPv1 and RIPv2 updates are slightly different and are shown as follows. Additionally, RIPv1 updates are sent to a broadcast address, RIPv2 updates can be either sent to a broadcast or multicast address (224.0.0.9). RIPv2 supports subnet masks, a feature that was not available in RIPv1.

A network address of 0.0.0.0 is considered a default route. A default route is used when it is not convenient to list every possible network in the RIP updates, and when one or more closely-connected gateways in the system are prepared to handle traffic to the networks that are not listed explicitly. These gateways create RIP entries for the address 0.0.0.0, as if it were a network to which they are connected.

3.1.1.1 RIP version types

The router allows you to specify the RIP version that is sent to RIP neighbors and RIP updates that is accepted and processed. The router allows the following combinations:

- Send only RIPv1 or send only RIPv2 to either the broadcast or multicast address or send no messages.

The default sends RIPv2 formatted messages to the broadcast address.

- Receive only RIPv1, receive only RIPv2, or receive both RIPv1 and RIPv2, or receive none.

The default receives both.

3.1.1.2 RIPv2 authentication

RIPv2 messages carry more information, which allows the use of a simple authentication mechanism to secure table updates. The router implementation enables the use of a simple password (plain text) or message digest (MD5) authentication.

3.1.1.3 Metrics

By default, RIP advertises all RIP routes to each peer every 30 seconds. RIP uses a hop count metric to determine the distance between the packet source and destination. The metric/cost values for a valid route is 1 through 15. A metric value of 16 (infinity) indicates that the route is no longer valid and should be removed from the router routing table.

Each router along the path increments the hop count value by 1. When a router receives a routing update with new or different destination information, the metric increments by 1.

The maximum number of hops in a path is 15. If a router receives a routing update with a metric of 15 and contains a new or modified entry, increasing the metric value by 1, it causes the metric to increment to 16 (infinity). Then, the destination is considered unreachable.

The router implementation of RIP uses split horizon with poison reverse to protect from problems such as “counting to infinity”. Split horizon with poison reverse means that routes learned from a neighbor through a specific interface are advertised in updates out of the same interface but with a metric of 16 (infinity).



Note: Ensure that routes imported into RIP from other protocols such as OSPF, IS-IS, BGP, and so on, set the acceptable metric to a value lower than 15. In other words, routes with metric greater than 15 are rejected by the RIP import policies.

3.1.1.4 Timers

RIP uses numerous timers to determine how often RIP updates are sent and how long routes are maintained.

- **Update**

Times the interval between periodic routing updates.

- **Timeout**

This timer is initialized when a route is established and any time an update message is received for the route. When this timer expires, the route is no longer valid. It is retained in the table for a short time, so that neighbors can be notified that the route has been dropped.

- **Flush**

When the flush timer expires, the route is removed from the tables.

3.1.1.5 Import and export policies


Routing policies can control the content of the routing tables, the routes that are advertised and the best route to take to reach a destination. Import route policies determine which routes are accepted from RIP neighbors. Export route policies determine which routes are exported from the route table to RIP. By default, RIP does not export routes it has learned to its neighbors.

There are no default routing policies. A policy must be created explicitly and applied to a RIP import or export command.

3.1.1.6 RIP packet format

The following figure shows the RIP packet format.

Figure 2: RIP packet format




A RIP packet consists of the following fields:

- **Command**
Indicates whether the packet is a request or a response message. The request asks the responding system to send all or part of its routing table. The response may be sent in response to a request, or it may be an unsolicited routing update generated by the sender.
- **Version**
The RIP version used. This field can signal different potentially incompatible versions.
- **Must be zero**
Not used in RIPv1. This field provides backward compatibility with pre-standard varieties of RIP. The default value is zero.
- **Address family identifier (AFI)**
The AFI is the type of address. RIP can carry routing information for several different protocols. Each entry in this field has an AFI to indicate the type of address being specified. The IP AFI is 2.
- **Address**
The IP address for the packet.
- **Metric**
Specifies the number of hops to the destination.
- **Mask**
Specifies the IP address mask.
- **Next hop**
Specifies the IP address of the next router along the path to the destination.

3.1.1.6.1 RIPv1 format

There can be between 1 and 25 (inclusive) RIP entries. The following figure shows RIPv1 format.


Figure 3: RIPv1 format



3.1.1.6.2 RIPv2 format

The following figure shows RIPv2 packet format.

Figure 4: RIPv2 format



The RIPv2 packets include the following fields:

- **Subnet mask**
The subnet mask for the entry. If this field is zero, no subnet mask has been specified for the entry.
- **Next hop**
The IP address of the next hop to forward packets.

3.1.2 Hierarchical levels

The minimum RIP configuration must define one group and one neighbor. The parameters configured on the global level are inherited by the group and neighbor levels. Parameters can be modified and overridden on a level-specific basis. RIP command hierarchy consists of three levels:

- global
- group
- neighbor


Many of the hierarchical RIP commands can be modified on different levels. The most specific value is used. That is, a RIP group-specific command takes precedence over a global RIP command. A neighbor-

specific statement takes precedence over a global RIP and group-specific command; for example, if you modify a RIP neighbor-level command default, the new value takes precedence over group- and global-level settings.

3.2 RIP configuration process overview

The following figure shows the process to configure RIP parameters.

Figure 5: RIP configuration and implementation flow



3.3 Configuration notes

This section describes RIP configuration caveats.

3.3.1 General

- Before RIP neighbor parameters can be configured, router interfaces must be configured.
- RIP must be explicitly created for each router interface. There are no default RIP instances on a router.

3.4 Configuring RIP with CLI

This section provides information to configure Routing Information Protocol (RIP) using the command line interface.

3.5 RIP configuration overview

3.5.1 Preconfiguration requirements

Before beginning the RIP configuration, it is optional to define policy statements in the **config>router>policy-options** context.

3.5.2 RIP hierarchy

RIP is configured in the **config>router>rip** context. RIP is not enabled by default. Three hierarchical levels are included in RIP configurations:

- global
- group
- neighbor

Commands and parameters configured on the global level are inherited by the group and neighbor levels although parameters configured on the group and neighbor levels take precedence over global configurations.

3.6 Basic RIP configuration

This section provides information to configure RIP and examples of common configuration tasks. For a router to accept RIP updates, in the **config>router>rip** context, you must define at least one group and one neighbor. A router ignores updates received from routers on interfaces not configured for RIP. Configuring other RIP commands and parameters are optional.

By default, the local router imports all routes from this neighbor and does not advertise routes. The router receives both RIPv1 and RIPv2 update messages with 25 to 255 route entries per message.

The RIP configuration commands have three primary configuration levels: **rip** for global configurations, group **group-name** for RIP group configurations, and **neighbor ip-int-name** for RIP neighbor configurations. Within the different levels, the configuration commands are identical. 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 RIP global settings.

The minimal RIP parameters that need to be configured in the **config>router>rip** context are:

- group
- neighbor

Example: Basic RIP configuration output

```
ALA-A>config>router>rip# info
-----
group "RIP-ALA-A"
        neighbor "to-ALA-4"
        exit
exit
-----
```

```
ALA-A>config>router>rip#
```

3.7 Common configuration tasks

About this task

This section provides a brief overview of the tasks that must be performed to configure RIP and provides the CLI commands.

Configure RIP hierarchically using the global level (applies to all peers), the group level (applies to all peers in peer-group), or the neighbor level (only applies to the specified interface). By default, group members inherit the group configuration parameters although a parameter can be modified on a per-member basis without affecting the group-level parameters.

Many of the hierarchical RIP commands can be used on different levels. The most specific value is used. That is, a RIP group-specific command takes precedence over a global RIP command. A neighbor-specific statement takes precedence over a global RIP or group-specific command.

All RIP instances must be explicitly created on each device. When created, RIP is administratively enabled.

To configure RIP, perform the following tasks:

Procedure

- Step 1.** Configure interfaces.
- Step 2.** Configure policy statements.
- Step 3.** Enable RIP.
- Step 4.** Configure group parameters.
- Step 5.** Configure neighbor parameters.

3.7.1 Configuring interfaces

The following command sequences create a logical IP interface. The logical interface can associate attributes like an IP address, port, Link Aggregation Group (LAG), or the system. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide* for more information about configuring interfaces.

Use the following syntax to configure a network interface.

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

Example: Command usage to configure a router interface

```
config>router> interface "to-ALA-4"
config>router>if$ address 10.10.12.1/24
config>router>if# port 1/1/1
config>router>if# exit
```

Example: IP configuration output showing the interface information

```
ALA-3>config>router# info
```

```
#-----  
echo "IP Configuration "  
#-----  
    interface "system"  
        address 10.10.10.103/32  
    exit  
    interface "to-ALA-4"  
        address 10.10.12.1/24  
        port 1/1/1  
    exit  
#-----  
ALA-3>config>router#
```

3.7.2 Configuring a route policy

The import route policy command allows you to filter routes being imported by the local router from its neighbors. If no match is found, the local router does not import any routes.

The export route policy command allows you to determine which routes are exported from the route table to RIP. By default, RIP does not export routes it has learned to its neighbors. If no export policy is specified, non-RIP routes are not exported from the routing table manager to RIP.

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.

This section only provides brief instructions to configure route policies. For more information, see the [Route policies](#) chapter.

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 and enables changes made to route policies during a session.
- The **abort** command discards changes that have been made to route policies during a session.

Use the following syntax to configure a policy to use for the RIP global, group, and neighbor **import** and **export** commands.

```
config>router>policy-options  
    begin  
    commit  
    abort  
    policy-statement name  
    description text  
    default-action {accept|reject}  
    entry entry-id  
        description text  
        action {accept|reject}  
        from  
        to
```

Use the following syntax to enter the edit mode.

```
config>router> policy-options  
    begin
```

Example

The following shows the command usage to configure a policy statement. Policy option commands are configured in the **config>router** context. Use the **commit** command to save the changes.

```
config>router>policy-options# begin
  policy-options# policy-statement "RIP-policy"
  policy-options>policy-statement$ description "this is a test RIP policy"
  policy-options>policy-statement>default# entry 1
  policy-options>policy-statement>entry$ action accept
  policy-options>policy-statement>entry# exit
  policy-options>policy-statement# default-action reject
  policy-options>policy-statement# exit
policy-options# commit
```

```
ALA-A>config>router>policy-options# info
-----
policy-statement "RIP-policy"
description "this is a test RIP policy"
entry 1
action accept
exit
exit
default-action reject
exit
-----
ALA-A>config>router>policy-options>policy-statement#
```

```
policy-statement "direct_to_RIP"
  description "Policy From direct To rip"
  entry 10
    description "Entry 10 - From Prot. direct To rip"
    from
      protocol direct
    exit
    to
      protocol rip
    exit
    action accept
    exit
  exit
  entry 20
    from
      protocol ospf
    exit
    to
      protocol rip
    exit
    action accept
    metric subtract 100
    exit
  exit
  default-action reject
exit
```

3.7.3 Configuring RIP parameters

Use the following syntax to configure the RIP parameters.

```
config>router
  rip
  authentication-key [authentication-key|hash-key [hash|hash2]
  authentication-type {none | password | message-digest | message-digest-20}
  check-zero {enable | disable}
  description string
  export policy-name [policy-name ...up to 5 max]
  import policy-name [policy-name ...up to 5 max]
  message-size number
  metric-in metric
  metric-out metric
  preference number
  receive {both | none | version-1 | version-2}
  send {broadcast | multicast | none | version-1 | both}
  no shutdown
  split-horizon {enable | disable}
  timers update timeout flush

  group group-name
    authentication-key [authentication-key|hash-key [hash|hash2]
    authentication-type {none|password|message-digest| message-digest-20}
    check-zero {enable|disable}
    description string
    export policy-name [policy-name ...up to 5 max]]
    import policy-name [policy-name ...up to 5 max]]
    message-size number
    metric-in metric
    metric-out metric
    preference number
    receive {both|none|version-1|version-2}
    send {broadcast|multicast|none|version-1}
    no shutdown
    split-horizon {enable|disable}
    timers update timeout flush

  neighbor ip-int-name
    authentication-key [authentication-key|hash-key [hash|hash2]
    authentication-type {none|password|message-digest| message-digest-20}
    check-zero {enable|disable}
    description string
    export policy-name [policy-name ...up to 5 max]]
    import policy-name [policy-name ...up to 5 max]]
    message-size number
    metric-in metric
    metric-out metric
    preference number
    receive {both|none|version-1|version-2}
    send {broadcast|multicast|none|version-1}
    split-horizon {enable|disable}
    timers update timeout flush
    no shutdown
```

3.7.4 Configuring global-level parameters

When the RIP protocol instance is created, the **no shutdown** command is not required because RIP is administratively enabled upon creation. Minimally, to enable RIP on a router, at least one group and one neighbor must be configured. There are no default groups or neighbors. Each group and neighbor must be explicitly configured.

**Note:**

Careful planning is essential to implement commands that can affect the behavior of global, group, and neighbor-levels. Because the RIP commands are hierarchical, analyze the values that can disable features on a particular level. Use the following CLI syntax to configure global-level RIP parameters:

```
config>router
  rip
  authentication-key [authentication-key|hash-key [hash|hash2]
  authentication-type {password | message-digest}
  check-zero {enable|disable}
  export policy-name [policy-name ...up to 5 max]
  import policy-name [policy-name ...up to 5 max]
  message-size number
  metric-in metric
  metric-out metric
  preference number
  receive {both|none|version-1|version-2}
  send {broadcast | multicast | none | version-1| both}
  no shutdown
  split-horizon {enable | disable}
  timers update timeout flush
```

Example: Command usage to configure the global RIP

```
config>router# rip
config>router>rip# authentication-type password
config>router>rip# authentication-key test123
config>router>rip# receive both
config>router>rip# split-horizon enable
config>router>rip# timers 300 600 600
config>router>rip>group# exit
```

Example: RIP group configuration output

```
ALA-A>config>router>rip# info
-----
      authentication-type simple
      authentication-key "ac1865lvz1d" hash
      timers 300 600 600
-----
ALA-A>config>router>rip#
```

3.7.5 Configuring group-level parameters

A group is a collection of related RIP 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 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. Use the following syntax to configure a group.

```
config>router# rip
  group group-name
    authentication-key [authentication-key|hash-key [hash|hash2]]
    authentication-type {password|message-digest}
    check-zero {enable|disable}
    description string
    export policy-name [policy-name ...]
    import policy-name [policy-name ...]
    message-size number
    metric-in metric
    metric-out metric
    preference number
    receive {both|none|version-1|version-2}
    send {broadcast|multicast|none|version-1|both}
    no shutdown
    split-horizon {enable|disable}
    timers update timeout flush
```

Example: Command usage to configure a display group

```
config>router# rip
config>router>rip# group headquarters
config>router>rip>group$ description "Mt. View"
config>router>rip>group# no shutdown
```

Example: RIP group configuration output

```
ALA-A>config>router>rip# info
-----
authentication-type simple
authentication-key "acl865lvzld" hash
timers 300 600 600
group "headquarters"
  description "Mt. View"
exit
-----
ALA-A>config>router>rip#
```

3.7.6 Configuring neighbor-level parameters

After you create a group name and assign options, add neighbor interfaces within the same group. 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.

Use the following syntax to add a neighbor to a group and define options that override the same group-level command value.

```
config>router# rip
  group group-name
    neighbor ip-int-name
      authentication-key [authentication-key|hash-key [hash|hash2]
      authentication-type {password|message-digest}
      check-zero {enable|disable}
      description string
      export policy-name [policy-name ...]
      import policy-name [policy-name ...]
      message-size number
      metric-in metric
      metric-out metric
      preference number
      receive {both|none|version-1|version-2}
      send {broadcast|multicast|none|version-1}
      split-horizon {enable|disable}
      timers update timeout flush
      no shutdown
```

Example: Command usage to configure a display neighbor

```
config>router# rip
  config>router>rip# group headquarters1
  config>router>rip>group# neighbor ferguson-274
  config>router>rip>group>neighbor$ preference 255
  config>router>rip>group>neighbor# send both
  config>router>rip>group>neighbor# split-horizon enable
  config>router>rip>group>neighbor# message-size 255
```

Example: Output of the neighbor configured in group "headquarters"

```
ALA-A>config>router>rip>group>neighbor# info
-----
                message-size 255
                preference 255
                split-horizon enable
                no timers
-----
ALA-A>config>router>rip>group>neighbor#
```

3.8 RIP configuration management tasks

The following section describes the syntax used to configure the RIP configuration management tasks.

3.8.1 Modifying RIP parameters

Modify, add or remove RIP parameters in the CLI. The changes are applied immediately. For the complete list of CLI commands, see [Configuring RIP parameters](#).

```
config>router# rip
  group group-name
  . . .
```



```
neighbor ip-int-name  
. . .
```

Example

```
config>router>rip# group "headquarters"  
config>router>rip>group# neighbor "ferguson-274"  
config>router>rip>group>neighbor# import RIPpolicy  
config>router>rip>group>neighbor# message-size 150
```

Example: Updated parameters

```
ALA-A>config>router>rip# info  
-----  
authentication-type simple  
authentication-key "ac1865lvz1d" hash  
timers 300 600 600  
group "headquarters"  
  description "Mt. View"  
  neighbor "ferguson-274"  
    import "RIPpolicy"  
    message-size 150  
    preference 255  
    split-horizon enable  
    no timers  
  exit  
exit  
-----  
ALA-A>config>router>rip#
```

3.8.2 Deleting a group

A group must be shut down first to delete it.

Use the following syntax to shut down and then delete a group.

```
config>router# rip  
[no] group group-name  
shutdown
```

Example

```
config>router# rip  
config>router>rip# group "RIP-ALA-3"  
config>router>rip>group# shutdown  
config>router>rip>group# exit  
config>router>rip# no group "RIP-ALA-33"
```

Example

If you try to delete the group without shutting it down first, the following message appears:

```
INFO: RIP #1204 group should be administratively down -  
virtual router index 1,group RIP-ALA-4
```

3.8.3 Deleting a neighbor

The neighbor must be shut down before it can be deleted.

Use the following syntax to delete a neighbor:

```
config>router# rip
[no] group group-name
[no] neighbor ip-int-name
shutdown
```

Example

```
config>router# rip
config>router>rip# group "RIP-ALA-4"
config>router>rip>group# neighbor "to-ALA-3"
config>router>rip>group>neighbor# shutdown
config>router>rip>group>neighbor# exit
config>router>rip>group# no neighbor "to-ALA-3"
```

Example

If you try to delete the neighbor before it is shut down, the following message appears:

```
INFO: RIP #1101 neighbor should be administratively down - virtual router index
```

3.9 RIP command reference

- [Command hierarchies](#)
- [Command descriptions](#)

3.9.1 Command hierarchies

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

3.9.1.1 Configuration commands

```
config
- router router-name
  - [no] rip
    - authentication-key [authentication-key | hash-key] [hash | hash2]
    - no authentication-key
    - authentication-type {none | password | message-digest | message-digest-20}
    - no authentication-type
    - check-zero {enable | disable}
    - no check-zero
```

```
- description description-string
- no description
- export policy-name [policy-name ... (up to 5 max)]
- no export
- export-limit number [log percentage]
- no export-limit
- import policy-name [policy-name ... (up to 5 max)]
- no import
- message-size max-num-of-routes
- no message-size
- metric-in metric
- no metric-in
- metric-out metric
- no metric-out
- preference preference
- no preference
- receive receive-type
- no receive
- send send-type
- no send
- [no] shutdown
- split-horizon {enable | disable}
- no split-horizon
- timers update timeout flush
- no timers
```

3.9.1.1.1 Group commands

```
config
- router router-name
- [no] rip
- [no] group group-name
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- authentication-type {none | password | message-digest | message-digest-20}
- no authentication-type
- check-zero {enable | disable}
- no check-zero
- description description-string
- no description
- export policy-name [policy-name ... (up to 5 max)]
- no export
- import policy-name [policy-name ... (up to 5 max)]
- no import
- message-size max-num-of-routes
- no message-size
- metric-in metric
- no metric-in
- metric-out metric
- no metric-out
- preference preference
- no preference
- receive receive-type
- no receive
- send send-type
- no send
- [no] shutdown
- split-horizon {enable | disable}
- no split-horizon
- timers update timeout flush
```

```
- no timers
```

3.9.1.1.2 Neighbor commands

```
config
- router router-name
- [no] rip
- [no] group group-name
- [no] neighbor ip-int-name
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- authentication-type {none | password | message-digest}
- no authentication-type
- check-zero {enable | disable}
- no check-zero
- description description-string
- no description
- export policy-name [policy-name ...(up to 5 max)]
- no export
- import policy-name [policy-name ...(up to 5 max)]
- no import
- message-size max-num-of-routes
- no message-size
- metric-in metric
- no metric-in
- metric-out metric
- no metric-out
- preference preference
- no preference
- receive receive-type
- no receive
- send send-type
- no send
- [no] shutdown
- split-horizon {enable | disable}
- no split-horizon
- timers update timeout flush
- no timers
```

3.9.1.2 Show RIP commands

```
show
- router
- rip
- database [ip-prefix [/mask] [longer] [peer ip-address] [detail [qos]]]
- group [name] [detail]
- neighbors [ip-int-name | ip-addr] [detail] [advertised-routes]
- peer [interface-name]
- statistics [ip-int-name | ip-addr]
```

3.9.1.3 Clear RIP commands

```
clear
- router
- rip
```

```
- database
- export
- statistics [neighbor ip-int-name | ip-address]
```

3.9.1.4 Debug RIP commands

```
debug
- router
  - rip
    - [no] auth [neighbor ip-int-name | ip-address]
    - [no] error [neighbor ip-int-name | ip-address]
    - [no] events [neighbor ip-int-name | ip-address]
    - [no] holddown [neighbor ip-int-name | ip-address]
    - [no] packets [neighbor ip-int-name | ip-address]
    - [no] request [neighbor ip-int-name | ip-address]
    - [no] trigger [neighbor ip-int-name | ip-address]
    - [no] updates [neighbor ip-int-name | ip-address]
```

3.9.2 Command descriptions

- [RIP configuration commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug RIP commands](#)

3.9.2.1 RIP configuration commands

3.9.2.1.1 Generic commands

description

Syntax

description *string*

no description

Context

config>router>rip>group

config>router>rip>group>neighbor

Platforms

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

Description

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

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

Parameters

string

Specifies the description character string. Allowed values are any string up to 80 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax

[no] shutdown

Context

config>router>rip

config>router>rip>group

config>router>rip>group>neighbor

Platforms

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

Description

This command administratively disables an entity. Shutting down an entity does not change, reset, or remove any configuration settings or statistics. Many objects must be shut down before they may be deleted.

The **shutdown** command administratively shuts down an entity. Administratively shutting down an entity changes the operational state of the entity to down and the operational state of any entities contained within the administratively down entity.

Unlike other commands and parameters where the default state 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

RIP Global

In the **config>router>rip** context, the **shutdown** command administratively enables or disables the RIP protocol instance. If RIP is globally shut down, all RIP group and neighbor interfaces transition to the operationally down state. Routes learned from a neighbor that is shut down are immediately removed from the RIP database and route table manager (RTM). A RIP protocol instance is administratively enabled by default.

RIP Group

In the **config>router>rip>group** context, the **shutdown** command administratively enables or disables the RIP group. If a RIP group is shut down, all member neighbor interfaces transition to the operationally down state. Routes learned from a neighbor that is shut down are immediately removed from the RIP database and route table manager (RTM). A RIP group is administratively enabled by default.

RIP Neighbor

In the **config>router>rip>group>neighbor** context, the **shutdown** command administratively enables or disables the RIP neighbor interface. If a RIP neighbor is shut down, the neighbor interface transitions to the operationally down state. Routes learned from a neighbor that is shut down are immediately removed from the RIP database and route table manager (RTM). A RIP neighbor interface is administratively enabled by default.

3.9.2.1.2 RIP commands

rip

Syntax

[no] rip

Context

config>router

Platforms

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

Description

Commands in this context configure the RIP protocol instance.

When a RIP instance is created, the protocol is enabled by default. To start or suspend execution of the RIP protocol without affecting the configuration, use the **[no] shutdown** command.

The **no** form of this command deletes the RIP protocol instance removing all associated configuration parameters.

Default

no rip

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

```
config>router>rip  
config>router>rip>group  
config>router>rip>group>neighbor
```

Platforms

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

Description

This command sets the authentication password that is passed between RIP neighbors.

The authentication type and authentication key must match exactly for the RIP message to be considered authentic and processed.

The **no** form of this command removes the authentication password from the configuration and disables authentication.

Default

no authentication-key

Parameters

authentication-key

Specifies the authentication key. Allowed values are any string up to 16 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 33 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

Keyword to specify that the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Keyword to specify that the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

authentication-type

Syntax

authentication-type {none | password | message-digest | message-digest-20}

no authentication-type

Context

```
config>router>rip  
config>router>rip>group  
config>router>rip>group>neighbor
```

Platforms

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

Description

This command sets the type of authentication that is used between RIP neighbors.

The type and password must match exactly for the RIP message to be considered authentic and processed.

The **no** form of this command removes the authentication type from the configuration and effectively disables authentication.

Default

no authentication-type

Parameters

none

Keyword that explicitly disables authentication at a specific level (global, group, neighbor). If the command does not exist in the configuration, the parameter is inherited.

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 configure 16-byte message digest for MD5 authentication. If this option is configured, at least one message-digest key must be configured.

message-digest-20

Keyword to configure 20-byte message digest for MD5 authentication in accordance with RFC 2082, *RIP-2 MD5 Authentication*. If this option is configured, at least one message-digest key must be configured.

check-zero

Syntax

```
check-zero {enable | disable}  
no check-zero
```

Context

```
config>router>rip
config>router>rip>group
config>router>rip>group>neighbor
```

Platforms

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

Description

This command enables checking for zero values in fields specified to be zero by the RIPv1 and RIPv2 specifications.

This command enables checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting of non-compliant RIP messages.

This command disables this check and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.

This command can be set at all RIP levels: global level (applies to all groups and neighbor interfaces), group level (applies to all neighbor interfaces in the group), or neighbor level (only applies to the specified neighbor interface). The most specific value is used. In particular, if no value is set (**no check-zero**), the setting from the less specific level is inherited by the lower level.

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

Special Cases

RIP Global

By default, check-zero is disabled at the global RIP instance level.

Parameters

enable

Keyword to reject RIP messages that do not have zero in the RIPv1 and RIPv2 mandatory fields.

disable

Keyword to allow the receipt of RIP messages that do not have the mandatory zero fields reset.

export

Syntax

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

no export

Context

```
config>router>rip
config>router>rip>group
```

```
config>router>rip>group>neighbor
```

Platforms

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

Description

This command specifies the export route policies used to determine which routes are exported to RIP.

If no export policy is specified, non-RIP routes are not exported from the routing table manager to RIP. RIP-learned routes are exported to RIP neighbors.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered overrides 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. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

The specified names must already be defined.

export-limit

Syntax

export-limit *number* [*log percentage*]

no export-limit

Context

```
config>router>rip
```

Platforms

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

Description

This command configures the maximum number of routes (prefixes) that can be exported into RIP from the route table.

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

Default

no export-limit

Parameters

number

Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.

Values 1 to 4294967295

log percentage

Specifies the percentage of the export-limit, at which a warning log message and SNMP notification are sent.

Values 1 to 100

group

Syntax

[no] **group** *group-name*

Context

config>router>rip

Platforms

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

Description

This command configures a RIP group of neighbor interfaces.

RIP groups are a way of logically associating RIP neighbor interfaces to facilitate a common configuration for RIP interfaces.

The **no** form of this command deletes the RIP neighbor interface group. Deleting the group also removes the RIP configuration of all the neighbor interfaces currently assigned to this group.

Default

no group

Parameters

group-name

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

import

Syntax

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

no import

Context

config>router>rip

config>router>rip>group

config>router>rip>group>neighbor

Platforms

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

Description

This command configures import route policies to determine which routes are accepted from RIP neighbors. If no import policy is specified, RIP accepts all routes from configured RIP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple import commands are issued, the last command entered overrides the previous command. A maximum of five policy names can be specified.

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

Default

no import

Parameters

policy-name

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

The specified names must already be defined.

message-size

Syntax

message-size *max-num-of-routes*

no message-size

Context

```
config>router>rip  
config>router>rip>group  
config>router>rip>group>neighbor
```

Platforms

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

Description

This command configures the maximum number of routes per RIP update message.
The **no** form of this command reverts to the default value.

Default

message-size 25

Parameters

max-num-of-routes

Specifies the maximum number of RIP routes per RIP update message, expressed as a decimal integer.

Values 25 to 255

metric-in

Syntax

```
metric-in metric  
no metric-in
```

Context

```
config>router>rip  
config>router>rip>group  
config>router>rip>group>neighbor
```

Platforms

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

Description

This command configures the metric added to routes received from a RIP neighbor.

When applying an export policy to a RIP configuration, the policy overrides the metric values determined through calculations involving the **metric-in** and **metric-out** values.

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

Default

metric-in 1

Parameters

metric

Specifies the value added to the metric of routes received from a RIP neighbor, expressed as a decimal integer.

Values 1 to 16

metric-out

Syntax

metric-out *metric*

no metric-out

Context

config>router>rip

config>router>rip>group

config>router>rip>group>neighbor

Platforms

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

Description

This command configures the metric assigned to routes exported into RIP and advertised to RIP neighbors.

When applying an export policy to a RIP configuration, the policy overrides the metric values determined through calculations involving the **metric-in** and **metric-out** values.

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

Default

metric-out 1

Parameters

metric

Specifies the value added to the metric for routes exported into RIP and advertised to RIP neighbors, expressed as a decimal integer.

Values 1 to 16

neighbor

Syntax

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

Context

config>router>rip>group

Platforms

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

Description

This command enables the context for configuring a RIP neighbor interface.

By default, interfaces are not activated in an interior gateway protocol, such as RIP, unless explicitly configured.

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

Default

no neighbor

Parameters

ip-int-name

Specifies the IP interface name. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters 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.

preference

Syntax

preference *preference*

no preference

Context

config>router>rip

config>router>rip>group


```
config>router>rip>group>neighbor
```

Platforms

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

Description

This command configures the preference for RIP routes.

A route can be learned by the router from different protocols, in which case the costs (metrics) are not comparable. When this occurs the preference is used to decide which route is used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is made according to the default preference table defined in [Table 33: Route preference defaults by route type](#) . If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of which route to use is determined by the configuration of the **ecmp** command in the **config>router** context.

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

Default

preference 100

Parameters

preference

Specifies the preference for RIP routes expressed as a decimal integer. The following table lists the defaults for different route types.

Table 33: Route preference defaults by route type

Route type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes

Values 0 to 255

receive

Syntax

receive {**both** | **none** | **version-1** | **version-2**}

no receive

Context

config>router>rip

config>router>rip>group

config>router>rip>group>neighbor

Platforms

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

Description

This command configures the types of RIP updates that are accepted and processed.

If **both** or **version-2** is specified, the RIP instance listens for and accepts packets sent to the broadcast and multicast (224.0.0.9) addresses.

If **version-1** is specified, the router only listens for and accept packets sent to the broadcast address.

This control can be issued at the global, group, or interface level. The default behavior is to accept and process both RIPv1 and RIPv2 messages.

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

Default

receive both

Parameters

both

Keyword to specify that RIP updates in either version 1 or version 2 format are accepted.

none

Keyword to specify that RIP updates are not accepted.

version-1

Keyword to specify that RIP updates in version 1 format only are accepted.

version-2

Keyword to specify that RIP updates in version 2 format only are accepted.

send

Syntax

send {**broadcast** | **multicast** | **none** | **version-1**}
no send

Context

config>router>rip
config>router>rip>group
config>router>rip>group>neighbor

Platforms

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

Description

This command specifies the type of RIP messages sent to RIP neighbors.

If **version-1** is specified, the router need only listen for and accept packets sent to the broadcast address.

This control can be issued at the global, group, or interface level.

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

Default

send broadcast

Parameters

broadcast

Keyword that sends RIPv2 formatted messages to the broadcast address.

multicast

Keyword that sends RIPv2 formatted messages to the multicast address.

none

Keyword that specifies not to send any RIP messages (that is, silent listener).

version-1

Keyword that sends RIPv1 formatted messages to the broadcast address.

split-horizon

Syntax

split-horizon {**enable** | **disable**}
no split-horizon

Context

```
config>router>rip
config>router>rip>group
config>router>rip>group>neighbor
```

Platforms

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

Description

This command enables the use of split-horizon.

RIP uses split-horizon with poison-reverse to avoid looping routes propagating through the network. Split-horizon with poison reverse means that routes learned from a neighbor through a specific interface are advertised in updates out of the same interface but with a metric of 16 (infinity).

The **split-horizon disable** command enables split horizon without poison reverse. This allows the routes to be readvertised on interfaces other than the interface that learned the route, with the advertised metric equaling an increment of the **metric-in** value.

This configuration parameter can be set at the following levels: global level (applies to all groups and neighbor interfaces), group level (applies to all neighbor interfaces in the group), or neighbor level (only applies to the specified neighbor interface). The most specific value is used. In particular, if no value is set (**no split-horizon**), the setting from the less specific level is inherited by the lower level.

The **no** form of this command disables split horizon, which allows the lower level to inherit the setting from an upper level.

Default

split-horizon enable

Parameters

enable

Keyword to enable split horizon and poison reverse.

disable

Keyword to disable split horizon, allowing routes to be readvertised on the same interface on which they were learned with the advertised metric incremented by the **metric-in** value.

timers

Syntax

```
timers update timeout flush
no timers
```

Context

```
config>router>rip
config>router>rip>group
```

```
config>router>rip>group>neighbor
```

Platforms

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

Description

This command configures values for the update, timeout, and flush RIP timers.

The RIP update timer determines how often RIP updates are sent.

If the route is not updated by the time the RIP timeout timer expires, the route is declared invalid but is maintained in the RIP database.

The RIP flush timer determines how long a route is maintained in the RIP database after it has been declared invalid. When the flush timer expires, the route is removed from the RIP database.

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

Default

timers 30 180 120

Parameters

update

Specifies the RIP update timer value in seconds, expressed as a decimal integer.

Values 1 to 600

timeout

Specifies the RIP timeout timer value in seconds, expressed as a decimal integer.

Values 1 to 1200

flush

Specifies the RIP flush timer value in seconds, expressed as a decimal integer.

Values 1 to 1200

3.9.2.2 Show commands

database

Syntax

database [*ip-prefix* [*/mask*] [**longer**] [**peer** *ip-address*] [**detail** [**qos**]]]

Context

show>router>rip

Platforms

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

Description

This command displays the routes in the RIP database.

Parameters

detail

Displays detailed RIP database information.

Output

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

Sample output

```
*A:dut-c>show>router>rip# database

=====
RIP Route Database
=====
Destination      Peer           Interface      Met  TTL  Valid
-----
10.33.33.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.34.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.35.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.36.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.37.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.38.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.39.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.40.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.41.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.42.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.43.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.44.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.45.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.46.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
10.33.47.0/24    10.1.4.2      to-ixia-Nw-4   1    172  Yes
=====
A:ALA-A#
```

Table 34: Output fields: RIP database

Label	Description
Destination	Displays the RIP destination for the route
Peer	Displays the router ID of the peer router
Interface	Displays the IP address of the interface
Metric	Displays the hop count to rate the value of different hops

Label	Description
TTL	Displays how many seconds the specific route remains in the routing table. When an entry reaches 0, it is removed from the routing table.
Valid	No — the route is not valid Yes — the route is valid

group

Syntax

group [*group-name*] [**detail**]

Context

show>router>rip

Platforms

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

Description

This command displays RIP group information.

Parameters

group-name

Displays RIP group information for the specified group.

detail

Displays detailed RIP group information.

Output

The following outputs are examples of RIP group information, and the corresponding tables describe the output fields.

- [Sample standard RIP group output, Table 35: Output fields: group](#)
- [Sample output — detailed, Table 36: Output fields: RIP group detail](#)

Sample standard RIP group output

```
A:ALA-A# show router rip group
=====
RIP Groups
=====
Group                Adm      Opr      Send Mode  Recv Mode  Metric In
-----
rip-group            Up       Down    BCast     Both       1
=====
```

A:ALA-A#

Table 35: Output fields: group

Label	Description
Group	Displays the RIP group name
Adm	Down — the RIP group is administratively down Up — the RIP group is administratively up
Opr	Down — the RIP group is operationally down Up — the RIP group is operationally up
Send Mode	Bcast — Specifies that RIPv2 formatted messages are sent to the broadcast address Mcast — Specifies that RIPv2 formatted messages are sent to the multicast address None — Specifies that no RIP messages are sent (that is, silent listener) RIPv1 — Specifies that RIPv1 formatted messages are sent to the broadcast address
Recv Mode	Both — Specifies that RIP updates in either version 1 or version 2 format are accepted None — Specifies that RIP updates are not accepted RIPv1 — Specifies that RIP updates in version 1 format only are accepted RIPv2 — Specifies that RIP updates in version 2 format only are accepted
Metric In	Displays the metric value added to routes received from a RIP neighbor

Sample output — detailed

```
*A:dut-c>show>router>rip# group detail

=====
RIP groups (Detail)
=====
-----
Group "test"
-----
Description      : No Description Available
Admin State      : Up
Oper State       : Down
Send Mode        : Broadcast
Receive Mode     : Both
Metric In        : 1
Metric Out       : 1
Split Horizon    : Enabled
Check Zero       : Disabled
Message Size     : 25
Preference       : 100
Auth. Type       : None
Update Timer     : 30
Timeout Timer    : 180
Flush Timer      : 120
```



```

Export Policies:
  direct_to_RIP
Import Policies:
  None
-----
Group "to-ixia"
-----
Description      : No Description Available
Admin State     : Up                               Oper State      : Up
Send Mode       : Broadcast                       Receive Mode    : Both
Metric In       : 1                               Metric Out      : 1
Split Horizon   : Enabled                         Check Zero     : Disabled
Message Size    : 25                             Preference     : 100
Auth. Type      : None                           Update Timer    : 30
Timeout Timer   : 180                            Flush Timer     : 120
Export Policies:
  direct_to_RIP
Import Policies:
  None
=====
*A:dut-c>show>router>rip#

```

Table 36: Output fields: RIP group detail

Label	Description
Description	Displays the RIP group description. No Description Available indicates no description is configured.
Admin State	Indicates whether the RIP group interface is administratively up or down
Oper State	Indicates whether the RIP group interface is operationally up or down
Send Mode	Bcast — specifies that RIPv2 formatted messages are sent to the broadcast address Mcast — specifies that RIPv2 formatted messages are sent to the multicast address None — specifies that no RIP messages are sent (silent listener) RIPv1 — specifies that RIPv1 formatted messages are sent to the broadcast address
Receive Mode	Both — specifies that RIP updates in either version 1 or version 2 format are accepted None — specifies that RIP updates are not accepted RIPv1 — specifies that RIP updates in version 1 format only are accepted RIPv2 — specifies that RIP updates in version 2 format only are accepted

Label	Description
Metric In	Displays the metric value added to routes received from a RIP neighbor
Metric Out	Displays the value added to routes exported into RIP and advertised to RIP neighbors
Split Horizon	Indicates whether split horizon and poison reverse is Enabled or Disabled for the RIP neighbor.
Check Zero	Disabled — the mandatory zero fields in RIP packets are not checked, allowing receipt of RIP messages even if mandatory zero fields are non-zero for the neighbor Enabled — mandatory zero fields in RIP packets are checked and non-compliant RIP messages are rejected
Message Size	Displays the maximum number of routes per RIP update message
Preference	Displays the preference of RIP routes from the neighbor
Auth. Type	Specifies the authentication type
Update Timer	Displays the current setting of the RIP update timer value expressed in seconds
Timeout Timer	Displays the current RIP timeout timer value expressed in seconds
Flush Timer	Displays the number of seconds after a route has been declared invalid that it is flushed from the route database
Export Policies	Displays the export route policy that is used to determine routes advertised to all peers
Import Policies	Displays the import route policy that is used to determine which routes are accepted from RIP neighbors

neighbors

Syntax

neighbors [*ip-addr* | *ip-int-name*] [**advertised-routes** | **detail**]

Context

show>router>rip

Platforms

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

Description

This command displays RIP neighbor interface information.

Parameters

ip-addr | ip-int-name

Displays information for the specified IP interface.

Default all neighbor interfaces

advertised-routes

Displays the routes advertised to RIP neighbors. If no neighbors are specified, all routes advertised to all neighbors are displayed. If a neighbor is specified, only routes advertised to the specific neighbor or interface are displayed.

Default Displays RIP information

Output

The following outputs are examples of RIP neighbor information, and the corresponding tables describe the output fields.

- [Sample output, Table 37: Output fields: neighbor standard](#)
- [Sample detailed output, Table 38: Output fields: neighbor detail](#)
- [Sample output with advertised routes](#)

Sample output

```
*A:dut-c>show>router>rip# neighbor

=====
RIP Neighbors
=====
Interface                Adm  Opr  Primary IP      Send  Recv  Metric
                        Mode Mode              Mode  Mode   In
-----
to-ixia-Nw-4              Up   Up   10.1.4.1        BCast Both    1
-----
No. of RIP Neighbors: 1
=====
*A:dut-c>show>router>rip#
```

Table 37: Output fields: neighbor standard

Label	Description
Neighbor	Displays the RIP neighbor interface name
Adm	Down — RIP neighbor interface is administratively down Up — RIP neighbor interface is administratively up
Opr	Down — RIP neighbor interface is operationally down Up — RIP neighbor interface is operationally up

Label	Description
Primary IP	Displays the primary IP address of the RIP neighbor interface
Send Mode	Bcast — Specifies that RIPv2 formatted messages are sent to the broadcast address Mcast — Specifies that RIPv2 formatted messages are sent to the multicast address None — Specifies that no RIP messages are sent (that is, silent listener) RIPv1 — Specifies that RIPv1 formatted messages are sent to the broadcast address
Recv Mode	Both — Specifies that RIP updates in either version 1 or version 2 format are accepted None — Specifies that RIP updates are not accepted RIPv1 — Specifies that RIP updates in version 1 format only are accepted RIPv2 — Specifies that RIP updates in version 2 format only are accepted
Metric In	Displays the metric added to routes received from a RIP neighbor

Sample detailed output

```
*A:dut-c>show>router>rip# neighbor detail

=====
RIP Neighbors (Detail)
=====
-----
Neighbor "to-ixia-Nw-4"
-----
Description      : No Description Available
Primary IP       : 10.1.4.1      Group       : to-ixia
Admin State      : Up           Oper State  : Up
Send Mode        : Broadcast    Receive Mode : Both
Metric In        : 1            Metric Out  : 1
Split Horizon    : Enabled      Check Zero  : Disabled
Message Size     : 25           Preference  : 100
Auth. Type       : None         Update Timer : 30
Timeout Timer    : 180          Flush Timer  : 120
Export Policies:
  direct_to_RIP
Import Policies:
  None
=====
*A:dut-c>show>router>rip#
```

Table 38: Output fields: neighbor detail

Label	Description
Neighbor	Displays the RIP neighbor name
Description	Displays the RIP neighbor description. No Description Available indicates no description is configured.
Primary IP	Displays the RIP neighbor interface primary IP address
Group	Displays the RIP group name of the neighbor interface
Admin State	Down — RIP neighbor interface is administratively down Up — RIP neighbor interface is administratively up
Oper State	Down — RIP neighbor interface is operationally down Up — RIP neighbor interface is operationally up
Send Mode	Bcast — Specifies that RIPv2 formatted messages are sent to the broadcast address Mcast — Specifies that RIPv2 formatted messages are sent to the multicast address None — Specifies that no RIP messages are sent (that is, silent listener) RIPv1 — Specifies that RIPv1 formatted messages are sent to the broadcast address
Recv Mode	Both — Specifies that RIP updates in either version 1 or version 2 format are accepted None — Specifies that RIP updates are not accepted RIPv1 — Specifies that RIP updates in version 1 format only are accepted RIPv2 — Specifies that RIP updates in version 2 format only are accepted
Metric In	Displays the metric value added to routes received from a RIP neighbor
Metric Out	Displays the value added to routes exported into RIP and advertised to RIP neighbors
Split Horizon	Disabled — split horizon disabled for the neighbor Enabled — split horizon and poison reverse enabled for the neighbor
Check Zero	Disabled — checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications are not checked allowing receipt of RIP

Label	Description
	messages even if mandatory zero fields are non-zero for the neighbor Enabled — checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting non-compliant RIP messages is enabled for the neighbor
Message Size	Displays the maximum number of routes per RIP update message
Preference	Displays the preference of RIP routes from the neighbor
Auth. Type	Specifies the authentication type
Update Timer	Displays the current setting of the RIP update timer value expressed in seconds
Timeout Timer	Displays the current RIP timeout timer value expressed in seconds
Export Policies	Displays the export route policy that is used to determine routes advertised to all peers
Import Policies	Displays the import route policy that is used to determine which routes are accepted from RIP neighbors

Sample output with advertised routes

```
A:ALA-A# show router rip neighbors interface advertised-routes
=====
RIP Advertised Routes
=====
Destination      Interface      NextHop      Metric  Tag      TTL
-----
10.0.0.2/32      10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.5/32      10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.8/32      10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.9/32      10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.10/32     10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.11/32     10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.12/32     10.1.8.12     0.0.0.0      1       0x0000   n/a
10.0.0.13/32     10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.14/32     10.1.8.12     0.0.0.0      16      0x0000   n/a
10.0.0.15/32     10.1.8.12     0.0.0.0      2       0x0000   n/a
10.0.0.16/32     10.1.8.12     0.0.0.0      3       0x0000   n/a
-----
No. of Advertised Routes: 11
=====
A:ALA-A#
```

peer

Syntax

peer [ip-int-name]

Context

show>router>rip

Platforms

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

Description

This command displays RIP peer information.

Parameters

ip-int-name

Displays peer information for peers on the specified IP interface.

Default Displays peers for all interfaces

Output

The following output is an example of RIP peer information, and [Table 39: Output fields: peer](#) describes the output fields.

Sample output

```
*A:dut-c>show>router>rip# peer

=====
RIP Peers
=====
Peer IP Addr      Interface Name      Version      Last Update
-----
10.1.4.2          to-ixia-Nw-4       RIPv2        25
-----
No. of Peers: 1
=====
*A:dut-c>show>router>rip#
```

Table 39: Output fields: peer

Label	Description
Peer IP Addr	Displays the IP address of the peer router
Interface Name	Displays the peer interface name
Version	Displays the version of RIP running on the peer

Label	Description
Last Update	Displays the number of days since the last update
No. of Peers	Displays the number of RIP peers

statistics

Syntax

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

Context

show>router>rip

Platforms

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

Description

This command displays interface level statistics for the RIP protocol.

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

If an IP address or interface name is specified, only data about the specified RIP interface is displayed.

Parameters

ip-addr* | *ip-int-name

Displays statistics for the specified IP interface.

Output

The following output is an example of RIP statistics information, and [Table 40: Output fields: statistics](#) describes the output fields.

Sample output

```
*A:dut-c>show>router>rip# statistics

=====
RIP Statistics
=====
Learned Routes      : 2,000          Timed Out Routes    : 0
Current Memory      : 1,944,096      Maximum Memory      : 4,456,640

-----
Interface "to-ixia-Nw-4"
-----
Primary IP          : 10.1.4.1        Update Timer        : 30
Timeout Timer       : 180             Flush Timer         : 120

Counter              Total              Last 5 Min          Last 1 Min
-----
Updates Sent         450983              656                 0
```



```

Triggered Updates          88          0          0
Bad Packets Received       0          0          0
RIPv1 Updates Received     0          0          0
RIPv1 Updates Ignored      0          0          0
RIPv1 Bad Routes           0          0          0
RIPv1 Requests Received    0          0          0
RIPv1 Requests Ignored     0          0          0
RIPv2 Updates Received     404218      640        80
RIPv2 Updates Ignored      0          0          0
RIPv2 Bad Routes           0          0          0
RIPv2 Requests Received    0          0          0
RIPv2 Requests Ignored     0          0          0
Authentication Errors       0          0          0

```

```

=====
*A:dut-c>show>router>rip#

```

Table 40: Output fields: statistics

Label	Description
Learned Routes	Displays the number of RIP-learned routes that were exported to RIP neighbors
Timed Out Routes	Displays the number of routes that have been timed out
Current Memory	Displays the amount of memory used by this RIP router instance
Maximum Memory	Displays the amount of memory allocated for this RIP router instance
Interface	Displays the name of each interface configured in RIP and associated RIP statistics.
Primary IP	Displays the interface IP address
Update Timer	Displays the current setting of the RIP update timer value expressed in seconds
Timeout Timer	Displays the current RIP timeout timer value expressed in seconds
Flush Timer	Displays the number of seconds after a route has been declared invalid that it is flushed from the route database
Updates Sent	Total — total number of RIP updates that were sent Last 5 Min — number of RIP updates that were sent in the last 5 minutes Last 1 Min — number of RIP updates that were sent in the last 1 minute
Triggered Updates	Total — total number of triggered updates sent. These updates are sent before the entire RIP routing table is sent Last 5 Min — number of triggered updates that were sent in the last 5 minutes

Label	Description
	Last 1 Min — number of triggered updates that were sent in the last 1 minute
Bad Packets Received	Total — total number of RIP updates received on this interface that were discarded as invalid Last 5 Min — number of RIP updates received on this interface that were discarded as invalid in the last 5 minutes Last 1 Min — number of RIP updates received on this interface that were discarded as invalid in the last 1 minute
RIPv1 Updates Received	Total — total number of RIPv1 updates received Last 5 Min — number of RIPv1 updates received in the last 5 minutes Last 1 Min — number of RIPv1 updates received in the last 1 minute
RIPv1 Updates Ignored	Total — total number of RIPv1 updates ignored Last 5 Min — number of RIPv1 updates ignored in the last 5 minutes Last 1 Min — number of RIPv1 updates ignored in the last 1 minute
RIPv1 Bad Routes	Total — total number of bad routes received from the peer Last 5 Min — number of bad routes received from the peer in the last 5 minutes Last 1 Min — number of bad routes received from the peer in the last minute
RIPv1 Requests Received	Total — total number of times the router received RIPv1 route requests from other routers Last 5 Min — number of times the router received RIPv1 route requests from other routers in the last 5 minutes Last 1 Min — The number of times the router received RIPv1 route requests from other routers in the last 1 minute
RIPv1 Requests Ignored	Total — total number of times the router ignored RIPv1 route requests from other routers Last 5 Min — number of times the router ignored RIPv1 route requests from other routers in the last 5 minutes Last 1 Min — number of times the router ignored RIPv1 route requests from other routers in the last 1 minute
RIPv2 Updates Received	Total — total number of RIPv2 updates received Last 5 Min — number of RIPv2 updates received in the last 5 minutes

Label	Description
	Last 1 Min — number of RIPv2 updates received in the last minute
RIPv2 Updates Ignored	Total — total number of RIPv2 updates ignored Last 5 Min — number of RIPv2 updates ignored in the last 5 minutes Last 1 Min — number of RIPv2 updates ignored in the last minute
RIPv2 Bad Routes	Total — total number of RIPv2 bad routes received from the peer Last 5 Min — number of RIPv2 bad routes received from the peer in the last 5 minutes Last 1 Min — number of RIPv2 bad routes received from the peer in the last minute
RIPv2 Requests Received	Total — total number of times the router received RIPv2 route requests from other routers Last 5 Min — number of times the router received RIPv2 route requests from other routers in the last 5 minutes Last 1 Min — number of times the router received RIPv2 route requests from other routers in the last minute
RIPv2 Requests Ignored	Total — total number of times the router ignored RIPv2 route requests from other routers Last 5 Min — number of times the router ignored RIPv2 route requests from other routers in the last 5 minutes Last 1 Min — number of times the router ignored RIPv2 route requests from other routers in the last minute
Authentication Errors	Total — total number of authentication errors to secure table updates Last 5 Min — number of authentication errors to secure table updates in the last 5 minutes Last 1 Min — number of authentication errors to secure table updates in the last minute

3.9.2.3 Clear commands

database

Syntax

database

Context

clear>router>rip

Platforms

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

Description

This command clears all routes in the RIP database.

export

Syntax

export

Context

clear>router>rip

Platforms

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

Description

This command re-evaluates all effective export policies.

statistics

Syntax

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

Context

clear>router>rip

Platforms

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

Description

This command clears statistics for RIP neighbors.

Parameters

neighbor *ip-int-name* | *ip-address*

Clears the statistics for the specified RIP interface.

Default Clears statistics for all RIP interfaces

3.9.2.4 Debug RIP commands

auth

Syntax

[no] auth [neighbor *ip-int-name* | *ip-addr*]

Context

debug>router>rip

Platforms

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

Description

This command enables debugging for RIP authentication.

Parameters

neighbor *ip-addr* | *ip-int-name*

Debugs the RIP authentication for the neighbor IP address or interface.

error

Syntax

[no] error [neighbor *ip-int-name* | *ip-addr*]

Context

debug>router>rip

Platforms

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

Description

This command enables debugging for RIP errors.

Parameters

neighbor *ip-addr* | *ip-int-name*

Debugs the RIP errors sent on the neighbor IP address or interface.

events

Syntax

[no] events [**neighbor** *ip-int-name* | *ip-addr*]

Context

debug>router>rip

Platforms

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

Description

This command enables debugging for RIP events.

Parameters

neighbor *ip-addr* | *ip-int-name*

Debugs the RIP events sent on the neighbor IP address or interface.

holddown

Syntax

[no] holddown [**neighbor** *ip-int-name* | *ip-addr*]

Context

debug>router>rip

Platforms

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

Description

This command enables debugging for RIP hold downs.

Parameters

neighbor *ip-addr* | *ip-int-name*

Debugs the RIP hold downs sent on the neighbor IP address or interface.

packets

Syntax

[no] packets [**neighbor** *ip-int-name* | *ip-addr*]

Context

debug>router>rip

Platforms

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

Description

This command enables debugging for RIP packets.

Parameters

neighbor *ip-addr* | *ip-int-name*

Debugs the RIP packets sent on the neighbor IP address or interface.

request

Syntax

[no] request [neighbor *ip-int-name* | *ip-addr*]

Context

debug>router>rip

Platforms

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

Description

This command enables debugging for RIP requests.

Parameters

neighbor *ip-addr* | *ip-int-name*

Debugs the RIP requests sent on the neighbor IP address or interface.

trigger

Syntax

[no] trigger [neighbor *ip-int-name* | *ip-addr*]

Context

debug>router>rip

Platforms

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

Description

This command enables debugging for RIP trigger updates.

Parameters

neighbor *ip-addr* | *ip-int-name*

Debugs the RIP updates sent on the neighbor IP address or interface.

updates

Syntax

[no] updates [neighbor *ip-int-name* | *ip-addr*]

Context

debug>router>rip

Platforms

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

Description

This command enables debugging for RIP updates.

Parameters

neighbor *ip-addr* | *ip-int-name*

Debugs the RIP updates sent on the neighbor IP address or interface.

4 OSPF

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



Note:

OSPF and OSPFv3 are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

4.1 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. The Nokia implementation of OSPF conforms to OSPF Version 2 specifications presented in RFC 2328, OSPF Version 2 and OSPF Version 3 specifications presented in RFC 2740, OSPF for IPv6. Routers running OSPF can be enabled with minimal configuration. All default and command parameters can be modified.

Changes between OSPF for IPv4 include the following:

- 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.
- Flooding scope for LSAs has been generalized.
- 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.
- Most field and packet-size limitations present in OSPF for IPv4 have been relaxed.
- Option handling has been made more flexible.

Key OSPF features are:

- backbone areas
- stub areas
- Not-So-Stubby areas (NSSAs)
- virtual links
- authentication
- route redistribution

- routing interface parameters
- OSPF-TE extensions (the Nokia implementation allows MPLS fast reroute)

The 7210 SAS supports IGP-LDP synchronization on OSPF routes. For information, see "IGP-LDP and Static Route-LDP Synchronization" in the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide*.



Note:

Static route-LDP synchronization is supported on all 7210 SAS platforms as described in this document.

4.1.1 OSPF areas

The hierarchical design of OSPF allows a collection of networks to be grouped into a logical area. An area topology is concealed from the rest of the AS which significantly reduces OSPF protocol traffic. With the correct 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.

4.1.1.1 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) then 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 6: Backbone area



4.1.1.2 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 6: 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.

4.1.1.3 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 6: Backbone area](#), area 0.0.0.3 could be configured as a NSSA area.

4.1.1.3.1 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, 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 7: 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 super backbone described previously apply and Type 3 LSAs should be generated. If the OSPF domain does not match, the route type is 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.

4.1.1.3.2 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 is not used. Even with a super backbone configuration it is regarded as a inter-area connection.

Figure 8: 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 required 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.

4.1.1.3.3 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 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 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, the OSPF super backbone is implemented.

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

4.1.1.3.4 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.

4.1.1.3.5 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 does 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 are ignored by any other PE routers that receive them.

DN-BIT loop avoidance is also supported.

4.1.1.3.6 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.

4.1.1.3.7 Sham links

A sham link is only required if a back door link (shown as the red link in [Figure 8: Sham links](#)) is present, otherwise configuring an OSPF super backbone will probably suffice.

4.1.2 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.

This feature is supported on IES and VPRN interfaces as well as on virtual links.

The re-keying procedure defined in RFC 4552 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.

4.1.3 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 6: 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.

4.1.4 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.

4.1.5 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.

4.1.6 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.

4.1.7 Authentication

All OSPF protocol exchanges can be authenticated. This means that only trusted routers can participate in autonomous system routing. The Nokia 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. The Nokia 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.

4.1.8 Multiple OSPF instances

The main route table manager (RTM) can create multiple instances of OSPF by extending the current creation of an instance. A specific interface can only be a member of a single OSPF instance. When an interface is configured in a specific domain and needs to be moved to another domain the interface must first be removed from the old instance and recreated in the new instance.

4.1.8.1 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, only routes installed by the base instance are 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.

4.1.8.2 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.

4.1.9 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).

4.1.10 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.

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



Note:

IP FRR is not supported on 7210 SAS devices. 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.

4.2.1 LFA configuration



Note:

IP FRR is not supported on 7210 SAS platforms. LFA is supported on 7210 SAS platforms 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.

4.2.1.1 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 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 fails.

4.2.2 ECMP considerations


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

4.2.3 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 9: Example topology 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 10: Example topology 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 since 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):

$$\text{Distance\_opt}(\text{R2}, \text{R5}) < \text{Distance\_opt}(\text{R2}, \text{R1}) + \text{Distance\_opt}(\text{R1}, \text{R5})$$

and,

$$\text{Distance\_opt}(\text{R2}, \text{R5}) < \text{Distance\_opt}(\text{R2}, \text{PN}) + \text{Distance\_opt}(\text{PN}, \text{R5})$$

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 picks 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, it picks one based on the least cost.
3. If more than one LFA next-hop with the same cost results from step (b), SPF selects the first one. This is not a deterministic selection and varies 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 picks 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, it picks one based on the least cost.
4. If more than one LFA next-hop with equal cost results from Step C, SPF selects the first one from the remaining set. This is not a deterministic selection and varies following each SPF calculation.
5. If no LFA next-hop results from Step D, SPF reruns Steps B-D 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.

4.3 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.

4.3.1 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 is still stored temporarily in the template module until the **commit** is executed under the **route-next-hop-policy** context. Any temporary parameter changes are lost if the user enters the **abort** command before the **commit** command.
- 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 has no effect on the prior deletion or creation of a template.

When the **commit** command is issued, IS-IS or OSPF reevaluates the templates and if there are any net changes, it schedules a new LFA SPF to recompute the LFA next-hop for the prefixes associated with these templates.

4.3.1.1 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 is 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 first attempts 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 simply prunes all links belonging to the specified admin group before making the LFA backup next-hop selection for a prefix.

If the same group name is part of both **include** and **exclude** statements, the **exclude** statement wins. 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.

4.3.1.2 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 that use separate fibers but are carried in the same fiber conduit. If the conduit is accidentally cut, all the fiber links are cut which means all IP interfaces using these fiber links 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 is 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 selects an 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.

4.3.1.3 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 is 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 the **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 IS-IS or OSPF interface to a group is performed in the **config>router>interface>if-attribute** context. This is used by IS-IS 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.

4.3.1.4 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 falls 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}

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

4.3.2 Application of route next-hop policy template to an interface

When the route next-hop policy template is configured with the needed 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 fails.

If the user excluded the interface from LFA using the **loopfree-alternate-exclude** command, 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 is not rejected but results in no action taken.

4.3.3 Excluding prefixes from LFA SPF

In the current SR OS 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-policy1* [*prefix-policy2...up to 5*]

config>router>ospf>loopfree-alternate-exclude prefix-policy *prefix-policy1* [*prefix-policy2...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
    action accept
    exit
```

```
exit
default-action reject
exit
```

The default action of the preceding **loopfree-alternate-exclude** command when not 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 which does not match any entry in the policy is accepted into LFA SPF.

4.3.4 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. Links that do not include one or more of the admin-groups in the **include-group** statements and links that belong to admin-groups that 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. The pruning applies only to IP next hops.



Note:

Only IP next hops (specified using the command **nh-type ip**) are supported on 7210 SAS and are considered for LFA selection.

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:

- Prune the IP LFA next-hops that 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**.

4.4 Segment routing in Shortest Path Forwarding

OSPF can be configured in 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.

4.4.1 LFA protection using segment routing backup node SID



Note:

- This feature is only supported on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode).
- The **config>router>ospf>segment-routing>backup-node-sid** command associated with the backup node SID is not supported or needed on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, or 7210 SAS-T. Backup node SID configuration must be completed on the 7750 SR AGN.


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 an 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 11: 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 AGN2 to push up to four labels, as shown in the following figure.

Figure 11: Label stack for remote LFA in ring topology

Label Location	Label Name	Assigned By	Protocol	Use
Label 1 (Bottom)	Service (PW, VC) Label	Remote PE	MP-BGP, T-LDP	Identifies Service on Remote PE
Label 2	Inter-Area Label	ABR1	BGP-LU	Identifies Path to Remote PE
Label 3	Intra-Area	AGN1	LDP, RSVP, SR	Identifies Path to ABR1
Label 4 (Top)	R-LFA Label	AGN3	LDP, RSVP, SR	Identifies Path to AGN5



The objective of the LFA protection with 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 attract packets normally destined for ABR1. These packets are forwarded by ABR2 to ABR1 via the inter-ABR link.


As a result, AGN2 pushes 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.

4.4.1.1 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 12: 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 does 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 uses 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 is not 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 is not 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.

4.4.1.2 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 is indicated in the following table.

Table 41: 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

Old LSA/LSP	New LSA/LSP			
	Backup node SID	Local backup node SID	Node SID	Local node SID
Node SID	Old	Old	Equal/Old ⁹	Equal/New ¹⁰
Local Node SID	Old	Old	Equal/Old ⁹	Equal/Old ⁹

4.4.1.3 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 42: OSPF control plane extension fields

Field	Value
Type	2
Length	variable
Flags	1 octet field

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.

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.

The following flags are defined; the "B" flag is new:

```

  0  1  2  3  4  5  6  7
+---+---+---+---+---+---+
|  NP|M |E |V |L |B|  |
+---+---+---+---+---+---+

```

The following table describes OSPF control plane extension flags.

Table 43: 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-D.filsfils-spring-segment-routing-ldp-interop.
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:

Flag	Description
	• a 32-bit index defining the offset in the SID/Label space advertised by this router • a 24-bit label where the 20 rightmost bits are used for encoding the label value

4.5 OSPF configuration process overview

The following figure shows the process to provision basic OSPF parameters.

Figure 13: OSPF configuration and implementation flow



4.6 Configuration notes

This section describes OSPF configuration restrictions.

4.6.1 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.

4.6.1.1 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.

4.7 Configuring OSPF with CLI

This section provides information to configure Open Shortest Path First (OSPF) using the command line interface.

4.8 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 are not automatically restarted with the new router ID. Shut down and restart the protocol to initialize the new router ID.

- **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.

4.9 Basic OSPF configuration

This section provides information to configure OSPF 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 system interface IP address is used.

- one or more areas
- interfaces (interface "system")

Example: Basic OSPF configuration output

```
ALA-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
-----
ALA-A>config>router>ospf#
```

4.9.1 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#
```

4.10 Configuring OSPF components

The following section describes the syntax used to configure the OSPF components.

4.10.1 Configuring OSPF parameters

Example: Basic OSPF 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
-----
A:ALA-49>config>router>ospf# ex
```

4.10.2 Configuring an OSPF 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 your network consists of multiple areas you must also configure a backbone area (0.0.0.0) on at least one router. The backbone consists 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
```

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:
```

4.10.3 Configuring a 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.

Use the following syntax to configure virtual links.

```
ospf
  area area-id
  stub
    default-metric metric
    summaries
```

Example: Stub configuration output

```
ALA-A>config>router>ospf>area># info
-----
...
        area 0.0.0.0
        exit
        area 0.0.0.20
  stub
        exit
        exit
...
-----
ALA-A>config>router>ospf#
```

4.10.4 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 stub areas.

```
ospf ospf-instance
  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: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
exit
area 0.0.0.20
  stub
  exit
exit
area 0.0.0.25
  nssa
  exit
exit
-----
A:ALA-49>config>router>ospf#
```

4.10.5 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 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 router ID, not the local router ID. The transit area cannot be a stub area or an NSSA.

Use the following syntax to configure stub areas.

```
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: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
exit
area 0.0.0.20
    stub
    exit
exit
area 0.0.0.25
    nssa
    exit
exit
area 1.2.3.4
    exit
-----
A:ALA-49>config>router>ospf#
```

4.10.6 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 ignores ingress OSPF protocol packets and not transmit any OSPF protocol packets.

An interface can be part of more than one area, as specified in RFC5185. 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
```

Example: 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
```

4.10.7 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 **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#

4.10.8 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 "3WErEDoxyzQ" hash
            exit
        exit
        area 1.2.3.4
        exit
-----
A:ALA-49>config>router>ospf#

```

4.10.9 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 a stub area](#) and [Configuring a Not-So-Stubby Area](#).

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: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#
```

4.10.10 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 44: Route preference defaults by route type

Route type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes ¹¹
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

¹¹ Preference for OSPF internal routes is configured with the **preference** command.

Route type	Preference	Configurable
IS-IS level 2 external	165	Yes
BGP	170	Yes

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](#).

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
-----
```

4.11 OSPF configuration management tasks

This section describes the OSPF configuration management tasks.

4.11.1 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#
```

4.11.2 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, the last 32 bits of the chassis MAC address is used as the router ID.

4.11.3 Modifying OSPF parameters

You can change or remove existing OSPF parameters in the CLI or NMS. The changes are applied immediately.

Example

The following example shows the command usage for an 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 shows the command usage for OSPF configuration with the modifications entered in the previous example.

```
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#

```

4.12 OSPF command reference

- [Command hierarchies](#)
- [Command descriptions](#)

4.12.1 Command hierarchies

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

4.12.1.1 Configuration commands for OSPF

```

config
- router
- [no] ospf [ospf-instance] [router-id]
- advertise-router-capability {link | area | as}
- no advertise-router-capability
- [no] area area-id
- area-range ip-prefix/mask [advertise | not-advertise]
- no area-range ip-prefix/mask
- [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
- lfa-policy-map route-nh-template template-name
- no lfa-policy-map

```

```

- loopfree-alternate-exclude
- 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 ip-prefix/mask [advertise | not-advertise]
  - no area-range ip-prefix/mask
  - 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
- [no] ldp-over-rsvp
- loopfree-alternate [remote-lfa]
- loopfree-alternate remote-lfa [max-pq-cost value]
- no loopfree-alternate

```

```

5)]
- loopfree-alternate-exclude prefix-policy prefix-policy [prefix-policy...(up to
- 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

```

4.12.1.2 Configuration commands for OSPF3

```

config
- router
- [no] ospf3
- [no] area area-id
- area-range ip-prefix/ipv6 [advertise | not-advertise]
- no area-range ip-prefix/ipv6
- [no] blackhole-aggregate
- [no] interface ip-int-name [secondary]
- authentication bidirectional sa-name
- authentication inbound sa-name outbound sa-name
- no authentication
- 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
- lfa-policy-map route-nh-template template-name
- no lfa-policy-map
- loopfree-alternate-exclude
- 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
- [no] nssa
- area-range ip-prefix/ipv6 [advertise | not-advertise]
- no area-range ip-prefix/ipv6
- 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 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
- [no] ldp-over-rsvp
- 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
- [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]]

```

4.12.1.3 Show commands

```
show
- router
  - ospf [ospf-instance]
  - 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 | ipv6-address] [detail]
    - interface [area area-id] [detail]
    - interface [ip-int-name | ip-address | ipv6-address] database [detail]
    - neighbor [remote ip-address] [detail]
    - neighbor [ip-int-name] [router-id] [detail]
    - opaque-database area area-id | as [adv-router router-id][ls-id] [detail]
    - prefix-sids [ip-prefix[/prefix-length]] [sid sid] [adv-router router-id]
    - range [area-id]
    - sham-link [interface-name] [detail]
    - sham-link interface-name remote ip-address [detail]
    - sham-link-neighbor [detail]
    - sham-link-neighbor interface-name remote ip-address [detail]
    - spf
    - statistics
    - status
    - virtual-link [detail]
    - virtual-neighbor [remote ip-address] [detail]
```

4.12.1.4 Clear commands

```
clear
- router
  - ospf [ospf-instance]
  - database [purge]
  - export
  - neighbor [ip-int-name | ip-address]
  - statistics
```

4.12.1.5 Debug commands

```
debug
- router
  - ospf [ospf-instance]
  - ospf3
    - 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] [interface-name] [ingress | egress] [detail]
- no packet
- rtm [ip-addr]
- no rtm
- spf [type] [dest-addr]
- no spf
- virtual-neighbor [ip-address]
- no virtual-neighbor
```

4.12.2 Command descriptions

- [Configuration commands](#)
- [Show commands](#)
- [Clear commands](#)
- [OSPF debug commands](#)

4.12.2.1 Configuration commands

4.12.2.1.1 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>interface
config>router>ospf3>area>interface
config>router>ospf3>area>virtual-link
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description



Note:

The **config>router>ospf>segment-routing** context is supported only on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode).

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.

OSPF Protocol Handling

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure router ospf** command instantiates the protocol in the **no shutdown** state and resources are allocated to enable the node to process the protocol.

To deallocate resources, issue the **configure router ospf shutdown** and **configure router no ospf** commands to allow the node to boot up correctly after the reboot. It is not sufficient to only issue a **configure router ospf shutdown** command.

Resources for OSPF are allocated when the OSPF context is enabled either in the base routing instance or the VPRN service instance. Resources are deallocated when the configuration of the last OSPF context under either base routing instances or VPRN service is removed or shut down.

OSPFv3 Protocol Handling

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure router ospf3** command instantiates the protocol in the **no shutdown** state and resources are allocated to enable the node to process the protocol.

To deallocate resources, issue the **configure router ospf3 shutdown** and **configure router no ospf3** commands to allow the node to boot up correctly after the reboot. It is not sufficient to only issue a **configure router ospf3 shutdown** command.

Resources are allocated when the OSPFv3 instance is enabled and resources are deallocated when the OSPFv3 instance is removed or shut down.

4.12.2.1.2 OSPF global commands

ospf

Syntax

[no] ospf [ospf-instance] [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 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 command 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, the last 32 bits of the chassis MAC address are used.

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 to reverts to the default value.



Note:

The number of OSPF instances supported on various 7210 SAS platforms are different. Contact a Nokia representative for information about the supported scaling limits.

Default

no ospf

Parameters

ospf-instance

Specifies a unique integer that identifies an 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 version 6 of the Internet Protocol (IPv6).

When an OSPF instance is created, the protocol is enabled. To start or suspend execution of the OSPF protocol without affecting the configuration, use the **no shutdown** command.

The **no** form of this command deletes the OSPF protocol instance removing all associated configuration parameters.

Default

no ospf

advertise-router-capability

Syntax

advertise-router-capability {link | area | as}

no advertise-router-capability

Context

config>router>ospf

Platforms

7210 SAS-Mxp

Description

This command enables the advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A router information (RI) LSA as defined in RFC 4970 advertises the following capabilities:

- OSPF graceful restart capable: no
- OSPF graceful restart helper: yes, when enabled
- OSPF stub router support: yes
- OSPF traffic engineering support: yes, when enabled
- OSPF point-to-point over LAN: yes
- OSPF experimental TE: no

The **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.

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, configure all routers in a domain 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 corresponding bit.

Figure 14: 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 is set in the AS-external LSA.

The **no** form of this command removes the ASBR status and withdraws the routes redistributed from the RTM into this instance of 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 RFC1583 and earlier RFCs.

RFC1583 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 is favorable to require all routers to run a more current compliance level, this command allows the router to use obsolete methods of calculation.

The **no** form of this command enables the post-RFC1583 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. The command then disables IGP-LDP synchronization for all interfaces. This command does not delete the interface configuration. The **no** form of this command must be entered to re-enable IGP-LDP synchronization for this routing protocol.

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

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. Export policies are only in effect if OSPF is configured as an ASBR.

If no export policy is specified, non-OSPF routes are not exported from the routing table manager to OSPF.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered overrides 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. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

The specified names must already be defined.

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 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

log percentage

Specifies the percentage of the export limit, at which point a warning log message and SNMP notification are sent.

Values 1 to 100

external-db-overflow

Syntax

external-db-overflow *limit interval*

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 enables limits on the number of non-default AS-external-LSA entries that can be stored in the LSDB and specifies a wait timer before processing these after the limit is exceeded.

The *limit* value specifies the maximum number of non-default AS-external-LSA entries that can be stored in the link-state database (LSDB). Placing a limit on the non-default AS-external-LSAs in the LSDB protects

the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reach or exceed the *limit*, the table is in an overflow state. When in an overflow state, the router will not originate any new AS-external-LSAs. In fact, it withdraws all the self-originated non-default external LSAs.

The *interval* 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** must be set identically on all routers attached to any regular OSPF area. OSPF stub areas and not-so-stubby areas (NSSAs) are excluded.

The **no** form of this command disables limiting the number of non-default AS-external-LSA entries.

Default

no external-db-overflow

Parameters

limit

Specifies the maximum number of non-default AS-external-LSA entries that can be stored in the LSDB before going into an overflow state, expressed as a decimal integer.

Values -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 external routes.

A route can be learned by the router from different protocols, in which case the costs (metrics) are not comparable. When this occurs the preference is used to decide which route is used.

Different protocols should not be configured with the same preference; if this occurs the tiebreaker is per the default preference table as defined in the [Table 45: Route preference defaults by route type](#). If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of which route to use is determined by the configuration of the **ecmp** command in the **config>router** context.

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

Default

external-preference 150

Parameters

preference

Specifies the preference for external routes expressed as a decimal integer. The following table lists the defaults for different route types.

Table 45: Route preference defaults by route type

Route type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes ¹²
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

Values 1 to 255

¹² Preference for OSPF internal routes is configured using the **preference** command.

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 reconverge 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), be a restarting router, or both. The 7210 SAS supports only helper mode. This facilitates the graceful restart of neighbors but does not act as a restarting router.

The **no helper-disable** command enables helper support and is the default when graceful-restart is enabled.

Default

disabled

ldp-over-rsvp

Syntax

[no] ldp-over-rsvp

Context

config>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables LDP-over-RSVP processing in this OSPF instance.

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 level.

When this command is enabled, it instructs 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 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 appending the **remote-lfa** option. When this option is enabled in an IGP instance, SPF performs the remote LFA additional computation

following the regular LFA next-hop calculation when the latter 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, that puts the packets back into the shortest path without looping them back to the node that forwarded them over the repair tunnel. The remote LFA node is referred to as a PQ node. A repair tunnel can, in theory, be an RSVP 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 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, it is not included in LFA calculation regardless of its priority. The prefix tag is, however, used in the main SPF. 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 of whether the user explicitly added the statement "default-action reject" to the prefix policy, a prefix that does not match an entry in the policy is 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, 32 characters maximum. The specified name must have been previously defined.

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 routing, but is not used for transit traffic. Traffic destined for directly attached interfaces continues to reach the router.

Enter a timeout value to put the IGP in an overload state. The IGP enters 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, when **overload-on-boot** is configured under OSPF with no timeout value, the router remains in the overload state indefinitely after a reboot.

The **no** form of this command reverts to the default value. 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 stub networks should be advertised with a maximum metric value when the system goes into overload state for any reason. When enabled, the system uses the maximum metric value. When this command is enabled and the router is in overload, all stub interfaces, including loopback and system interfaces, are advertised at the maximum metric.

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

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 boot-up in the overload state until one of the following events occur:

- timeout timer expires
- manual override of the current overload state is entered with the **no overload** command

The **no overload** command does not affect the **overload-on-boot** function.

The default timeout value is 60 seconds, which means that after 60 seconds in overload status the 7210 SAS recovers (changes back to non-overload status). However, when the **overload-on-boot** command is configured under OSPF with no *timeout* value, the router remains in the overload state indefinitely after a reboot.

The **no** form of this command removes the overload-on-boot functionality from the configuration.

Parameters

timeout seconds

Specifies the number of seconds to reset overloading.

Values 60 to 1800

Default indefinitely in overload

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 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 is used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker decision is made according to the default preferences defined in [Table 46: 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. The following table lists defaults for different route types.

Table 46: Route preference defaults by route type

Route type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes ¹³
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

Values 1 to 255

reference-bandwidth

Syntax

reference-bandwidth *reference-bandwidth*
no reference-bandwidth

Context

config>router>ospf
config>router>ospf3

¹³ Preference for OSPF internal routes is configured using the **preference** command.

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:

$\text{cost} = \text{reference bandwidth} / \text{bandwidth}$

The default reference bandwidth is 100 000 000 kb/s or 100 Gb/s; therefore 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** *ip-int-name* context.

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

Default

reference-bandwidth 100000000

Parameters

reference-bandwidth

Specifies the reference bandwidth in kilobits per second, expressed as a decimal integer.

Values 1 to 1000000000

router-id

Syntax

router-id *ip-address*

no router-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 ID for the OSPF instance.

Configuring the router ID 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. If the router ID is not configured in the **config>router** context, the following applies.

- The system uses the system interface address (which is also the loopback address).
- If a system interface address is not configured, the system uses the last 32 bits of the chassis MAC address.

This is a required command when configuring multiple instances and the instance being configured is not the base instance.

When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for the instance, or reboot the entire router.

By default, the value for non-base instances is 0.0.0.0 and is invalid, in this case the instance of OSPF does not start and when running a show command an error is displayed

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

Parameters

ip-address

Specifies a 32-bit, unsigned integer uniquely identifying the router in the AS.

segment-routing

Syntax

segment-routing

no segment-routing

Context

config>router>ospf

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

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 interpret the SID and label range information and use it for calculating the label to swap to or push for a specific 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 datapath.
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

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

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.

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 fails 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

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

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 as 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 is 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 maximum transmission unit (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

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

Description

This command configures the TTM preference of 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 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

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 timers. Timers control the delay between the receipt of a link state advertisement (LSA) requiring a Dijkstra (Shortest Path First (SPF)) calculation and the minimum time between successive SPF calculations.

Changing the timers affects CPU utilization and network reconvergence times. Lower values reduce convergence time but increase CPU utilization. Higher values reduce CPU utilization but increase 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 **lsa-generate** *lsa-second-wait* interval for the neighbors be equal to or greater than the *lsa-arrival-time* value.

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 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.

Nokia recommends configuring the *lsa-arrival-time* to be equal to or less than the *lsa-second-wait* interval configured in the **lsa-generate** command.

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 LSAs generated, 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 another topology change occurs within the specified *lsa-initial-wait* period, 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 (that is, two times 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) occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, the next SPF runs after 2000 milliseconds, and the next SPF runs 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*.

The timer must be entered in increments of 100 milliseconds. Values entered that do not match this requirement are 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

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.

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

4.12.2.1.3 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

Commands in this context configure an OSPF or OSPF3 area. An area is a collection of network segments within an AS that have been administratively grouped together. The area ID can be specified in dotted decimal notation or as a 32-bit decimal integer.

The **no** form of this command deletes the specified area from the configuration. Deleting the area also removes the OSPF configuration of all the interfaces, virtual links, address ranges, and so on, that are currently assigned to this area.

Default

no area

Parameters

area-id

Specifies the OSPF area ID, expressed in dotted decimal notation or as a 32-bit decimal integer.

Values 0.0.0.0 to 255.255.255.255 (dotted decimal)
 0 to 4294967295 (decimal integer)

area-range

Syntax

```
area-range ip-prefix/mask [advertise | not-advertise]  
no area-range ip-prefix/mask  
no area-range ipv6-prefix/prefix-length  
area-range ipv6-prefix/prefix-length [advertise | not-advertise]
```

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

In the NSSA context, the option 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

ip-prefix	Specifies the IP prefix in dotted decimal notation for the range used by the ABR.		
Values	ip-prefix/mask:	ip-prefix a.b.c.d (host bits must be 0)	
	ipv6-prefix:	x:x:x:x:x:x:x (eight 16-bit pieces)	
		x:x:x:x:x:d.d.d.d	
		x: [0 to FFFF]H	
		d: [0 to 255]D	
	prefix-length:	0 to 128	
mask	Specifies the subnet mask for the range.		
Values	0 to 32 (mask length)		
advertise not-advertise	Specifies whether to advertise the summarized range of addresses into other areas. The advertise keyword indicates the range is advertised, and the not-advertise keyword indicates the range is not advertised.		
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 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

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the metric used by the 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

lfa-policy-map

Syntax

lfa-policy-map route-nh-template *template-name*

no lfa-policy-map

Context

config>router>ospf>area>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 interface. When a route next-hop policy template is applied to an interface in OSPF, it is applied in all areas.

The command in an OSPF interface context can only be executed under the area in which the specified interface is primary, and then applied in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command fails.

If the user excluded the interface from LFA using the command **loopfree-alternate-exclude**, the LFA policy, if applied to the interface, has no effect.

If the user applied a route next-hop policy template to a loopback interface or to the system interface, the command is not rejected, but the policy has no effect on the interface.

The **no** form of this command deletes the mapping of a route next-hop policy template to an OSPF interface.

Parameters

template-name

Specifies the name of the template, 32 characters maximum.

loopfree-alternate-exclude

Syntax

[no] **loopfree-alternate-exclude**

Context

config>router>ospf>area

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 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 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 fail.

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

Default

no loopfree-alternate-exclude

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 OSPF3 Not So Stubby Area (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 an NSSA has the capability to flood external routes that it learns throughout its area and via an ABR to the entire OSPF or OSPF3 domain.

Existing virtual links of a non-stub or NSSA 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.

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]

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 (3 or 7) into an NSSA by an NSSA ABR or ASBR.

When configuring an NSSA with no summaries, the ABR injects a type 3 LSA default route into the NSSA. Some older implementations expect a type 7 LSA default route.

The **no** form of this command disables origination of a default route.

Default

no originate-default-route

Parameters

type-7

Specifies that a type-7 LSA should be used for the default route.

Configure this 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 **originate-default-route** without the **type-7** parameter.

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 an 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 OSPF3 areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an ABR to the entire OSPF or OSPF3 domain.

The **no** form of this command disables the default behavior to automatically redistribute external routes into the NSSA 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

This command adds or removes the stub designation from an OSPF or OSPF3 area. Commands in this context configure the OSPF or OSPF3 stub 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 area or NSSA are 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>stub

```
config>router>ospf>area>nssa  
config>router>ospf3>area>stub  
config>router>ospf3>area>nssa
```

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 command is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub area or NSSA.

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

4.12.2.1.4 Interface and virtual link commands

authentication

Syntax

authentication [*inbound sa-name outbound sa-name*]

authentication bidirectional *sa-name*

no authentication

Context

```
config>router>ospf3>area>interface
```

```
config>router>ospf3>area>virtual-link>if
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the password used by the OSPF3 interface or virtual link to send and receive OSPF3 protocol packets on the interface when simple password authentication is configured.

All neighboring routers must use the same type of authentication and password for correct protocol communication.

By default, no authentication key is configured.

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

Default

no authentication

Parameters

inbound *sa-name*

Specifies the inbound sa-name for OSPF3 authentication, up to 32 characters.

outbound *sa-name*

Specifies the outbound sa-name for OSPF3 authentication, up to 32 characters.

bidirectional *sa-name*

Specifies bidirectional OSPF3 authentication, up to 32 characters.

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

config>router>ospf>area>interface

config>router>ospf>area>virtual-link>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the password used by the OSPF interface or virtual-link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.

All neighboring routers must use the same type of authentication and password for correct 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

Keyword that 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

Keyword that 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

Keyword that 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>ospf>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 interface.

Both simple **password** and **message-digest** authentication are supported.

By default, authentication is not enabled on an interface.

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

Default

no authentication

Parameters

password

Keyword that enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple **password** authentication is enabled.

message-digest

Keyword that enables message digest MD5 authentication in accordance with RFC1321. 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

config>router>ospf3>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 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.



Note:

- BFD is not supported for IPv6 interfaces.
- For more information about the protocols and platforms that support BFD, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide*.

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

Default

no bfd-enable

Parameters

remain-down-on-failure

Keyword that forces adjacency down on BFD failure.

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 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 Interface

If the configured **dead-interval** value applies to an interface, all nodes on the subnet must have the same dead interval.

Virtual Link

If the configured **dead-interval** value 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.

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 no export policy is defined, non OSPF 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.

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

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

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 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.

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

Default

hello-interval 10

Special Cases

OSPF Interface

If the configured **hello-interval** value applies to an interface, all nodes on the subnet must have the same hello interval.

Virtual Link

If the configured **hello-interval** value 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 interface.

By default, interfaces are not activated in any interior gateway protocol, such as OSPF, unless explicitly configured.

The **no** form of this command deletes the OSPF interface configuration for this interface. The **shutdown** command in the **config>router>ospf>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 **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 is moved to this area.



Note:

The IPv6 address is not present for OSPF on the node. It is supported only for OSPF3.

secondary

Keyword that allows 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 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 that configures 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 that configures the interface to maintain this link as a point-to-point link.

message-digest-key

Syntax

message-digest-key *key-id* **md5** [*key* | *hash-key*] [*hash* | *hash2*]

no message-digest-key *key-id*

Context

config>router>ospf>area>interface

config>router>ospf>area>virtual-link

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* value.

Parameters

key-id

Specifies the key ID, 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

Keyword that 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

Keyword that 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 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 packet size used on this interface. If this parameter is not configured, OSPF derives the MTU value from the MTU configured (default or explicitly) in the **config>port>ethernet** context.

If this parameter is configured, the smaller value between the value configured here and the MTU configured (default or explicitly) in any of the previous contexts is used.

To determine the actual packet size, add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured in this command.

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

Default

no mtu

Parameters

bytes

Specifies the MTU to be used by OSPF for this logical interface, in bytes.

Values 512 to 9198

node-sid

Syntax

node-sid *index value*

node-sid *label value*

no node-sid

Context

config>router>ospf>area>interface

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-Sx/S 1/10GE

Description

This command assigns a node SID index or label value to the prefix representing the primary address of an IPv4 network interface of the loopback type. 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 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** command allows an interface to be advertised as an OSPF interface without running the OSPF protocol.

While in passive mode, the interface ignores ingress OSPF protocol packets and does not transmit OSPF protocol packets.

By default, service interfaces defined in the **config>router>service-prefix** context are passive. All other interfaces are not passive.

The **no** form of this command removes the passive property from the OSPF interface.

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 interface that is used in an election of the designated router on the subnet.

This parameter is used only 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 (DR) or backup designated router (BDR).

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

Default

priority 1

Parameters

number

Specifies the interface priority, expressed as a decimal integer. A value of 0 indicates the router is not eligible to be the DR or BDR on the interface subnet.

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

config>router>ospf3>area>virtual-link

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the length of time, in seconds, that OSPF waits 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 acknowledgment has been received, the LSA is retransmitted.

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

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, to transmit an LSA on the interface or virtual link. The **no** form of this command reverts to the default value.

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 via 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 (see area 0.0.0.2 in the following figure), the area border routers (routers 1 and 2 in the following figure) must be connected using a virtual link. The two area border routers form a point-to-point-like adjacency across the transit area. (area 0.0.0.1 in the following figure). A virtual link can be configured only 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 an NSSA.

The **no** form of this command deletes the 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 that identifies the transit area that links the backbone area with the area that has no physical connection with the backbone.

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, the area border routers (such as routers Y and Z) must be connected using a virtual link. The two area border routers form a point-to-point-like adjacency across the transit area (area 0.0.0.4).

4.12.2.2 Show commands

```
ospf
```

Syntax

```
ospf [ospf-instance]
```

Context

```
show>router
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context display OSPF information.



Note:

The number of OSPF instances that can be configured depends on the 7210 SAS platform. Contact Nokia technical support for information about the supported scaling limits.

Parameters

ospf-instance

Specifies the OSPF instance.

Values 0 to 31

```
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 are 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 47: Output fields: OSPF area](#) describes the standard and detailed 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
Active IFs        : 3
Area Bdr Rtrs     : 0
SPF Runs          : 7
Router LSAs       : 3
Summary LSAs      : 0
Nssa ext LSAs     : 0
Total LSAs        : 9
Blackhole Range   : True
Total Nbrs        : 2
Total IFs         : 3
AS Bdr Rtrs       : 0
Last SPF Run      : 10/26/2006 10:09:18
Network LSAs      : 3
Asbr-summ LSAs    : 0
Area opaque LSAs  : 3
LSA Cksum Sum     : 0x28b62
Unknown LSAs      : 0

```

```
*A:ALU_SIM11>show>router>ospf# area detail
```

OSPF Areas (Detailed)

```
Area Id: 0.0.0.0
```

```

Area Id      : 0.0.0.0
Virtual Links : 0
Active IFs    : 2
Area Bdr Rtrs : 1
SPF Runs      : 5
Router LSAs   : 2
Summary LSAs  : 1
Nssa ext LSAs : 0
Total LSAs    : 3
Blackhole Range : True
Type          : Standard
Total Nbrs    : 1
Total IFs     : 2
AS Bdr Rtrs   : 0
Last SPF Run  : 07/06/2010 10:36:45
Network LSAs  : 0
Asbr-summ LSAs : 0
Area opaque LSAs : 0
LSA Cksum Sum : 0x15668
Unknown LSAs  : 0

```

Sample Output for OSPF3

```
*A:Dut-A# show router ospf3 area detail
```

OSPF Areas (Detailed)

```
Area Id: 0.0.0.0
```

```

Area Id      : 0.0.0.0
Key Rollover Int.: 10
Virtual Links : 0
Active IFs    : 3
Area Bdr Rtrs : 0
SPF Runs      : 8
Router LSAs   : 3
IE Pfx LSAs   : 0
Nssa ext LSAs : 0
Total LSAs    : 12
Blackhole Range : True
Type          : Standard
Total Nbrs    : 2
Total IFs     : 3
AS Bdr Rtrs   : 0
Last SPF Run  : 10/09/2012 13:54:11
Network LSAs  : 3
IE Rtr LSAs   : 0
IA Pfx LSAs   : 6
LSA Cksum Sum : 0x67bc8
Unknown LSAs  : 0

```

```
*A:Dut-A# show router ospf3 area 0.0.0.0 detail
```

```
OSPF Area (Detailed) : 0.0.0.0
```

Configuration

```

Area Id      : 0.0.0.0
Key Rollover Int.: 10
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          : 8          Last SPF Run   : 10/09/2012 13:54:11
Router LSAs       : 3          Network LSAs   : 3
IE Pfx LSAs       : 0          IE Rtr LSAs    : 0
Nssa ext LSAs     : 0          IA Pfx LSAs    : 6
Total LSAs        : 12         LSA Cksum Sum  : 0x67bc8
Blackhole Range   : True       Unknown LSAs    : 0
=====
*A:Dut-A#

```

Table 47: Output fields: OSPF area

Label	Description
Area Id	Displays a 32-bit integer uniquely identifying an area
Type	NSSA — area is configured as an NSSA area Standard — area is configured as a standard area (not NSSA or stub) Stub — 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
LSA Count	Displays the total number of LSAs 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

Label	Description
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 AS external LSAs
Blackhole Range	False — no blackhole route is installed for aggregates configured in this area True — a lowest priority 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 link state database (LSDB).

When no options are specified, this command displays brief output for all database entries

Parameters

type keyword

Specifies that the OSPF LSDB information should be filtered based on the type specified by *keyword*.

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-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 OSPF link state database information, and [Table 48: 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	0xcfb9
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  Link Count     : 14
Nbr Rtr Id (1) : 180.0.0.13     I/F Address (1) : 180.0.22.2
No of TOS (1)  : 0              Metric-0 (1)   : 25
Link Type (2)  : Stub Network   Mask (2)       : 255.255.255.0
Network (2)    : 180.0.22.0     Metric-0 (2)   : 25
No of TOS (2)  : 0
Link Type (3)  : Point To Point I/F Address (3) : 180.0.5.2
Nbr Rtr Id (3) : 180.0.0.12     Metric-0 (3)   : 25
No of TOS (3)  : 0
Link Type (4)  : Stub Network   Mask (4)       : 255.255.255.0
Network (4)    : 180.0.5.0      Metric-0 (4)   : 25
No of TOS (4)  : 0
Link Type (5)  : Point To Point I/F Address (5) : 180.0.13.2
Nbr Rtr Id (5) : 180.0.0.8      Metric-0 (5)   : 6
No of TOS (5)  : 0
Link Type (6)  : Stub Network   Mask (6)       : 255.255.255.0
Network (6)    : 180.0.13.0     Metric-0 (6)   : 6
No of TOS (6)  : 0
Link Type (7)  : Point To Point I/F Address (7) : 180.0.14.2
Nbr Rtr Id (7) : 180.0.0.5      Metric-0 (7)   : 6
No of TOS (7)  : 0
Link Type (8)  : Stub Network   Mask (8)       : 255.255.255.0
Network (8)    : 180.0.14.0     Metric-0 (8)   : 6
No of TOS (8)  : 0
Link Type (9)  : Point To Point I/F Address (9) : 180.0.17.2
Nbr Rtr Id (9) : 180.0.0.11     Metric-0 (9)   : 25
No of TOS (9)  : 0
Link Type (10) : Stub Network   Mask (10)      : 255.255.255.0
Network (10)   : 180.0.17.0     Metric-0 (10)  : 25
No of TOS (10) : 0
Link Type (11) : Stub Network   Mask (11)      : 255.255.255.255
Network (11)   : 180.0.0.2      Metric-0 (11)  : 1
No of TOS (11) : 0
Link Type (12) : Stub Network   Mask (12)      : 255.255.255.0
Network (12)   : 180.0.18.0     Metric-0 (12)  : 24
No of TOS (12) : 0
Link Type (13) : Point To Point I/F Address (13) : 180.0.3.2
Nbr Rtr Id (13) : 180.0.0.10    Metric-0 (13)  : 25
No of TOS (13) : 0
Link Type (14) : Stub Network   Mask (14)      : 255.255.255.0
Network (14)   : 180.0.3.0      Metric-0 (14)  : 25
No of TOS (14) : 0
-----

```

```

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#

```

Sample Output for OSPF3

```
*A:Dut-A# show router ospf3 database
```

```

=====
OSPF Link State Database (Type : All)
=====
Type      Area Id      Link State Id  Adv Rtr Id    Age  Sequence  Cksum
-----
Router    0.0.0.0          0.0.0.0       1.1.1.1      116  0x8000006c 0x555a
Router    0.0.0.0          0.0.0.0       3.3.3.3       78   0x80000003 0x3fd0
Router    0.0.0.0          0.0.0.0       6.6.6.6      115  0x80000004 0x6c83
Network   0.0.0.0          0.0.0.3       1.1.1.1      116  0x80000001 0xac65
Network   0.0.0.0          0.0.0.2       6.6.6.6      768  0x80000001 0x668c
Network   0.0.0.0          0.0.0.3       6.6.6.6      118  0x80000001 0xc029
IA Pfx    0.0.0.0          0.0.0.0       1.1.1.1      116  0x80000075 0x6885
IA Pfx    0.0.0.0          0.0.117.51    1.1.1.1      116  0x80000001 0xedf0
IA Pfx    0.0.0.0          0.0.0.0       3.3.3.3       78   0x80000003 0xb994
IA Pfx    0.0.0.0          0.0.0.0       6.6.6.6      115  0x80000009 0xc769
IA Pfx    0.0.0.0          0.0.117.50    6.6.6.6      769  0x80000001 0x3e7b
IA Pfx    0.0.0.0          0.0.117.51    6.6.6.6      118  0x80000002 0x9114
-----

```

```
No. of LSAs: 12
```

```
*A:Dut-A#
```

```
*A:Dut-A# show router ospf3 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   : 1.1.1.1
Link State Id  : 0.0.0.0 (0)
LSA Type       : Router
Sequence No    : 0x8000006c            Checksum        : 0x555a
Age            : 147                   Length          : 56
Options        : --R--EV6
Flags          :
Link Type (1)  : Transit Network       Link Count      : 2
I/F Index (1)  : 2                     DR Rtr ID (1)   : 6.6.6.6
Metric (1)     : 100                   DR I/F Index (1): 2
Link Type (2)  : Transit Network       DR Rtr ID (2)   : 1.1.1.1
I/F Index (2)  : 3                     DR I/F Index (2): 3
Metric (2)     : 100

```

```

-----
Router LSA for Area 0.0.0.0
-----
Area Id       : 0.0.0.0           Adv Router Id   : 3.3.3.3
Link State Id  : 0.0.0.0 (0)
LSA Type      : Router
Sequence No   : 0x80000003        Checksum        : 0x3fd0
Age           : 109               Length           : 56
Options       : --R--EV6
Flags         :
Link Type (1) : Transit Network   Link Count       : 2
I/F Index (1) : 2                 DR Rtr ID (1)   : 1.1.1.1
Metric (1)    : 100               DR I/F Index (1) : 3
Link Type (2) : Transit Network   DR Rtr ID (2)   : 6.6.6.6
I/F Index (2) : 3                 DR I/F Index (2) : 3
Metric (2)    : 100
-----
Router LSA for Area 0.0.0.0
-----
Area Id       : 0.0.0.0           Adv Router Id   : 6.6.6.6
Link State Id  : 0.0.0.0 (0)
LSA Type      : Router
Sequence No   : 0x80000004        Checksum        : 0x6c83
Age           : 146               Length           : 56
Options       : --R--EV6
Flags         :
Link Type (1) : Transit Network   Link Count       : 2
I/F Index (1) : 2                 DR Rtr ID (1)   : 6.6.6.6
Metric (1)    : 100               DR I/F Index (1) : 2
Link Type (2) : Transit Network   DR Rtr ID (2)   : 6.6.6.6
I/F Index (2) : 3                 DR I/F Index (2) : 3
Metric (2)    : 100
-----
Network LSA for Area 0.0.0.0
-----
Area Id       : 0.0.0.0           Adv Router Id   : 1.1.1.1
Link State Id  : 0.0.0.3 (3)
LSA Type      : Network
Sequence No   : 0x80000001        Checksum        : 0xac65
Age           : 148               Length           : 32
Options       : --R--EV6          No of Adj Rtrs  : 2
Router Id (1) : 1.1.1.1           Router Id (2)   : 3.3.3.3
-----
Network LSA for Area 0.0.0.0
-----
Area Id       : 0.0.0.0           Adv Router Id   : 6.6.6.6
Link State Id  : 0.0.0.2 (2)
LSA Type      : Network
Sequence No   : 0x80000001        Checksum        : 0x668c
Age           : 801               Length           : 32
Options       : --R--EV6          No of Adj Rtrs  : 2
Router Id (1) : 6.6.6.6           Router Id (2)   : 1.1.1.1
-----
Network LSA for Area 0.0.0.0
-----
Area Id       : 0.0.0.0           Adv Router Id   : 6.6.6.6
Link State Id  : 0.0.0.3 (3)
LSA Type      : Network
Sequence No   : 0x80000001        Checksum        : 0xc029
Age           : 150               Length           : 32
Options       : --R--EV6          No of Adj Rtrs  : 2
Router Id (1) : 6.6.6.6           Router Id (2)   : 3.3.3.3
-----
IA Pfx LSA for Area 0.0.0.0

```

```
-----
Area Id       : 0.0.0.0          Adv Router Id   : 1.1.1.1
Link State Id : 0.0.0.0 (0)
LSA Type      : IA Pfx
Sequence No   : 0x80000075      Checksum         : 0x6885
Age           : 148              Length            : 52
Ref Ls Type   : 2001            Ref Ls Id         : 0
Ref Adv Rtr   : 1.1.1.1        No of Pfxs        : 1
Prefix (1)    : 1001::1/128
Options (1)   : LA              Metric (1)       : 0
-----
```

IA Pfx LSA for Area 0.0.0.0

```
-----
Area Id       : 0.0.0.0          Adv Router Id   : 1.1.1.1
Link State Id : 0.0.117.51 (30003)
LSA Type      : IA Pfx
Sequence No   : 0x80000001      Checksum         : 0xedf0
Age           : 148              Length            : 44
Ref Ls Type   : 2002            Ref Ls Id         : 3
Ref Adv Rtr   : 1.1.1.1        No of Pfxs        : 1
Prefix (1)    : 2013::/64
Options (1)   :                 Metric (1)       : 0
-----
```

IA Pfx LSA for Area 0.0.0.0

```
-----
Area Id       : 0.0.0.0          Adv Router Id   : 3.3.3.3
Link State Id : 0.0.0.0 (0)
LSA Type      : IA Pfx
Sequence No   : 0x80000003      Checksum         : 0xb994
Age           : 110              Length            : 52
Ref Ls Type   : 2001            Ref Ls Id         : 0
Ref Adv Rtr   : 3.3.3.3        No of Pfxs        : 1
Prefix (1)    : 1001::3/128
Options (1)   : LA              Metric (1)       : 0
-----
```

IA Pfx LSA for Area 0.0.0.0

```
-----
Area Id       : 0.0.0.0          Adv Router Id   : 6.6.6.6
Link State Id : 0.0.0.0 (0)
LSA Type      : IA Pfx
Sequence No   : 0x80000009      Checksum         : 0xc769
Age           : 148              Length            : 52
Ref Ls Type   : 2001            Ref Ls Id         : 0
Ref Adv Rtr   : 6.6.6.6        No of Pfxs        : 1
Prefix (1)    : 1001::2/128
Options (1)   : LA              Metric (1)       : 0
-----
```

IA Pfx LSA for Area 0.0.0.0

```
-----
Area Id       : 0.0.0.0          Adv Router Id   : 6.6.6.6
Link State Id : 0.0.117.50 (30002)
LSA Type      : IA Pfx
Sequence No   : 0x80000001      Checksum         : 0x3e7b
Age           : 801              Length            : 44
Ref Ls Type   : 2002            Ref Ls Id         : 2
Ref Adv Rtr   : 6.6.6.6        No of Pfxs        : 1
Prefix (1)    : 2012::/64
Options (1)   :                 Metric (1)       : 0
-----
```

IA Pfx LSA for Area 0.0.0.0

```
-----
Area Id       : 0.0.0.0          Adv Router Id   : 6.6.6.6
Link State Id : 0.0.117.51 (30003)
LSA Type      : IA Pfx
-----
```

```
Sequence No      : 0x80000002      Checksum         : 0x9114
Age              : 151              Length            : 44
Ref Ls Type      : 2002             Ref Ls Id         : 3
Ref Adv Rtr      : 6.6.6.6          No of Pfxs        : 1
Prefix (1)       : 2023::/64
Options (1)      :                  Metric (1)          : 0
=====
```

*A:Dut-A#

Table 48: Output fields: OSPF database

Label	Description
Area Id	Displays the OSPF area identifier
Type LSA Type	Router — LSA type of router (OSPF) Network — LSA type of network (OSPF) Summary — LSA type of summary (OSPF) ASBR Summary — LSA type of ASBR summary (OSPF) Nssa-ext — LSA area-specific, NSSA external (OSPF) Area opaque — LSA type of area opaque (OSPF) router — LSA type of router (OSPF3) Network — LSA type of network (OSPF3) IE Pfx — LSA type of IE prefix LSA type of IE Pfx (OSPF3)IE Rtr — LSA type of IE Rtr (OSPF3) IA Pfx — LSA type of IA Pfx (OSPF3) Nssa-ext — NSSA area-specific AS external (OSPF3)
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

Label	Description
	R — If clear, a node can participate in OSPF topology distribution without being used to forward transit traffic N — type 7 LSA support MC — multicast support E — external routes support
Prefix Options	P — propagate NSSA LSA MC — multicast support
Flags	None — no flags set V — router is an endpoint for one or more fully adjacent Virtual Links having the described area as the transit area E — router is an AS Boundary Router B — router is an Area Border Router
Link Count	Displays the number of links advertised in the LSA
Link Type (n)	Displays the link type of the <i>n</i> th link in the LSA
Network (n)	Displays the network address of the <i>n</i> th link in the LSA
Metric-0 (n)	Displays the cost metric of the <i>n</i> th link in the LSA

interface

Syntax

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

interface [*area area-id*] [**detail**]

interface [*ip-int-name* | *ip-address* | *ipv6-address* **database** [**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 interface, which 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 large amount of data. Nokia recommends using the **detail** option only when specifying a particular interface.

Parameters

ip-addr

Display only the interface identified by this IP address.

ip-int-name

Display only the interface identified by this interface name.

ipv6-address

Displays only the interface identified by the this IPv6 address.

area area-id

Display all interfaces configured in this area.

database

Keyword to display database information.

detail

Displays detailed information about the interface.

Output

The following outputs are examples of OSPF interface information, and the associated tables describe the output fields.

- Standard output: [Sample output](#), [Table 49: Output fields: OSPF interface](#)
- Detailed output: [Sample output — detailed](#), [Table 50: Output fields: OSPF interface detail](#)

Sample output

```
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
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 Dead Int.      : 0
Bad Versions       : 0
LSA Count          : 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#

*A:ALU_SIM11>show>router>ospf# interface 6.6.6.2 detail

=====
OSPF Interface (Detailed) : 6.6.6.2
=====
-----
Configuration
-----
IP Address         : 6.6.6.2
Area Id            : 0.0.0.0
Hello Intrvl       : 10 sec
Retrans Intrvl     : 5 sec
Cfg Metric         : 0
Transit Delay      : 1
Passive            : False
LFA                : Exclude
Priority           : 1
Rtr Dead Intrvl    : 40 sec
Poll Intrvl        : 120 sec
Advert Subnet      : True
Auth Type          : None
Cfg MTU            : 0
-----
State
-----
Admin Status       : Enabled
Designated Rtr     : 0.0.0.0
IF Type            : Point To Point
Oper MTU           : 1564
Oper Metric         : 100
Te Metric          : 100
Admin Groups       : None
Ldp Sync           : outOfService
Ldp Timer State    : Disabled
Oper State         : Point To Point
Backup Desig Rtr   : 0.0.0.0
Network Type       : Transit
Last Enabled       : 07/06/2010 10:34:11
Bfd Enabled        : No
Te State           : Down
Ldp Sync Wait      : Disabled
Ldp Tm Left        : 0
-----
Statistics
-----
Nbr Count          : 1
Tot Rx Packets     : 353
Rx Hellos           : 314
Rx DBDs            : 31
Rx LSRs            : 1
Rx LSUs            : 4
Rx LS Acks         : 3
Retransmits        : 1
Bad Networks       : 0
Bad Areas          : 6
Bad Auth Types     : 0
Bad Neighbors      : 0
Bad Lengths        : 0
Bad Dead Int.      : 0
Bad Versions       : 0
LSA Count          : 0
If Events           : 7
Tot Tx Packets     : 348
Tx Hellos          : 309
Tx DBDs            : 30
Tx LSRs            : 1
Tx LSUs            : 4
Tx LS Acks         : 4
Discards           : 6
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
=====

```



```
*A:7210-SAS>show>router>ospf# interface C_Port detail

=====
OSPF Interface (Detailed) : C_Port
=====
-----
Configuration
-----
IP Address       : 10.26.26.2
Area Id          : 0.0.0.2
Hello Intrvl     : 10 sec
Retrans Intrvl  : 5 sec
Cfg Metric       : 0
Transit Delay    : 1
Passive          : False
LFA              : Exclude
Priority          : 1
Rtr Dead Intrvl : 40 sec
Poll Intrvl      : 120 sec
Advert Subnet    : True
Auth Type        : None
Cfg MTU          : 0
-----
State
-----
Admin Status     : Enabled
Designated Rtr   : 0.0.0.0
IF Type          : Point To Point
Oper MTU         : 9198
Oper Metric      : 100
Te Metric        : 100
Admin Groups     : None
Ldp Sync         : outOfService
Ldp Timer State  : Disabled
Oper State       : Point To Point
Backup Desig Rtr : 0.0.0.0
Network Type     : Transit
Last Enabled     : 12/14/2010 09:48:30
Bfd Enabled      : No
Te State         : Up
Ldp Sync Wait    : Disabled
Ldp Tm Left      : 0
-----
Statistics
-----
Nbr Count        : 1
Tot Rx Packets   : 22391
Rx Hellos        : 8641
Rx DBDs          : 20
Rx LSRs          : 0
Rx LSUs          : 13531
Rx LS Acks       : 199
Retransmits      : 26
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
If Events        : 1
Tot Tx Packets   : 22273
Tx Hellos        : 8640
Tx DBDs          : 19
Tx LSRs          : 0
Tx LSUs          : 13611
Tx LS Acks       : 3
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
=====
```

```
*A:7210-SAS>show>router>ospf#
```

Sample Output for OSPF3

```
*A:Dut-A# show router ospf3 interface detail
```

```
=====
OSPF Interfaces (Detailed)
=====
-----
Interface : system
-----
IP Address       : 2001:db8::1
Area Id          : 0.0.0.0
Hello Intrvl     : 10 sec
Priority          : 1
Rtr Dead Intrvl : 40 sec
```

Retrans Intrvl	: 5 sec	Poll Intrvl	: 120 sec
Cfg Metric	: 0	Advert Subnet	: True
Transit Delay	: 1	Auth Type	: None
Passive	: True	Cfg MTU	: 0
IPsec InStatSA	:	IPsec OutStatSA	:
IPsec InStatSATmp:			
Admin Status	: Enabled	Oper State	: Designated Rtr
Designated Rtr	: 1.1.1.1	Backup Desig Rtr	: 0.0.0.0
IF Type	: Broadcast	Network Type	: Stub
Oper MTU	: 1500	Last Enabled	: 10/09/2012 13:41:23
Oper Metric	: 0	Bfd Enabled	: No
Te Metric	: 0	Te State	: Down
Admin Groups	: None		
Ldp Sync	: outOfService	Ldp Sync Wait	: Disabled
Ldp Timer State	: Disabled	Ldp Tm Left	: 0
Nbr Count	: 0	If Events	: 2
Tot Rx Packets	: 0	Tot Tx Packets	: 0
Rx Hellos	: 0	Tx Hellos	: 0
Rx DBDs	: 0	Tx DBDs	: 0
Rx LSRs	: 0	Tx LSRs	: 0
Rx LSUs	: 0	Tx LSUs	: 0
Rx LS Acks	: 0	Tx LS Acks	: 0
Retransmits	: 0	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
LSA Count	: 0	LSA Checksum	: 0x0

Interface : to_b

IP Address	: FE80::8E90:D3FF:FEBC:8F5A-"to_b"		
Area Id	: 0.0.0.0	Priority	: 1
Hello Intrvl	: 10 sec	Rtr Dead Intrvl	: 40 sec
Retrans Intrvl	: 5 sec	Poll Intrvl	: 120 sec
Cfg Metric	: 0	Advert Subnet	: True
Transit Delay	: 1	Auth Type	: None
Passive	: False	Cfg MTU	: 0
IPsec InStatSA	:	IPsec OutStatSA	:
IPsec InStatSATmp:			
Admin Status	: Enabled	Oper State	: Backup Desig Rtr
Designated Rtr	: 6.6.6.6	Backup Desig Rtr	: 1.1.1.1
IF Type	: Broadcast	Network Type	: Transit
Oper MTU	: 9198	Last Enabled	: 10/09/2012 13:42:16
Oper Metric	: 100	Bfd Enabled	: No
Te Metric	: 100	Te State	: Down
Admin Groups	: None		
Ldp Sync	: outOfService	Ldp Sync Wait	: Disabled
Ldp Timer State	: Disabled	Ldp Tm Left	: 0
Nbr Count	: 1	If Events	: 4
Tot Rx Packets	: 449	Tot Tx Packets	: 339
Rx Hellos	: 96	Tx Hellos	: 96
Rx DBDs	: 5	Tx DBDs	: 3
Rx LSRs	: 1	Tx LSRs	: 1
Rx LSUs	: 235	Tx LSUs	: 230
Rx LS Acks	: 112	Tx LS Acks	: 9
Retransmits	: 1	Discards	: 1
Bad Networks	: 0	Bad Virt Links	: 0
Bad Areas	: 0	Bad Dest Addrs	: 0

```

Bad Auth Types      : 0
Bad Neighbors       : 0
Bad Lengths         : 0
Bad Dead Int.       : 0
Bad Versions        : 0
LSA Count           : 2
Auth Failures       : 0
Bad Pkt Types       : 0
Bad Hello Int.      : 0
Bad Options         : 0
Bad Checksums       : 0
LSA Checksum        : 0x16aa9

-----
Interface : to_c
-----
IP Address          : FE80::8E90:D3FF:FEBE:8F5A-"to_c"
Area Id             : 0.0.0.0
Hello Intrvl        : 10 sec
Retrans Intrvl      : 5 sec
Cfg Metric          : 0
Transit Delay       : 1
Passive             : False
IPsec InStatSA      :
IPsec InStatSATmp   :
Admin Status        : Enabled
Designated Rtr      : 1.1.1.1
IF Type             : Broadcast
Oper MTU            : 9198
Oper Metric         : 100
Te Metric           : 100
Admin Groups        : None
Ldp Sync            : outOfService
Ldp Timer State     : Disabled
Nbr Count           : 1
Tot Rx Packets      : 117
Rx Hellos           : 97
Rx DBDs             : 2
Rx LSRs             : 1
Rx LSUs             : 13
Rx LS Acks          : 4
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           : 2
Priority            : 1
Rtr Dead Intrvl     : 40 sec
Poll Intrvl         : 120 sec
Advert Subnet       : True
Auth Type           : None
Cfg MTU             : 0
IPsec OutStatSA     :
Oper State          : Designated Rtr
Backup Desig Rtr    : 3.3.3.3
Network Type        : Transit
Last Enabled        : 10/09/2012 13:42:14
Bfd Enabled         : No
Te State            : Down
Ldp Sync Wait       : Disabled
Ldp Tm Left         : 0
If Events           : 77
Tot Tx Packets      : 118
Tx Hellos           : 97
Tx DBDs             : 3
Tx LSRs             : 1
Tx LSUs             : 10
Tx LS Acks          : 7
Discards            : 75
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        : 0x17644

=====
*A:Dut-A#

```

Table 49: 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.

Label	Description
D Rtr Id	Displays the IP interface address of the router identified as the Designated Router (DR) for the network in which this interface is configured. Set to 0.0.0.0 if there is no DR.
BD Rtr Id	Displays the IP interface address of the router identified as the backup DR for the network in which this interface is configured. Set to 0.0.0.0 if there is no backup DR.
Adm	Dn — OSPF on this interface is administratively shut down Up — OSPF on this interface is administratively enabled
Opr	Down — the initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable. Wait — the router is trying to determine the identity of the (backup) DR for the network PToP — the interface is operational, and connects either to a physical point-to-point network or to a virtual link DR — this router is the DR for this network BDR — this router is the backup DR for this network ODR — interface is operational and part of a broadcast or NBMA network on which another router has been selected to be the DR
No. of OSPF Interfaces	Displays the number of interfaces listed
Bfd Enabled	Indicates whether BFD is enabled
LFA	Displays the interface LFA status (included in LFA computation or excluded in LFA computations)

Sample output — detailed

```
*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 Intrvl     : 5 sec
Retrans Intrvl   : 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
Priority         : 1
Rtr Dead Intrvl  : 15 sec
Poll Intrvl      : 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
```

```

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

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#

*A:7210-SAS>show>router>ospf# interface detail

=====

OSPF Interfaces (Detailed)

=====

Interface : system

IP Address	: 10.1.1.4	Priority	: 1
Area Id	: 0.0.0.2	Rtr Dead Intrvl	: 40 sec
Hello Intrvl	: 10 sec	Poll Intrvl	: 120 sec
Retrans Intrvl	: 5 sec	Advert Subnet	: True
Cfg Metric	: 0	Auth Type	: None
Transit Delay	: 1	Cfg MTU	: 0
Passive	: True	Oper State	: Point To Point
Admin Status	: Enabled	Backup Desig Rtr	: 0.0.0.0
Designated Rtr	: 0.0.0.0	Network Type	: Stub
IF Type	: Point To Point	Last Enabled	: 12/14/2010 09:47:33
Oper MTU	: 1500	Bfd Enabled	: No
Oper Metric	: 0	Te State	: Down
Te Metric	: 0		
Admin Groups	: None		
Ldp Sync	: outOfService	Ldp Sync Wait	: Disabled
Ldp Timer State	: Disabled	Ldp Tm Left	: 0
Nbr Count	: 0	If Events	: 1
Tot Rx Packets	: 0	Tot Tx Packets	: 0
Rx Hellos	: 0	Tx Hellos	: 0
Rx DBDs	: 0	Tx DBDs	: 0
Rx LSRs	: 0	Tx LSRs	: 0
Rx LSUs	: 0	Tx LSUs	: 0
Rx LS Acks	: 0	Tx LS Acks	: 0
Retransmits	: 0	Discards	: 0
Bad Networks	: 0	Bad Virt Links	: 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
Bad Versions	: 0	Bad Checksums	: 0
LSA Count	: 0	LSA Checksum	: 0x0

Interface : F_Port

IP Address	: 10.1.1.2	Priority	: 1
Area Id	: 0.0.0.2	Rtr Dead Intrvl	: 40 sec
Hello Intrvl	: 10 sec		

Retrans Intrvl	: 5 sec	Poll Intrvl	: 120 sec
Cfg Metric	: 0	Advert Subnet	: True
Transit Delay	: 1	Auth Type	: None
Passive	: False	Cfg MTU	: 0
Admin Status	: Enabled	Oper State	: Point To Point
Designated Rtr	: 0.0.0.0	Backup Desig Rtr	: 0.0.0.0
IF Type	: Point To Point	Network Type	: Transit
Oper MTU	: 9198	Last Enabled	: 12/14/2010 09:48:07
Oper Metric	: 100	Bfd Enabled	: Yes
Te Metric	: 100	Te State	: Up
Admin Groups	: None		
Ldp Sync	: outOfService	Ldp Sync Wait	: Disabled
Ldp Timer State	: Disabled	Ldp Tm Left	: 0
Nbr Count	: 1	If Events	: 1
Tot Rx Packets	: 21739	Tot Tx Packets	: 20709
Rx Hellos	: 8630	Tx Hellos	: 8629
Rx DBDs	: 19	Tx DBDs	: 20
Rx LSRs	: 0	Tx LSRs	: 11
Rx LSUs	: 12782	Tx LSUs	: 1872
Rx LS Acks	: 308	Tx LS Acks	: 10177
Retransmits	: 13	Discards	: 0
Bad Networks	: 0	Bad Virt Links	: 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
Bad Versions	: 0	Bad Checksums	: 0
LSA Count	: 0	LSA Checksum	: 0x0

Interface : F_Lag

IP Address	: 10.1.1.2	Priority	: 1
Area Id	: 0.0.0.2	Rtr Dead Intrvl	: 40 sec
Hello Intrvl	: 10 sec	Poll Intrvl	: 120 sec
Retrans Intrvl	: 5 sec	Advert Subnet	: True
Cfg Metric	: 0	Auth Type	: None
Transit Delay	: 1	Cfg MTU	: 0
Passive	: False	Oper State	: Point To Point
Admin Status	: Enabled	Backup Desig Rtr	: 0.0.0.0
Designated Rtr	: 0.0.0.0	Network Type	: Transit
IF Type	: Point To Point	Last Enabled	: 12/14/2010 09:48:09
Oper MTU	: 9198	Bfd Enabled	: Yes
Oper Metric	: 50	Te State	: Up
Te Metric	: 50		
Admin Groups	: None	Ldp Sync Wait	: Disabled
Ldp Sync	: outOfService	Ldp Tm Left	: 0
Ldp Timer State	: Disabled	If Events	: 1
Nbr Count	: 1	Tot Tx Packets	: 22347
Tot Rx Packets	: 21885	Tx Hellos	: 8634
Rx Hellos	: 8629	Tx DBDs	: 20
Rx DBDs	: 19	Tx LSRs	: 5
Rx LSRs	: 0	Tx LSUs	: 13152
Rx LSUs	: 13221	Tx LS Acks	: 536
Rx LS Acks	: 16	Discards	: 0
Retransmits	: 16	Bad Virt Links	: 0
Bad Networks	: 0	Bad Dest Adrs	: 0
Bad Areas	: 0	Auth Failures	: 0
Bad Auth Types	: 0	Bad Pkt Types	: 0
Bad Neighbors	: 0	Bad Hello Int.	: 0
Bad Lengths	: 0	Bad Options	: 0
Bad Dead Int.	: 0	Bad Checksums	: 0
Bad Versions	: 0	LSA Checksum	: 0x0
LSA Count	: 0		

```

-----
Interface : C_Lag
-----
IP Address      : 10.1.1.1
Area Id         : 0.0.0.2
Hello Intrvl    : 10 sec
Retrans Intrvl  : 5 sec
Cfg Metric      : 0
Transit Delay   : 1
Passive         : False
Admin Status    : Enabled
Designated Rtr  : 0.0.0.0
IF Type         : Point To Point
Oper MTU        : 9198
Oper Metric     : 50
Te Metric       : 50
Admin Groups    : None
Ldp Sync        : outOfService
Ldp Timer State : Disabled
Nbr Count       : 1
Tot Rx Packets  : 22578
Rx Hellos       : 8628
Rx DBDs         : 20
Rx LSRs         : 12
Rx LSUs         : 13883
Rx LS Acks      : 35
Retransmits     : 23
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
Cfg MTU         : 0
Oper State      : Point To Point
Backup Desig Rtr : 0.0.0.0
Network Type    : Transit
Last Enabled    : 12/14/2010 09:48:33
Bfd Enabled     : Yes
Te State        : Up
Ldp Sync Wait   : Disabled
Ldp Tm Left     : 0
If Events       : 1
Tot Tx Packets  : 21802
Tx Hellos       : 8634
Tx DBDs         : 19
Tx LSRs         : 1
Tx LSUs         : 12831
Tx LS Acks      : 317
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 : C_Port
-----
IP Address      : 10.26.26.2
Area Id         : 0.0.0.2
Hello Intrvl    : 10 sec
Retrans Intrvl  : 5 sec
Cfg Metric      : 0
Transit Delay   : 1
Passive         : False
Admin Status    : Enabled
Designated Rtr  : 0.0.0.0
IF Type         : Point To Point
Oper MTU        : 9198
Oper Metric     : 100
Te Metric       : 100
Admin Groups    : None
Ldp Sync        : outOfService
Ldp Timer State : Disabled
Nbr Count       : 1
Tot Rx Packets  : 22380
Rx Hellos       : 8632
Rx DBDs         : 20
Rx LSRs         : 0
Rx LSUs         : 13531
Rx LS Acks      : 197
Retransmits     : 26
Bad Networks    : 0
Bad Areas       : 0
Priority        : 1
Rtr Dead Intrvl : 40 sec
Poll Intrvl     : 120 sec
Advert Subnet   : True
Auth Type       : None
Cfg MTU         : 0
Oper State      : Point To Point
Backup Desig Rtr : 0.0.0.0
Network Type    : Transit
Last Enabled    : 12/14/2010 09:48:30
Bfd Enabled     : No
Te State        : Up
Ldp Sync Wait   : Disabled
Ldp Tm Left     : 0
If Events       : 1
Tot Tx Packets  : 22262
Tx Hellos       : 8631
Tx DBDs         : 19
Tx LSRs         : 0
Tx LSUs         : 13609
Tx LS Acks      : 3
Discards        : 0
Bad Virt Links  : 0
Bad Dest Addrs  : 0

```

```

Bad Auth Types   : 0
Bad Neighbors    : 0
Bad Lengths      : 0
Bad Dead Int.    : 0
Bad Versions     : 0
LSA Count        : 0
Auth Failures    : 0
Bad Pkt Types    : 0
Bad Hello Int.   : 0
Bad Options      : 0
Bad Checksums    : 0
LSA Checksum     : 0x0
=====

```

Sample Output for OSPF3

A:Dut-A# show router ospf3 interface area 0 detail

```

=====
OSPF Interfaces in Area (Detailed) : 0
=====

```

```

-----
Interface : system
-----

```

```

IP Address       : 2001:db8::1
Area Id          : 0.0.0.0
Hello Intrvl     : 10 sec
Retrans Intrvl   : 5 sec
Cfg Metric       : 0
Transit Delay    : 1
Passive          : True
IPsec InStatSA   :
IPsec InStatSAmp:
Admin Status     : Enabled
Designated Rtr   : 1.1.1.1
IF Type          : Broadcast
Oper MTU         : 1500
Oper Metric      : 0
Te Metric        : 0
Admin Groups     : None
Ldp Sync         : outOfService
Ldp Timer State  : Disabled
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
Cfg MTU          : 0
IPsec OutStatSA  :
Oper State       : Designated Rtr
Backup Desig Rtr : 0.0.0.0
Network Type     : Stub
Last Enabled     : 10/09/2012 13:41:23
Bfd Enabled      : No
Te State         : Down
Ldp Sync Wait    : Disabled
Ldp Tm Left      : 0
If Events        : 2
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 Adrs    : 0
Auth Failures    : 0
Bad Pkt Types    : 0
Bad Hello Int.   : 0
Bad Options      : 0
Bad Checksums    : 0
LSA Checksum     : 0x0

```

```

-----
Interface : to_b
-----

```

```

IP Address       : FE80::8E90:D3FF:FE8F5A-"to_b"
Area Id          : 0.0.0.0
Hello Intrvl     : 10 sec
Retrans Intrvl   : 5 sec
Cfg Metric       : 0
Transit Delay    : 1
Priority          : 1
Rtr Dead Intrvl  : 40 sec
Poll Intrvl      : 120 sec
Advert Subnet    : True
Auth Type        : None

```


Passive	: False	Cfg MTU	: 0
IPsec InStatSA	:	IPsec OutStatSA	:
IPsec InStatSATmp:			
Admin Status	: Enabled	Oper State	: Backup Desig Rtr
Designated Rtr	: 6.6.6.6	Backup Desig Rtr	: 1.1.1.1
IF Type	: Broadcast	Network Type	: Transit
Oper MTU	: 9198	Last Enabled	: 10/09/2012 13:42:16
Oper Metric	: 100	Bfd Enabled	: No
Te Metric	: 100	Te State	: Down
Admin Groups	: None		
Ldp Sync	: outOfService	Ldp Sync Wait	: Disabled
Ldp Timer State	: Disabled	Ldp Tm Left	: 0
Nbr Count	: 1	If Events	: 4
Tot Rx Packets	: 456	Tot Tx Packets	: 346
Rx Hellos	: 103	Tx Hellos	: 103
Rx DBDs	: 5	Tx DBDs	: 3
Rx LSRs	: 1	Tx LSRs	: 1
Rx LSUs	: 235	Tx LSUs	: 230
Rx LS Acks	: 112	Tx LS Acks	: 9
Retransmits	: 1	Discards	: 1
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
LSA Count	: 2	LSA Checksum	: 0x16aa9

Interface : to_c

IP Address	: FE80::8E90:D3FF:FEBE:8F5A-"to_c"	Priority	: 1
Area Id	: 0.0.0.0	Rtr Dead Intrvl	: 40 sec
Hello Intrvl	: 10 sec	Poll Intrvl	: 120 sec
Retrans Intrvl	: 5 sec	Advert Subnet	: True
Cfg Metric	: 0	Auth Type	: None
Transit Delay	: 1	Cfg MTU	: 0
Passive	: False	IPsec OutStatSA	:
IPsec InStatSA	:		
IPsec InStatSATmp:			
Admin Status	: Enabled	Oper State	: Designated Rtr
Designated Rtr	: 1.1.1.1	Backup Desig Rtr	: 3.3.3.3
IF Type	: Broadcast	Network Type	: Transit
Oper MTU	: 9198	Last Enabled	: 10/09/2012 13:42:14
Oper Metric	: 100	Bfd Enabled	: No
Te Metric	: 100	Te State	: Down
Admin Groups	: None		
Ldp Sync	: outOfService	Ldp Sync Wait	: Disabled
Ldp Timer State	: Disabled	Ldp Tm Left	: 0
Nbr Count	: 1	If Events	: 77
Tot Rx Packets	: 124	Tot Tx Packets	: 125
Rx Hellos	: 104	Tx Hellos	: 104
Rx DBDs	: 2	Tx DBDs	: 3
Rx LSRs	: 1	Tx LSRs	: 1
Rx LSUs	: 13	Tx LSUs	: 10
Rx LS Acks	: 4	Tx LS Acks	: 7
Retransmits	: 0	Discards	: 75
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 Versions       : 0
LSA Count          : 2
Bad Options        : 0
Bad Checksums      : 0
LSA Checksum       : 0x17644
=====

```

```

*A:Dut-A#

```

Table 50: 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 DR 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 a router's 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	False — When a point-to-point interface is configured as false, the subnet is not advertised and the endpoints are advertised as host routes True — When a point-to-point interface is configured as true, the subnet is advertised
Transit Delay	Displays the estimated number of seconds it takes to transmit a link-state update packet over this interface
Auth Type	Identifies the authentication procedure to be used for the packet

Label	Description
	None — Routing exchanges over the network or subnet are not authenticated 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 MD5 — A shared secret key is configured in all routers attached to a common network or subnet. For each OSPF protocol packet, the key is used to generate or verify a "message digest" that is appended to the end of the OSPF packet
Passive	False — This interface operates as a normal OSPF interface with regard to adjacency forming and network or link behavior.
	True — No OSPF Hellos will be sent out on this interface, and the router advertises this interface as a stub network or link in its router LSAs.
MTU	The wanted size of the largest packet that can be sent or received on this OSPF interface, specified in octets. This size DOES include the underlying IP header length, but not the underlying layer headers or trailers.
Admin Status	Disabled — OSPF on this interface is administratively shut down
	Enabled — OSPF on this interface is administratively enabled
Oper State	Down — The initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable
	Waiting — The router is trying to determine the identity of the (backup) DR for the network
	Point To Point — The interface is operational and connects either to a physical point-to-point network or to a virtual link
	Designated Rtr — This router is the DR for this network
	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 DR
	Backup Desig Rtr — This router is the backup DR for this network
DR-Id	Displays the IP interface address of the router identified as the DR for the network in which this interface is configured. Set to 0.0.0.0 if there is no DR.
BDR-Id	Displays 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.

Label	Description
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; only traffic sourced or destined for this network is routed to this network Transit — OPSF has established at least one neighbor relationship with any other OSPF router on this network; traffic en route to other networks may be routed via this network
Oper MTU	Displays the operational size of the largest packet that can be sent or received on this OSPF interface, in octets. This size DOES include the underlying IP header length, but not the underlying layer headers or 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
Tx LSRs	Displays the total number of LSRs transmitted on this interface since this interface was last enabled

Label	Description
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 for 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
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

Label	Description
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	Indicates 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	Indicates the MPLS interface TE status from OSPF standpoint
Admin Groups	Indicates the bit-map inherited from MPLS interface that identifies the admin groups to which this interface belongs

neighbor

Syntax

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

neighbor [**remote** | *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, select the neighbors on a specific interface using address or name.

The **detail** option produces a large amount of data. Nokia recommends using the **detail** option only when specifying a particular neighbor.

Parameters

remote

Keyword that specifies the remote OSPF neighbor.

ip-address

Displays neighbor information for the neighbor identified by the specified IP address.

ip-int-name

Displays neighbor information for neighbors of the interface identified by the interface name.

Output

The following outputs are examples of OSPF and OSPFv3 neighbor information, and the associated tables describe the output fields.

- [Sample output, Table 51: Output fields: OSPF neighbor](#)
- [Sample output — detailed, Table 52: Output fields: OSPF/OSPF3 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#

Sample Output for OSPF3
*A:Dut-A# show router ospf3 neighbor
=====
OSPF Neighbors
=====
Interface-Name          Rtr Id          State    Pri  RetxQ  TTL
Area-Id
-----
to_b                    6.6.6.6         Full     1    0      33
  0.0.0.0
to_c                    3.3.3.3         Full     1    0      35
  0.0.0.0
-----
No. of Neighbors: 2
=====
*A:Dut-A#
```

Table 51: Output fields: OSPF neighbor

Label	Description
Nbr IP Addr	Displays the IP address this neighbor is using in its IP source address. On links with no addresses, this will not be 0.0.0.0, but the address of another of the neighbor's interfaces.

Label	Description
Nbr Rtr Id	Displays a 32-bit integer to uniquely identify the neighbor router in the AS
Nbr State	Down — 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, a 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). Two Way — In this state, communication between the two routers is bidirectional. 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 upon the initial database descriptor sequence number. Exchange — In this state, the router is describing its entire link state database by sending database description packets to the neighbor. 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. Full — In this state, the neighboring routers are fully adjacent. These adjacencies now appear in router-LSAs and network-LSAs.
Priority	Displays the priority of this neighbor in the designated router election algorithm. A value of 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 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   : 10.13.7.165           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#
```

Sample output for OSPF3

```
*A:Dut-A# show router ospf3 neighbor detail

=====
OSPF Neighbors
=====
-----
Neighbor Rtr Id : 6.6.6.6 Interface: to_b
-----
Neighbor IP Addr : FE80::225:BAFF:FE0D:1E90-"to_b"
Local IF IP Addr : FE80::8E90:D3FF:FEBE:8F5A-"to_b"
Area Id          : 0.0.0.0
Designated Rtr   : 6.6.6.6      Backup Desig Rtr : 1.1.1.1
Neighbor State    : Full        Priority           : 1
Retrans Q Length : 0           Options            : --R--EV6
Events           : 6           Last Event Time    : 10/09/2012 13:43:08
Up Time          : 0d 00:17:21  Time Before Dead   : 34 sec
GR Helper        : Not Helping  GR Helper Age      : 0 sec
GR Exit Reason    : None        GR Restart Reason  : Unknown
Bad Nbr States   : 3           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 : 3.3.3.3 Interface: to_c
-----
Neighbor IP Addr : FE80::8E90:D3FF:FEAA:35F-"to_c"
Local IF IP Addr : FE80::8E90:D3FF:FEBE:8F5A-"to_c"
Area Id          : 0.0.0.0
Designated Rtr   : 1.1.1.1      Backup Desig Rtr : 3.3.3.3
Neighbor State    : Full        Priority           : 1
Retrans Q Length : 0           Options            : --R--EV6
Events           : 5           Last Event Time    : 10/09/2012 13:53:59
Up Time          : 0d 00:05:41  Time Before Dead   : 36 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
=====
*A:Dut-A#
```

Table 52: Output fields: OSPF/OSPF3 neighbor detail

Label	Description
Neighbor IP Addr	Displays the IP address this neighbor is using in its IP source address. On links with no addresses, 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.

Label	Description
Designated Rtr	Displays the IP Interface address of the router identified as the DR for the network in which this interface is configured. Set to 0.0.0.0 if there is no DR.
Neighbor Rtr Id	Displays a 32-bit integer uniquely identifying the neighboring router in the AS
Neighbor 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, a 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). Two Way — In this state, communication between the two routers is bidirectional. 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. Exchange — In this state the router is describing its entire link state database by sending database description packets to the neighbor. 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. Full — In this state, the neighboring routers are fully adjacent. These adjacencies now appear in router-LSAs and network-LSAs.
Priority	Displays the priority of this neighbor in the designated router election algorithm. A value of 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	E — external routes support MC — multicast support N/P — type 7 LSA support EA — external attribute LSA support DC — demand circuit support

Label	Description
	O — opaque LSA support
Backup Desig Rtr	Displays 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.
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	This value represents 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 because of 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 a LS update was received with an option mismatch since this interface was last enabled

Label	Description
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

opaque-database

Syntax

opaque-database [*link link-id* | *area area-id* *as*] [*adv-router router-id*] [*ls-id*] [*detail*]

Context

show>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays OSPF opaque database information.

Output

The following output is an example of OSPF opaque database information, and [Table 53: Output fields: OSPF opaque-database](#) describes the output fields.

Sample output

```
A:ALA-A# show router ospf opaque-database
=====
OSPF Opaque Link State Database (Type : All)
=====
Area Id      Type  Link State Id  Adv Rtr Id  Age  Sequence  Cksum
-----
0.0.0.0      Area  1.0.0.1       10.0.0.2    205  0x8000007e 0xb1b2
0.0.0.0      Area  1.0.0.1       10.0.0.5    617  0x80000084 0xb1a6
0.0.0.0      Area  1.0.0.1       10.0.0.8    1635 0x80000081 0xc391
0.0.0.0      Area  1.0.0.1       10.0.0.9    1306 0x80000082 0xc58c
0.0.0.0      Area  1.0.0.1       10.0.0.10   53   0x80000082 0xc986
0.0.0.0      Area  1.0.0.1       10.0.0.11   577  0x8000007e 0xd57c
0.0.0.0      Area  1.0.0.1       10.0.0.12   1628 0x80000080 0xd578
0.0.0.0      Area  1.0.0.1       10.0.0.13   581  0x80000080 0xd972
0.0.0.0      Area  1.0.0.1       10.0.0.22   1006 0x80000080 0xfd3c
0.0.0.0      Area  1.0.0.1       10.0.0.23   1238 0x80000083 0xfb39
0.0.0.0      Area  1.0.0.1       10.0.0.27   55   0x80000083 0xc21
0.0.0.0      Area  1.0.0.1       10.0.0.28   389  0x80000083 0x101b
0.0.0.0      Area  1.0.0.1       10.0.0.29   1658 0x80000082 0x1614
0.0.0.0      Area  1.0.0.1       10.0.0.30   976  0x80000083 0x180f
0.0.0.0      Area  1.0.0.2       10.0.0.2    45   0x800000a0 0x2f60
0.0.0.0      Area  1.0.0.2       10.0.0.5    1357 0x80000084 0x7038
0.0.0.0      Area  1.0.0.2       10.0.0.8    1960 0x80000084 0x3472
...
-----
No. of Opaque LSAs: 88
```

```

=====
A:ALA-A#

*A:Dut-A# show router ospf opaque-database adv-router 10.20.1.1 detail
=====
OSPF Opaque Link State Database (Type : All) (Detailed)
=====
-----
Opaque LSA
-----
Area Id       : 0.0.0.0          Adv Router Id   : 10.20.1.1
Link State Id  : 1.0.0.1          LSA Type        : Area Opaque
Sequence No    : 0x800000028      Checksum        : 0xb136
Age            : 192              Length           : 28
Options        : E
Advertisement   :
                ROUTER-ID TLV (0001) Len  4 : 10.20.1.1
-----
Opaque LSA
-----
Area Id       : 0.0.0.0          Adv Router Id   : 10.20.1.1
Link State Id  : 1.0.0.2          LSA Type        : Area Opaque
Sequence No    : 0x80000000d      Checksum        : 0x17f3
Age            : 678              Length           : 164
Options        : E
Advertisement   :
                LINK INFO TLV (0002) Len 140 :
                  Sub-TLV: 1      Len: 1      LINK_TYPE   : 2
                  Sub-TLV: 2      Len: 4      LINK_ID      : 10.10.1.2
                  Sub-TLV: 3      Len: 4      LOC_IP_ADDR   : 10.10.1.1
                  Sub-TLV: 4      Len: 4      REM_IP_ADDR   : 0.0.0.0
                  Sub-TLV: 5      Len: 4      TE_METRIC     : 1000
                  Sub-TLV: 6      Len: 4      MAX_BDWTH     : 100000 Kbps
                  Sub-TLV: 7      Len: 4      RSRVBL_BDWTH  : 800000 Kbps
                  Sub-TLV: 8      Len: 32     UNRSRVD_CLS0 :
                    P0: 80000 Kbps P1: 320000 Kbps P2: 320000 Kbps P3: 320000 Kbps
                    P4: 400000 Kbps P5: 400000 Kbps P6: 400000 Kbps P7: 80000 Kbps
                  Sub-TLV: 9      Len: 4      ADMIN_GROUP  : 0 None
                  Sub-TLV: 17     Len: 36     TELK_BW_CONST:
                    BW Model : MAM
                    BC0: 80000 Kbps BC1: 0 Kbps BC2: 320000 Kbps BC3: 0 Kbps
                    BC4: 0 Kbps BC5: 400000 Kbps BC6: 0 Kbps BC7: 0 Kbps
=====
*A:Dut-A#

```

Table 53: Output fields: OSPF opaque-database

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.
Type	NSSA — This area is configured as an NSSA area Area — This area is configured as a standard area (not NSSA or stub) Stub — This area is configured as an NSSA

Label	Description
Link State Id	Displays the link state ID is 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	Displays the router identifier of the router advertising the LSA
Age	Displays the age of the link state advertisement, in seconds
Sequence	Displays the signed 32-bit integer sequence number
Cksum	Displays the 32-bit unsigned sum of the link-state advertisements LS checksums

prefix-sids

Syntax

prefix-sids [*ip-prefix[/prefix-length]*] [*sid sid*] [**adv-router** *router-id*]

Context

show>router>ospf

Platforms

7210 SAS-Mxp

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 54: Output fields: prefix SIDs](#) describes the output fields.

Sample output

```
*A:Dut-F# show router ospf prefix-sids
```

Rtr Base OSPFv2 Instance 0 Prefix-Sids			
Prefix	Area Adv-Rtr	RtType	SID Flags
10.0.11.1/32	0.0.0.0	INTER-AREA	4
	10.20.1.2		NnP
10.0.11.1/32	0.0.0.1	INTRA-AREA	4
	10.20.1.1		NnP
10.0.11.1/32	0.0.0.1	INTRA-AREA	999
	10.20.1.3		NnPB
10.0.22.2/32	0.0.0.0	INTER-AREA	5
	10.20.1.2		NnPA
10.0.22.2/32	0.0.0.1	INTRA-AREA	5
	10.20.1.2		NnP
10.0.22.2/32	0.0.0.1	INTRA-AREA	996
	10.20.1.6		NnPB
10.0.33.3/32	0.0.0.0	INTER-AREA	0
	10.20.1.2		NnP
10.0.33.3/32	0.0.0.1	INTRA-AREA	0
	10.20.1.3		NnP
10.0.33.3/32	0.0.0.1	INTRA-AREA	998
	10.20.1.1		NnPB
10.0.44.4/32	0.0.0.0	INTRA-AREA	1
	10.20.1.4		NnP
10.0.44.4/32	0.0.0.0	INTRA-AREA	994
	10.20.1.5		NnPB
10.0.44.4/32	0.0.0.1	INTER-AREA	1
	10.20.1.2		NnP
10.0.55.5/32	0.0.0.0	INTRA-AREA	2
	10.20.1.5		NnP
10.0.55.5/32	0.0.0.0	INTRA-AREA	995
	10.20.1.4		NnPB
10.0.55.5/32	0.0.0.1	INTER-AREA	2
	10.20.1.2		NnP
10.0.66.6/32	0.0.0.0	INTER-AREA	3
	10.20.1.2		NnP
10.0.66.6/32	0.0.0.1	INTRA-AREA	3
	10.20.1.6		NnP
10.0.66.6/32	0.0.0.1	INTRA-AREA	997
	10.20.1.2		NnPB
10.20.1.1/32	0.0.0.0	INTER-AREA	10
	10.20.1.2		NnP
10.20.1.1/32	0.0.0.1	INTRA-AREA	10
	10.20.1.1		NnP
10.20.1.2/32	0.0.0.0	INTRA-AREA	11
	10.20.1.2		NnP
10.20.1.2/32	0.0.0.1	INTER-AREA	11
	10.20.1.2		NnPA
10.20.1.3/32	0.0.0.0	INTER-AREA	6
	10.20.1.2		NnP
10.20.1.3/32	0.0.0.1	INTRA-AREA	6
	10.20.1.3		NnP
10.20.1.4/32	0.0.0.0	INTRA-AREA	7
	10.20.1.4		NnP
10.20.1.4/32	0.0.0.1	INTER-AREA	7
	10.20.1.2		NnP
10.20.1.5/32	0.0.0.0	INTRA-AREA	8
	10.20.1.5		NnP
10.20.1.5/32	0.0.0.1	INTER-AREA	8
	10.20.1.2		NnP
10.20.1.6/32	0.0.0.0	INTRA-AREA	9
	10.20.1.6		NnP


```

10.20.1.6/32          0.0.0.1    INTER-AREA  9
                    10.20.1.2    NnP
-----
No. of Prefix/SIDs: 30
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
            B = Backup flag
=====
*A:Dut-F#

```

```

*A:Dut-C# show router ospf prefix-sids
=====
Rtr Base OSPFv2 Instance 0 Prefix-Sids
=====
Prefix                Area      RtType  SID
Adv-Rtr              Active   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 54: 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
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 55: 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 55: 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

Label	Description
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

sham-link

Syntax

sham-link [*interface-name*] [*detail*]

sham-link *interface-name* **remote** *ip-address* [*detail*]

Context

show>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information about the sham links.

Output

The following output is an example of sham link information, and [Table 56: Output fields: OSPF sham-link](#) describes the output fields.

Sample output

```
A:7210SAS# show router 1 ospf sham-link

=====
OSPF Sham-links
=====
If Name                               Nbr IP                               Metric Adm Oper
-----
to-ixia-1                             20.1.1.1                             1    Up   PToP
to-ixia-1                             111.11.1.1                           1    Up   PToP
-----
No. of OSPF Sham-links: 2
*A:7210SAS>show>router>ospf#
```

Table 56: Output fields: OSPF sham-link

Label	Description
If Name	Displays the IP Interface name
Nbr IP	Displays the IP address of the neighbor
Metric	Displays the metric associated with the interface
Adm	Displays the administrative state of the IP interface
Oper	Displays the operational state of the interface
No. of OSPF Sham-links	Displays the number of sham links configured

sham-link-neighbor

Syntax

sham-link-neighbor [**detail**]

sham-link-neighbor *interface-name* **remote** *ip-address* [**detail**]

Context

show>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information about the OSPF neighbor. The user is provided with an option to select the neighbor or the interface to restrict output to specific neighbors. The **detail** option displays a large amount of output.

Output

The following is an example of sham-link-neighbor information, and [Table 57: Output fields: OSPF sham-link neighbor](#) describes the output fields.

Sample output

```
A:duta# show router 1 ospf sham-link-neighbor

=====
OSPF Sham-link Neighbors
=====
Interface Name           Neighbor IP      State    RetxQ    DeadTi
                        me
-----
to-ixia-1                111.11.1.1     Full      0        34
-----
```

```
No. of Neighbors: 1
# show router 1 ospf sham-link-neighbor

=====
OSPF Sham-link Neighbors
=====
Interface Name           Neighbor IP      State    RetxQ    DeadTi
                        me
-----
to-ixia-1                111.11.1.1     Full     0        34
-----

No. of Neighbors: 1
*A:7210SAS>show>router>ospf#
```

Table 57: Output fields: OSPF sham-link neighbor

Label	Description
Interface Name	Displays the IP interface name
Neighbor IP	Displays the neighbor IP address
State	Specifies the operational state of the virtual link to the neighboring router
RetxQ	Displays the current length of the retransmission queue
DeadTime	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 OSPF Neighbors	Displays the number of adjacent OSPF neighbors on this interface

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 OSPF SPF information, and [Table 58: Output fields: OSPF SPF](#) 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 58: Output fields: OSPF SPF

Label	Description
Total SPF Runs	Displays the total number of incremental SPF runs triggered by new or updated LSAs
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

Label	Description
	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 statistics.

Output

The following output is an example of OSPF statistics information, and [Table 59: 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 Versions       : 0
Failed SPF Attempts: 0
CSPF Requests      : 0
CSPF Path Found    : 0
Total SPF Runs     : 1
Bad Options        : 0
Bad Checksums      : 0
CSPF Request Drops : 0
CSPF Path Not Found: 0
Total LFA SPF Runs : 1

```

A:ALA-A#

Sample Output for OSPF3

*A:Dut-A# show router ospf3 statistics

OSPF Statistics

```

Rx Packets      : 606
Rx Hellos       : 233
Rx DBDs         : 7
Rx LSRs         : 2
Rx LSUs         : 248
Rx LS Acks      : 116
New LSAs Recvd  : 0
Ext LSAs Count  : 0
No of Interfaces : 3
Retransmits     : 1
Bad Networks    : 0
Bad Areas       : 0
Bad Auth Types  : 0
Bad Neighbors   : 0
Bad Lengths     : 0
Bad Dead Int.   : 0
Bad Versions    : 0
Failed SPF Attempts: 0
CSPF Requests   : 0
CSPF Path Found : 0
Tx Packets      : 497
Tx Hellos       : 233
Tx DBDs         : 6
Tx LSRs         : 2
Tx LSUs         : 240
Tx LS Acks      : 16
New LSAs Orig   : 30
No of Areas     : 1
No of Neighbors : 2
Discards        : 76
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
CSPF Request Drops : 0
CSPF Path Not Found: 0

```

*A:Dut-A#

Table 59: 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

Label	Description
Tx DBDs	Displays the total number of OSPF database description packets transmitted on all OSPF enabled interfaces
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 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 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 for 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

Label	Description
Bad Dest Addr	Displays the total number of OSPF packets received on all OSPF-enabled interfaces with the incorrect IP destination address
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 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
Total SPF Runs	Displays the number of times the SPF algorithm has been run to calculate the best path to a destination
Total LFA SPF Runs	Displays the number of times the SPF algorithm has been run to calculate the LFA (backup path to a destination)

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.

Output

The following output is an example of OSPF status information, and [Table 60: 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#

*A:ALU_SIM11>show>router>ospf# status

=====
OSPF Status
=====
OSPF Cfg Router Id       : 0.0.0.0
OSPF Oper Router Id     : 1.1.1.2
OSPF Version             : 2
OSPF Admin Status       : Enabled
OSPF Oper Status        : Enabled
Graceful Restart         : Disabled
GR Helper Mode          : Disabled
Preference               : 10
External Preference     : 150
Backbone Router         : True
Area Border Router      : False
AS Border Router        : False
Opaque LSA Support      : True
Traffic Engineering Support : False
RFC 1583 Compatible     : True
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          : 1000 msec
Sec SPF Delay           : 1000 msec
Max SPF Delay           : 10000 msec
Min LS Arrival Interval : 1000 msec
Init LSA Gen Delay      : 5000 msec
Sec LSA Gen Delay       : 5000 msec
Max LSA Gen Delay       : 5000 msec
Last Ext SPF Run        : Never
Ext LSA Cksum Sum       : 0x0
OSPF Last Enabled       : 07/06/2010 10:34:11
Multicast Import        : False
Export Policies         : None
Import Policies         : None
Lfa Policies            : pol1
                        : pol2
                        : pol3
                        : pol4
                        : pol5
OSPF Ldp Sync Admin Status : Enabled
LDP-over-RSVP           : Disabled
=====
```

Sample Output for OSPF3

```
*A:Dut-A# show router ospf3 status

=====
OSPF Status
=====
OSPF Cfg Router Id       : 0.0.0.0
OSPF Oper Router Id     : 1.1.1.1
OSPF Version             : 3
OSPF Admin Status       : Enabled
```

```

OSPF Oper Status      : Enabled
Graceful Restart      : Disabled
GR Helper Mode        : Disabled
Preference            : 10
External Preference   : 150
Backbone Router       : True
Area Border Router    : False
AS Border Router      : False
Traffic Engineering 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        : 1000 msec
Sec SPF Delay         : 1000 msec
Max SPF Delay         : 10000 msec
Min LS Arrival Interval : 1000 msec
Init LSA Gen Delay    : 5000 msec
Sec LSA Gen Delay     : 5000 msec
Max LSA Gen Delay     : 5000 msec
Last Ext SPF Run      : Never
Ext LSA Cksum Sum     : 0x0
OSPF Last Enabled     : 10/09/2012 13:41:23
Multicast Import      : False
Export Policies       : None
OSPF Ldp Sync Admin Status : Enabled
LDP-over-RSVP        : Disabled
RSVP-Shortcut        : Disabled
Advertise-Tunnel-Link : Disabled
Export Limit         : 0
Export Limit Log Percent : 0
Total Exp Routes      : 0

```

Table 60: Output fields: OSPF status

Label	Description
OSPF Router Id	Displays a 32-bit integer uniquely identifying the router in the AS. By default, this is the system IP address or, if not configured, this is 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 — the OSPF process is disabled on all interfaces Enabled — the OSPF process is active on at least one interface
OSPF Oper Status	Disabled — the OSPF process is not operational on all interfaces Enabled — 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

Label	Description
Backbone Router	False — indicates that this router is not configured as an OSPF back bone router True — 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 AS border router True — this router is configured as an AS 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 virtual links.

Parameters

detail

Provides 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 61: 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 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
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 Addrs : 0
Bad Auth Types : 0             Auth Failures  : 0
Bad Neighbors  : 0             Bad Pkt Types  : 0
```



```

Bad Lengths      : 0
Bad Dead Int.    : 0
Retrans Intrvl   : 5 sec
Last Event       : 11/07/2006 17:12:00
=====
Bad Hello Int.   : 0
Bad Options      : 0
Transit Delay    : 1 sec
Authentication    : MD5
=====
A:ALA-A#

```

Table 61: 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
Authentication	Specifies whether authentication is enabled for the interface or virtual link
Hello Intrvl	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

Label	Description
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 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
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

Label	Description
	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
Retrans Intrvl	Displays the length of time, in seconds, that OSPF waits before retransmitting an unacknowledged LSA to an OSPF neighbor
Transit Delay	Displays the time, in seconds, that it takes to transmit an 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*

Specifies the router ID. This reduces the amount of output displayed.

detail

Displays detailed information about the virtual neighbor. This option produces a large amount of data. Nokia recommends using the detail keyword only when requesting information for a specific neighbor.

Output

The following output is an example of OSPF virtual neighbor information, and [Table 62: 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 62: Output fields: OSPF virtual-neighbor

Label	Description
Nbr IP Addr	Displays the IP address this neighbor is using in its IP source address. On links with no addresses, this will not be 0.0.0.0, but the address of another of the neighbor's interfaces.
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
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 because of a resource allocation issue since the OSPF admin status was last enabled
Bad Seq Num	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

Label	Description
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

4.12.2.3 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.

Parameters

ospf-instance

Clears the specified OSPF instance.

Values 0 to 31

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

Keyword that 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 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 neighbor, router, interface, SPF, and global statistics for this OSPF instance.

4.12.2.4 OSPF debug commands

ospf

Syntax

```
ospf
```

Context

```
debug>router
```


Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command indicates the OSPF instance for debugging purposes.

Parameters

ospf-instance

Specifies the OSPF instance.

Values 0 to 31

area

Syntax

area [*area-id*]

no area

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 area.

Parameters

area-id

Specifies the OSPF area ID, expressed in dotted decimal notation or as a 32-bit decimal integer.

Values 0 to 4294967295

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 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 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 OSPF and OSPF3 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 and OSPF3 interface.

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 (#, \$, spaces, and so on), 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 link-state database (LSDB).

Parameters

type

Specifies the OSPF 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 a 32-bit integer uniquely identifying an area.

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 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 OSPF3 neighbor.

Parameters

ip-int-name

Specifies the neighbor interface name.

ip-address

Specifies 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 packets.

Parameters

detail

Displays detailed packet information.

egress
Keyword that specifies an egress packet.

ingress
Keyword that specifies an ingress packet.

interface-name
Specifies the interface name to debug.

Values

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

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 RTM.

Parameters
ip-address
Specifies the IP address to debug.

Values ipv4-address: a.b.c.d

spf

Syntax

spf [*type*] [*dest-addr*]
no spf

Context

debug>router>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for OSPF SPF. Information about overall SPF start and stop times is displayed. To display detailed information about 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

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables debugging for an OSPF virtual neighbor.

Parameters

ip-address

Specifies the IP address of the virtual neighbor.

5 IS-IS

This chapter provides information to configure Intermediate System to Intermediate System (IS-IS).



Note:

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

5.1 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 15: IS-IS routing domain



OSRG033

5.1.1 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 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 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 LSPDUs 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.

5.1.2 IS-IS frequently used terms

- **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**

The subset of the local routing information base pertinent to a single local Subnetwork Point of Attachments (SNPAs).

- **link**

The communication path between two neighbors. A link is up when communication is possible between the two SNPAs.

- **designated 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 (instead of $n-1$ 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 multi-access subnetwork that supports the capability of addressing a group of attached systems with a single PDU.

- **general topology sub-network**

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**

Level 2 sub-domain is the set of all Level 2 intermediate systems in a routing domain.

5.1.3 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**

A variable length field between 1 and 13 bytes long. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.

- **system ID**

A six-byte system identification. 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.

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 16: Using area addresses to form adjacencies



5.1.3.1 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.

5.1.3.2 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.

5.1.4 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 is the smallest metric of all the more specific IPv4 routes

5.1.5 IS-IS multi-topology for IPv6

IS-IS IPv6 TLVs for IPv6 routing is supported in 7210 SAS. This 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 the 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

5.1.6 IS-IS administrative tags

IS-IS administrative 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

5.1.6.1 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

5.1.6.2 Using route tags

The IS-IS administrative tags configured on an IS-IS router (or neighbor) do not have an 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
```

5.1.7 Segment routing in Shortest Path Forwarding



Note:

This feature is supported only on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode).

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 and LDP VPLS
- BGP-AD VPLS when the **use-provisioned-sdp** option is enabled in the binding to the PW template
- EVPN-VPLS and EVPN-VPWS (only on platforms that support these services)
- 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:

- The 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE do not support remote LFA when services use BGP 3107 labeled route tunnels. The push stack depth in this case exceeds the allowed limit.
- The EXP-to-profile mapping for SR is defined in the **mpls-lsp-exp-profile-map** policy.
- On the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE Nokia recommends the use of BGP3107 services using SR to advertise the loopback for BGP 3107 labeled routes to the IGP. This prevents a three-label pop on egress LER.
- On the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE when using RLFA with services (and without the use of BGP 3107), if the PQ node is the segment termination, the SR OS always uses the PQ node SID and does not use additional SIDs. Therefore, when the 7210 SAS is configured as an eLER, it always requires two labels to pop and terminate the service packet.
- On the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE the maximum label push depth is three MPLS labels, and the maximum label pop depth is two MPLS labels (both push and pop exclude the PW hash label).

5.1.7.1 Segment routing operational procedures

5.1.7.1.1 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 more 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 interpret 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 datapath.
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 datapath 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 datapath; 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.

5.1.7.1.2 Error and resource exhaustion handling


When the prefix corresponding to a node SID is being resolved, the following procedures are followed.

5.1.7.1.2.1 Procedure 1: Providing support for multiple topologies for the same destination prefix

The 7210 SAS supports assigning different prefix-SID indexes and labels to the same prefix in different IGP instances. While other routers that receive these prefix SIDs program a single route into the RTM, based on the winning instance ID as per RTM route type preference, the 7210 SAS adds two tunnels to this destination prefix in the TTM. This provides support for multiple topologies for the same destination prefix.

For example, in two different instances (L2, IS-IS instance 1 and L1, IS-IS instance 2 — see the following figure), Router D has the same prefix destination with different SIDs (SIDx and SIDy).

Figure 17: Programming multiple tunnels to the same destination



Assume 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 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:
 - prefix N
 - nhop1 = B
 - nhop2 = C
 - preference 18
 - For prefix N, the SR tunnel TTM entry is:
 - 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 C only.
 - For prefix N, the TTM entry is:
 - prefix N
 - nhop1 = B
 - preference 15RTM prefers the L1 route over the L2 route
 - For prefix N, there are two SR tunnel entries in the TTM:
 - SR entry for L2:
 - tunnel-id 1: prefix N-SIDx
 - nhop1 = B
 - nhop2 = C
 - tunl-pref 11 (tunl-pref 10 for OSPF)
 - SR entry for L1:
 - tunnel-id 2: prefix N-SIDy

5.1.7.1.2.2 Procedure 2: 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 system 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 system 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 toward all ABRs uses the selected SID.



Note:

The 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode) supports only LSR ECMP. It does not support LER ECMP.

5.1.7.1.2.3 Procedure 3: Checking for SID error before 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.

5.1.7.1.2.4 Procedure 4: 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 datapath
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 datapath
 - 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.


5.1.7.1.2.5 Procedure 5: 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 datapath.

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 18: 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 L1 route over the L2 route.
 - For prefix N, there is one SR tunnel entry for L2 in the TTM:
 - tunnel-id 1: prefix N-SIDx
 - nhop1 = B
 - nhop2 = C
 - tunl-pref 11 (tunl-pref 10 for OSPF)

5.1.7.1.2.6 Procedure 6: 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.

5.1.7.1.2.7 Procedure 7: 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.

5.1.7.2 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 datapath 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 is 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.



Note:

The 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE support only LSR ECMP; LER ECMP is not supported on these platforms.

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.

```
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 BPG 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.

5.1.7.2.1 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:

```
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:

$$\text{SR\_Tunnel\_MTU} = \text{MIN} \{ \text{Cfg\_SR\_MTU}, \text{IGP\_Tunnel\_MTU} - (1 + \text{frr-overhead}) * 4 \}$$

Where:

- Cfg_SR_MTU is the MTU configured by the user for all SR tunnels within a specific IGP instance using 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.

5.1.7.3 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 19: 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 19: 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 19: 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 20: Remote LFA next-hop in segment routing



al_0648

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** CLI 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.

5.1.7.4 Datapath 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 63: Datapath 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 datapath 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. • The BGP label may be swapped to another BGP label.
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.



Note:

Four label push is not supported on the 7210 SAS-Mxp.

5.1.7.4.1 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-Mxp, 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-Sx/S 1/10GE hash labels are supported only with specific services. See the *7210 SAS-Mxp, S, Sx, T Services Guide* for more information about the services supported with hash label.

5.1.7.5 Control protocol changes

This section describes the [IS-IS control protocol changes](#) and [OSPF control protocol changes](#).

5.1.7.5.1 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 toward 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 7210 SAS 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.

5.1.7.5.2 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 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 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.

5.1.7.6 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.



Note:

The 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode) do not support RFLA with BGP 3107 labeled tunnel routes.

```
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.

5.1.7.7 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 does 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 (**tl dp**), 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 see the *7210 SAS-Mxp, S, Sx, T Services Guide* for more information about the VPRN **auto-bind-tunnel** CLI command.

The following service contexts are supported with SR tunnels:

- VLL and LDP VPLS
- 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

5.1.7.8 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.

5.1.8 IGP-LDP synchronization

The 7210 SAS supports IGP-LDP synchronization on IS-IS routes. See to the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide* for more information.



Note:

Static route-LDP synchronization is supported on all 7210 SAS platforms as described in this document.

5.1.9 IS-IS import policy on the 7210 SAS-Mxp

IS-IS import policies block routes and prevent them from being installed in the routing table, but do not prevent the prefixes from being propagated in Link State PDUs (LSPs).



Caution:

Nokia recommends that the user must exercise caution when configuring this feature. Failure to do so may cause unintended connectivity loss (black holes).

The following match conditions are supported for IS-IS import policies:

- **from level**
- **from prefix-list**
- **from protocol isis** [**all** | **instance** *instance*]
- **from tag**

The following actions are supported with IS-IS import policies:


- **accept** (installs the matched prefixes into the routing table)
- **reject** (blocks the matched prefixes from being added to the routing table)
- **next-policy**

Actions that modify the metric value, RTM preference, or route tag value of an accepted IS-IS route may be present in a route policy but are ignored when that policy is applied as an IS-IS import policy.

5.2 IS-IS configuration process overview

The following figure shows the process to provision basic IS-IS parameters.

Figure 21: IS-IS configuration and implementation flow



sw0491

5.3 Configuration notes

The following describes IS-IS configuration restrictions:

- 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.

5.4 Configuring IS-IS with CLI

This section provides information to configure intermediate-system-to-intermediate-system (IS-IS) using the command line interface.

5.5 IS-IS configuration overview

5.5.1 Router levels

The router level capability can be configured globally and on a per-interface basis. The interface-level parameters specify the interface 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.

5.5.2 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 long. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.

- **system ID**

A six-byte system identification. 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 example displays 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

5.5.3 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 describes configuration combinations and the potential adjacencies that can be formed.

Table 64: 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

5.5.4 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 specific 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.

5.6 Basic IS-IS configuration

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:

- Enable IS-IS (specifying the instance ID of multi-instance IS-IS is to be enabled).
- If necessary, modify the level capability on the global level (default is level-1/2).
- Define area addresses.
- Configure IS-IS interfaces.

Example: IS-IS default values configuration output

```
A:Dut-A>config>router>isis$ info detail
-----
    level-capability level-1/2
    no graceful-restart
    area-id 01
    no authentication-key
    no authentication-type
    authentication-check
    csnp-authentication
    lsp-lifetime 1200
    no export
    hello-authentication
    psnp-authentication
    traffic-engineering
    no reference-bandwidth
    no disable-ldp-sync
    ipv4-routing
spf-wait 10 1000 1000
lsp-wait 5 0 1
level 1
```



```
no authentication-key
no authentication-type
csnp-authentication
external-preference 160
hello-authentication
preference 15
psnp-authentication
no wide-metrics-only
exit
level 2
no authentication-key
no authentication-type
csnp-authentication
external-preference 165
hello-authentication
preference 18
psnp-authentication
no wide-metrics-only
exit
no shutdown
-----
A:Dut-A>config>router>isis$
```

5.7 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 addresses.
- Configure IS-IS interface-specific parameters.

5.8 Configuring IS-IS components

The following section describes the syntax used to configure the IS-IS components.

5.8.1 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.

The following shows the command usage to configure IS-IS on a router.

```
isis
```

Example

```
config>router# isis
```

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 to the **config>router>isis** command.

5.8.2 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.

Example

The following shows the command usage to configure the router level.

```
config>router# isis
  level-capability {level-1|level-2|level-1/2}
  level {1|2}
```

```
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#
```

5.8.3 Configuring ISO area addresses

Use the following syntax to configure an area ID also called an address. A maximum of three area IDs can be configured.

```
config>router# isis
      area-id area-address
```

Example: command usage to configure the router 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#
```

5.8.4 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 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 "H5KBAAWAAQU" hash
      authentication-type password
      overload timeout 90
      traffic-engineering
-----
A:ALA-A>config>router>isis#
```

5.8.5 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 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

```
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
config>router>isis>if# mesh-group 101
config>router>isis>if# exit
config>router>isis#
```

Example: Global and interface-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
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#

```

5.8.5.1 Example: configuring a Level 1 area




Note:

Interfaces are configured in the **configure router interface** context.

The following figure shows the configuration of a Level 1 area.

Figure 22: 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#
```

5.8.5.2 Example: modifying a router level capability

In the previous example, 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 a Level 1/2 area.

Figure 23: Configuring a Level 1/2 area



Example

The following shows the 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
```

5.9 IS-IS configuration management tasks

This section describes the IS-IS configuration management tasks.

5.9.1 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
```

5.9.2 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
```

5.9.3 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

The following shows the 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#

```

5.9.4 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

The following shows the command usage for 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

```

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#
```

5.9.5 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 about only 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 about 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.

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, apply the policy to leak routes from Level 2 into 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#

```

Example

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, 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
```

Example

```
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#
```

5.9.6 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 [Route policies](#) for more information.

Example

The following is a sample 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#
```

5.10 IS-IS command reference

- [Command hierarchies](#)
- [Command descriptions](#)

5.10.1 Command hierarchies

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

5.10.1.1 Configuration commands

5.10.1.2 Global commands

```
config
- router
- [no] isis [isis-instance]
- [no] advertise-passive-only
- advertise-router-capability {area | as}
- no advertise-router-capability
- [no] all-l1isis ieee-address
- [no] 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] 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
- [no] ignore-attached-bit
- import policy-name [policy-name...(up to 5 max)]
- no import
- [no] ipv4-routing
- [no] ipv6-routing {native | mt}
- ldp-over-rsvp
- no ldp-over-rsvp
- loopfree-alternate [remote-lfa]
- loopfree-alternate remote-lfa [max-cq-cost value]
- no loopfree-alternate
- loopfree-alternate-exclude
- no loopfree-alternate-exclude
- level {1 | 2}
-   authentication-key [authentication-key | hash-key] [hash | hash2]
-   no authentication-key
-   authentication-type {password | message-digest}
-   no authentication-type
-   [no] csnp-authentication
-   [no] default-ipv6-unicast-metric 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
- lsp-refresh-interval [seconds]
- no lsp-refresh-interval
- [no] lsp-wait lsp-wait [lsp-initial-wait [lsp-second-wait]]
- multi-topology
- no multi-topology
-   ipv6-unicast
-   no ipv6-unicast
- overload [timeout seconds]
- no overload
- overload-on-boot [timeout seconds]
- no overload-on-boot
- [no] psnp-authentication
- reference-bandwidth reference-bandwidth
- no reference-bandwidth
- 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
- [no] spf-wait spf-wait [spf-initial-wait [spf-second-wait]]
- [no] strict-adjacency-check
- summary-address {ip-prefix/mask | ip-prefix [netmask]} level [tag tag]
- no summary-address {ip-prefix/mask | ip-prefix [netmask]}
- [no] traffic-engineering
```

5.10.1.3 Interface commands

```
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
- ipv6-unicast-disable
- no ipv6-unicast-disable
- 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 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}
- lsp-pacing-interval milli-seconds
- no lsp-pacing-interval
- mesh-group [value | blocked]
- no mesh-group
- [no] multi-topology
- [no] ipv6-unicast
- 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
```

5.10.1.4 Show commands

```
show
- router
  - isis all
  - isis [isis-instance]
    - adjacency [ip-address | ip-int-name | nbr-system-id] [detail]
    - database [system-id | lsp-id] [detail] [level level]
    - hostname
    - interface [ip-int-name | ip-address] [detail]
    - 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 [detail]
    - spf-log [detail]
    - statistics
    - status
    - summary-address [ip-prefix[/prefix-length]]
    - topology [[ipv4-unicast | ipv6-unicast | mt mt-id-number] [detail]]
```

5.10.1.5 Clear commands

```
clear
- router
  - isis [isis-instance]
    - adjacency [system-id]
    - database [system-id]
    - export
    - spf-log
    - statistics
```

5.10.1.6 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
    - [no] spf [level-number] [system-id]
```


5.10.2 Command descriptions

- [IS-IS configuration commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

5.10.2.1 IS-IS configuration commands

5.10.2.1.1 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 number of IS-IS instances supported on different 7210 platforms are different. Contact a Nokia representative about the supported scaling limits.

Parameters

isis-instance

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



Note:

The **config>router>isis>segment-routing** context is supported only on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode).

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables an entity.

Default

no shutdown

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 Protocol Handling

On all 7210 SAS platforms, IS-IS is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

- The **config>router>isis** command instantiates the protocol in the **no shutdown** state, and resources are allocated to enable the node to process the protocol.

- To deallocate resources, you must issue the **configure router isis shutdown** and **configure router no isis** commands to allow the node to boot up correctly after the reboot. It is not sufficient to issue only a **configure router isis shutdown** command.
- The resources are allocated when the first instance of IS-IS is configured, and resources are deallocated when the last instance of the configuration for IS-IS is removed or shut down.

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**).

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.

Values 1 to 4294967295

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 configures 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 at the global level, configure this command in the **config>router>isis** context. When this parameter is configured at 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

Keyword to specify 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

Keyword to specify 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 in 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 entered.

Configure the authentication type at the global level in the **config>router>isis** context. Configure or override the global setting by configuring the authentication type using the **config>router>isis>level** context

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

Default

no authentication-type

Parameters

password

Keyword to specify that simple password (plain text) authentication is required.

message-digest

Keyword to specify 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-Mxp, R6, R12, S, Sx, T Router Configuration Guide*.

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

Default

no bfd-enable ipv4

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 IS-IS packets of 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.

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

Default

csnp-interval 10 — CSN PDUs are sent every 10 seconds for LAN interfaces.

csnp-interval 5 — CSN PDUs are sent every 5 seconds for point-to-point interfaces.

Parameters

seconds

Specifies the time interval, in seconds, between successive CSN PDUs sent from this interface, expressed as a decimal integer.

Values 1 to 65535

default-ipv6-unicast-metric

Syntax

default-ipv6-unicast-metric *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.

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 **ipv6-unicast-metric** CLI command 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

metric

Specifies the metric assigned for this level on this interface.

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 OSPF or IS-IS routing protocol.

When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces which have the IGP-LDP synchronization enabled if the currently advertised cost is different. It 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.

The **no** form of this command reverts to the default value and re-enables IGP-LDP synchronization on all interfaces participating in the OSPF or IS-IS routing protocol and for which the **ldp-sync-timer** is configured.

Default

no disable-ldp-sync

export

Syntax

[no] **export** *policy-name* [*policy-name*...up to 32 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 five policy names can be specified.

export-limit

Syntax

export-limit *number* [**log** *percentage*]

no export-limit

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 IS-IS 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 are 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. This command configures the preference level of either IS-IS level 1 or IS-IS level 2 external routes.

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 to use.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is based on the default preferences. The following table lists the default preferences.

Table 65: 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 ¹⁴
IS-IS Level 2 internal	18	Yes ¹⁴
OSPF external	150	Yes
IS-IS Level 1 external	160	Yes
IS-IS Level 2 external	165	Yes
BGP	170	Yes

If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of the route to use is determined by the configuration of the **ecmp** in the **config>router** context.

Parameters

external-preference

Specifies the preference for external routes at this level, expressed as a decimal integer.

Values 1 to 255

¹⁴ Change internal preferences using the **preference** command in the **config>router>isis>level** context.

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 acts 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, then 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

no graceful-restart

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 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 does not act as a restarting router (meaning that the router does not help the neighbors to restart).

The **no** form of this command enables helper support and is the default when the **graceful-restart** command is enabled.

Default

no helper-disable

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 from the SPF LFA computation. The LFA SPF calculation can therefore be run only where it is not needed.

If an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2.

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

Default

no loopfree-alternate-exclude

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

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 packets of the Hello type.

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 Instance Identifier (IID) TLV is 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 the hello authentication type on a segment must match. The **hello-authentication-type** must be specified.

To configure the hello authentication key in the interface context use the **hello-authentication-key** command in the **config>router>isis>interface** context.

To configure or override the hello authentication key for a specific level, configure the **hello-authentication-key** in the **config>router>isis>interface>level** context.

If both IS-IS and hello authentication are configured, hello messages are validated using hello authentication. If only IS-IS authentication is configured, it is used to authenticate all IS-IS (including hello) protocol PDUs.

When the hello authentication key is configured in the **config>router>isis>interface** context, it applies to all levels configured for the interface.

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

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 352 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

Keyword to specify 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

Keyword to specify 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 **hello-authentication-type** in the **config>router>isis>interface** context.

To configure or override the hello authentication setting for a specific level, configure the **hello-authentication-type** in the **config>router>isis>interface>level** context.

The **no** form of this command disables hello authentication.

Default

no hello-authentication-type

Parameters

password

Keyword to specify the simple password (plain text) authentication is required.

message-digest

Keyword to specify that MD5 authentication in accordance with RFC2104 (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 to 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 > 3 to reduce routing instability.

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

Default

hello-multiplier 3

Parameters

multiplier

Specifies the multiplier for the hello interval, expressed as a decimal integer.

Values 2 to 100

ipv6-unicast-metric

Syntax

ipv6-unicast-metric *metric*

no 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.

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

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

Commands in this context configure an IS-IS interface.

When an area is defined, the interfaces belong to that area. Interfaces cannot belong to separate areas.

When the interface is a POS channel, the OSINCP is enabled when the interface is created and removed when the interface is deleted.

The **no** form of this command removes IS-IS from the interface.

The **shutdown** command in the **config>router>isis>interface** context administratively disables IS-IS on the interface without affecting the IS-IS configuration.

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.

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 IS-IS interface type as either broadcast or point-to-point.

Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the designated IS-IS overhead if the link is used as a point-to-point.

If the interface type is not known at the time the interface is added to IS-IS, and subsequently the IP interface is bound (or moved) to a different interface type, this command must be entered manually.

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

Default

interface-type point-to-point — for IP interfaces on SONET channels

interface-type broadcast — for IP interfaces on Ethernet or unknown type physical interfaces

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

Keyword to configure the interface to maintain this link as a broadcast network.

point-to-point

Keyword to configure the interface to maintain this link as a point-to-point link.

ipv6-unicast-disable

Syntax

[no] ipv6-unicast-disable

Context

config>router>isis>interface

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command disables IS-IS IPv6 unicast routing for the interface.

By default, IPv6 unicast on all interfaces is enabled. However, IPv6 unicast routing on IS-IS is in effect when the **config>router>isis>ipv6-routing mt** command is configured.

The **no** form of this command enables IS-IS IPv6 unicast routing for the interface.

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

Keyword that enables IS-IS IPv6 TLVs for IPv6 routing and enables support for native IPv6 TLVs.

mt

Keyword that enables IS-IS multi-topology TLVs for IPv6 routing. When this parameter is specified, the support for native IPv6 TLVs is disabled.

ldp-over-rsvp

Syntax

[no] **ldp-over-rsvp**

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables LDP over RSVP processing in IS-IS.

The **no** form of this command disables LDP over RSVP processing.

Default

no ldp-over-rsvp

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 or interface level parameters to the default, the following commands must be entered independently:

- **level>no hello-authentication-key**
- **level>no hello-authentication-type**
- **level>no hello-interval**
- **level>no hello-multiplier**
- **level>no metric**
- **level>no passive**
- **level>no priority**

Default

level 1 or level 2

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/or 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 level 2.
The following table displays configuration combinations and the potential adjacencies that can be formed.

Table 66: Potential adjacency capabilities

Global level	Interface level	Potential adjacency
L 1/2	L 1/2	Level 1 and/or Level 2

Global level	Interface level	Potential adjacency
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

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

Keyword to specify the router or interface can operate at level 1 only.

level-2

Keyword to specify the router or interface can operate at level 2 only.

level-1/2

Keyword to specify the router or 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 specifies the time interval that LSPs originated by the router are considered valid by other routers in the domain.

Each LSP received is maintained in an LSP database until the LSP lifetime expires, unless the originating router refreshes the LSP. By default, each router refreshes its LSPs every 20 minutes (1200 seconds), to ensure that other routers do not age out the LSP.

The LSP refresh timer is derived using the following formula: **lsp-lifetime** value / 2

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

Default

lsp-lifetime 1200

Parameters

seconds

Specifies the interval, in seconds, that LSPs originated by the router are considered valid by other routers in the domain.

Values 350 to 65535

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 MTU size is changed from the default using the CLI or SNMP, IS-IS must be restarted for the change to take effect. This can be done by performing a **shutdown** command and then a **no shutdown** command in the **config>router>isis** context.



Note:

If the MTU size is changed from the default value by using the **exec** command to execute a configuration file with the changed value, IS-IS automatically bounces before the change takes 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-refresh-interval

Syntax

lsp-refresh-interval [*seconds*]
no lsp-refresh-interval

Context

config>router>isis

Description

This command configures the IS-IS LSP refresh timer interval. The value specified for **lsp-lifetime** must be considered when configuring the **lsp-refresh-interval**. The LSP refresh interval cannot be greater than 90% of the LSP lifetime.

The **no** form of this command reverts to the default (600 seconds); however, if the configured value is greater than 90% of the LSP lifetime, the command is rejected. For example, if the LSP lifetime is 400, the **no lsp-refresh-interval** command is rejected.

Default

lsp-refresh-interval 600

Parameters

seconds

Specifies the refresh interval.

Values 150 to 65535

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 customizes the throttling of IS-IS LSP-generation. Timers that determine when to generate the first, second, and subsequent LSPs can be controlled using this command.

Subsequent LSPs are generated at increasing intervals of the *lsp-second-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 for each 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.



Caution:

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 this or any other router that is part of this mesh group.

Values 1 to 2000000000

blocked

Keyword that prevents 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 multi-topology support.

Default

no multi-topology

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.

The **no** form of this command disables multi-topology TLVs.

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

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 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

7210 SAS-Mxp

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** parameters 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.

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 layer 1 IS-IS routers. The MAC address should be a multicast address. Run the **/no shutdown** command on the IS-IS instance to make the change operational.

Default

no all-l1isis

(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.)

Parameters

ieee-address

Specifies the destination MAC address for all layer 1 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 layer 2 IS-IS routers. The MAC address should be a multicast address. Run the **shutdown/no shutdown** command on the IS-IS instance to make the change operational.

Default

no all-l2isis

(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.)

Parameters

ieee-address

Specifies the destination MAC address for all layer 2 IS-IS neighbors on the link for this IS-IS 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 configures the area ID portion of NSAP addresses, which identifies a point of connection to the network, such as a router interface, and is called a Network Service Access Point (NSAP). Addresses in the IS-IS protocol are based on the ISO NSAP addresses and Network Entity Titles (NETs), not IP addresses. This command was previously named the **net network-entity-title** command.

A maximum of 3 area addresses can be configured.

NSAP addresses are divided into the following parts (only the area ID is configurable):

- **Area ID**

This is a field of variable length, between 1 and 13 bytes. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.

- **System ID**

This is a six-byte system identification. This value is not configurable. The system ID is derived from the system or router ID.

- **Selector ID**

This 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.

The NET is constructed like an NSAP, but the selector byte contains a 00 value. NET addresses are exchanged in hello and LSP PDUs. All net addresses configured on the node are advertised to its neighbors.

For level 1 interfaces, neighbors can have different area IDs, but they must have at least one area ID (AFI + area) in common. Sharing a common area ID, they become neighbors and area merging between the potentially different areas can occur.

For level 2 (only) interfaces, neighbors can have different area IDs. However, if they have no area IDs in common, they become only level 2 neighbors, and level 2 LSPs are exchanged.

For level 1 and level 2 interfaces, neighbors can have different area IDs. If they have at least one area ID (AFI + area) in common, they become neighbors. In addition to exchanging level 2 LSPs, area merging between potentially different areas can occur.

If multiple **area-id** commands are entered, the system ID of all subsequent entries must match the first area address.

The **no** form of this command removes the area address.

Parameters

area-address

Specifies 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 or, at most, 13 bytes. If less than 13 bytes are entered, the rest is padded with zeros.

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 or indefinitely.

During normal operation, the router may be forced to enter an overload state because of a lack of resources. When in the overload state, the router is only used if the destination is reachable by the router and is not used for other transit traffic.

If a time period is specified, the overload state persists for the configured length of time. If no time is specified, the overload state operation is maintained indefinitely.

The **overload** command 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 overload state.

Values 60 to 1800

Default infinity (overload state maintained indefinitely)

overload-on-boot

Syntax

overload-on-boot [timeout *seconds*]

no overload-on-boot

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

When 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 using the **config router isis no overload** command

The **no overload** command does not affect the **overload-on-boot** command function.

If no timeout is specified, IS-IS goes into overload indefinitely after a reboot. After the reboot, the IS-IS status displays a permanent overload state:

- Layer 1 LSDB Overload: Manual on boot (Indefinitely in overload)
- Layer 2 LSDB Overload: Manual on boot (Indefinitely in overload)

This state can be cleared using the **no overload** command.

When specifying a timeout value, IS-IS goes into overload for the configured timeout after a reboot. After the reboot, the IS-IS status displays the remaining time the system stays in overload:

- Layer 1 LSDB Overload: Manual on boot (Overload Time Left: 17)
- Layer 2 LSDB Overload: Manual on boot (Overload Time Left: 17)

The overload state can be cleared before the timeout expires using the **no overload** command.

Use the **show router ospf status** or **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 the router remains in the overload state after rebooting.

Values 60 to 1800

ipv4-node-sid

Syntax

ipv4-node-sid index *value*

ipv4-node-sid label *value*

no ipv4-node-sid

Context

config>router>isis>interface

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-Sx/S 1/10GE

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. 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 using the global mode of operation, a new segment routing module checks that the same index or label value is not assigned to more than one loopback interface address. When using the per-instance mode of operation, this check is not required because the index and 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

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, with which the interface is advertised as an IS-IS interface without running the IS-IS protocol. Normally, only interface addresses that are configured for IS-IS are advertised as IS-IS interfaces at the level that they are configured.

When the passive mode is enabled, the interface or the interface at the level ignores ingress IS-IS protocol PDUs and does not transmit IS-IS protocol PDUs.

The **no** form of this command removes the passive attribute.

Default

passive — service interfaces are passive
no passive — all other interfaces are not passive

Special Cases

Service Interfaces

Service interfaces (defined using the service-prefix command in the **config>router** context) 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.

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 is used.

Different protocols should not be configured with the same preference. If this occurs the tiebreaker is based on the default preferences listed in the following table.

Table 67: Route preference defaults by route type

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 ¹⁵
IS-IS level 2 external	165	Yes ¹⁵
BGP	170	Yes

If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs

¹⁵ External preferences are changed using the **external-preference** command in the **config>router>isis>level** context.

(metrics) are equal, the decision of what route to use is determined by the configuration of ECMP in the **config>router** context.

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

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.

This priority is included in hello PDUs transmitted by the interface on a multi-access network. The router with the highest priority is the preferred designated router. The designated router is responsible for sending LSPs with regard to this network and the routers that are attached to it.

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

Default

priority 64

Parameters

number

Specifies the priority for this interface at this level.

Values 0 to 127

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 configures 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 reference-bandwidth  
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:

$\text{cost} = \text{reference-bandwidth} \# \text{bandwidth}$

If the reference bandwidth is configured as 10 Gbytes (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 the [wide-metrics-only](#) command).

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

reference-bandwidth

Specifies the reference bandwidth in kilobits per second, expressed as a decimal integer.

Values 0 to 1000000000

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

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

Description

Commands in this context configure segment routing (SR) parameters within an IGP instance.

SR 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 SR 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.

SR 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 SR is configured in the IS-IS instance, the router performs the following operations.

1. Advertises the SR capability sub-TLV to routers in all areas and levels of this IGP instance. However, only neighbors with which it established an adjacency interpret the SID/label range information and use it to calculate the label to swap to or push for a specific 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 SR module then programs the incoming label map (ILM) with a pop operation for each local node SID in the datapath.
3. Assigns and advertises automatically an adjacency SID label for each formed adjacency over a network IP interface in the new adjacency SID sub-TLV. The SR 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. Resolves received prefixes, and if a prefix SID sub-TLV exists, the SR 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 SR is enabled 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

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

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 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 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 SR 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 SR context of an IGP instance to change the SID index or label range of that IGP instance using the **prefix-sid-range** command.

Any range change fails 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 SR 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

global

Keyword to enable global operation mode.

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

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

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, IGP determines the MTU.

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, the SR tunnel MTU is 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

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

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. Either use the global TTM preference or list the tunnel types to use. When listing the tunnel types, the TTM preference is used to select one type over the other. In both cases, a fall-back 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 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 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 defines the maximum interval between two consecutive SPF calculations.

Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled using this command. Subsequent SPF runs (if required) occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, the next SPF runs after 2000 milliseconds, and the next SPF runs 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 0 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 do not come up unless both routers have exactly the same address families configured. An existing adjacency that has unmatched address families is torn down. This command is used to prevent black-holing traffic when IPv4 and IPv6 topologies are different. When disabled (**no strict-adjacency-check**) a BFD session failure for either IPv4 or Ipv6 causes the routes for the other address family to be removed as well.

When 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 enables 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 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]

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 IS-IS tag to routes matching the entry.

Values	Accepts decimal or hex formats: 1 to 4294967295

OSPF and IS-IS:

[0x0 to 0xFFFFFFFF]H

ignore-attached-bit

Syntax

[no] ignore-attached-bit

Context

config>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures IS-IS to suppress the installation of default routes.

The **no** form of this command removes suppression of default route installation.

Default

no ignore-attached-bit

import

Syntax

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

no import

Context

config>router>isis

Platforms

7210 SAS-Mxp

Description

This command specifies up to five route polices as IS-IS import policies.

When a prefix received in an IS-IS LSP is accepted by an entry in an IS-IS import policy, it is installed in the routing table if it is the most preferred route to the destination.

When a prefix received in an IS-IS LSP is rejected by an entry in an IS-IS import policy, it is not installed in the routing table, even if it has the lowest preference value among all the routes to that destination.

The flooding of LSPs is unaffected by IS-IS import policy actions.

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

Default

no import

Parameters

policy-name

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

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.

The **no** form of this command disables traffic-engineered route calculations.

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 for the adjacency and the IP prefix.

By default, both sets of TLVs are generated. When the **wide-metrics-only** command 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.

5.10.2.2 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

Commands in this context display information for a specified IS-IS instance.



Note:
The number of IS-IS instances supported is different for each 7210 SAS platform. Contact a Nokia representative for information about the supported scaling limits.

Parameters

- all
Keyword to display information for all IS-IS instances on the specified router.
- instance-id
Specifies the instance ID for an IS-IS instance.

Values 0 to 31

Default 0

adjacency

Syntax

adjacency [ip-address | ip-int-name | 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. When no *ip-address*, *ip-int-name*, or *nbr-system-id* values are specified, all adjacencies display.

Parameters

ip-address

Displays only adjacencies with the specified interface.

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

ip-int-name

Displays only adjacencies with the specified interface.

nbr-system-id

Displays only the adjacency with the specified ID.

detail

Displays all output displays in the detailed format.

Output

The following output is an example of IS-IS neighbor adjacency information, and [Table 68: Output fields: IS-IS 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                MT Enab
-----
asbr_east                L2    Up    25   if2/5                            Yes
-----
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 68: Output fields: IS-IS adjacency

Label	Description
Interface	Displays the interface name associated with the neighbor
System-id	Displays the neighbor system ID
Level	1 — Layer 1 only; 2 — Layer 2 only; 3 — Layer 1 and Layer 2
State	Up, down, new, one-way, initializing, or rejected
Hold	Displays the hold time remaining for the adjacency

Label	Description
SNPA	Displays the subnetwork point of attachment, MAC address of the next hop
Circuit type	Displays the level on the interface: Layer 1, Layer 2, or both
Expires In	Displays the number of seconds until the 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 time since the last transition change
Speaks	Displays the 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 • MT#0, MT#2 => "Topology : Unicast" • Native IPv4 Not supported MTIDs => Topology line suppressed

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.

Parameters

system-id

Displays only the LSPs related to the specified *system-id*. If no *system-id* or *lsp-id* are specified, all database entries are displayed.

- lsp-id**
Displays only the specified LSP (hostname). If no *system-id* or *lsp-id* are specified, all database entries are displayed.
- detail**
Displays all output in the detailed format.
- level**
Displays only the specified IS-IS protocol level attributes.

Output

The following output is an example of IS-IS link state database information, and [Table 69: Output fields: 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
Sequence    : 0x6                               Checksum    : 0xb7c4
Version     : 1                               Pkt Type    : 20
Attributes: L1L2                             Max Area    : 3
SysID Len   : 6                               Used Len     : 311
Level       : L2                               Lifetime    : 1153
Pkt Ver     : 1                               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 69: Output fields: IS-IS database

Label	Description
LSP ID	LSP IDs are auto-assigned by the originating IS-IS node. The LSP ID consists of three sections. The first 6 bytes represent the system ID for that node, followed by a single byte value for the pseudonode generated by that router, and finally a fragment byte that starts at zero. 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 pseudonode LSP is created. Usually

Label	Description
	the internal circuit ID is used to determine the ID assigned to the pseudonode. For example, for circuit 4, a LSP pseudonode with ID 1800.0000.0029.04-00 is created. The router learns hostnames and uses the hostname in place of the system ID. Example of LDP IDs are the following: acc_ar1.00-00 acc_ar1.00-01 acc_ar1.04-00
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	OV — the overload bit is set L1 — specifies a level 1 IS type L2 — specifies a level 2 IS type 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.
LSP Count	Displays a 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 or protocol ID extension. This value is always set to 1.
Pkt Type	Displays the PDU type number
Pkt Ver	Displays the version or 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

Label	Description
Supp Protocols	Displays the data protocols that are supported
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 70: Output fields: 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 70: Output fields: IS-IS hostname

Label	Description
System-id	Displays the system identifier mapped to the hostname
Hostname	Displays the host name for the specific system ID
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 displays IS-IS interface information. When no *ip-address* or *ip-int-name* is specified, all interfaces are listed.

Parameters

ip-address

Displays only the interface information associated with the specified IP address.

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

ip-int-name

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

detail

Displays all output specified in the detailed format.

Output

The following output is an example of IS-IS interface information, and [Table 71: 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#

*A:7210-SAS>show>router>isis# interface detail

=====
Router Base ISIS Instance 1 Interfaces
=====
-----
Interface      : abcd                      Level Capability: L1L2
Oper State     : Down                      Admin State      : Up
Auth Type      : None
Circuit Id     : 7                        Retransmit Int.  : 5
Type           : Broadcast                 LSP Pacing Int. : 100
Mesh Group     : Inactive                  CSNP Int.        : 10
Bfd Enabled    : No
Te Metric      : 0                        Te State         : Down
Admin Groups   : None
Ldp Sync       : outOfService              Ldp Sync Wait    : Disabled
Ldp Timer State: Disabled Ldp Tm Left      : 0
Route Tag      : None LFA : Included

Level          : 1                        Adjacencies      : 0
Desg. IS       : 0000.0000.0000
Auth Type      : None                      Metric           : 10
Hello Timer    : 9                        Hello Mult.      : 3
Priority        : 64                       Passive          : No

Level          : 2                        Adjacencies      : 0
Desg. IS       : 0000.0000.0000
Auth Type      : None                      Metric           : 10
Hello Timer    : 9                        Hello Mult.      : 3
Priority        : 64                       Passive          : No
```

```

=====
*A:7210-SAS>show>router>isis#

*A:7210-SAS>show>router>isis# interface abcd detail

=====
Router Base ISIS Instance 1 Interfaces
=====
-----
Interface      : abcd                      Level Capability: L1L2
Oper State     : Down                      Admin State      : Up
Auth Type      : None
Circuit Id     : 7                        Retransmit Int.  : 5
Type           : Broadcast                 LSP Pacing Int.  : 100
Mesh Group     : Inactive                  CSNP Int.        : 10
Bfd Enabled    : No
Te Metric      : 0                        Te State         : Down
Admin Groups   : None
Ldp Sync       : outOfService              Ldp Sync Wait    : Disabled
Ldp Timer State: Disabled                  Ldp Tm Left      : 0
Route Tag      : None LFA : Included

Level          : 1                        Adjacencies      : 0
Desg. IS       : 0000.0000.0000           Metric           : 10
Auth Type      : None                     Hello Mult.      : 3
Hello Timer    : 9                        Passive          : No
Priority        : 64

Level          : 2                        Adjacencies      : 0
Desg. IS       : 0000.0000.0000           Metric           : 10
Auth Type      : None                     Hello Mult.      : 3
Hello Timer    : 9                        Passive          : No
Priority        : 64

=====
*A:7210-SAS>show>router>isis#

```

Table 71: 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	Specifies whether LFA is included or excluded

prefix-sids

Syntax

prefix-sids [**ipv4-unicast**] [*ip-prefix[/prefix-length]*] [**sid** *sid*] [**adv-router** *system-id | hostname*]

Context

show>router>isis

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-Sx/S 1/10GE

Description

This command displays IS-IS prefix SIDs.

Parameters

ipv4-unicast

Displays information for 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-length	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 hostname, up to 38 characters.

Output

The following output is an example of IS-IS prefix SID information, and [Table 72: 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  AdvRtr
                               MT      Flags
-----
10.0.0.1/32                         1        2/Int.   N     Dut-B
```


10.0.0.1/32	1	2/Int.	N	0	RnNP
					Dut-C
10.0.0.1/32	1	1/Int.	N	0	RnNP
					Dut-D
10.0.0.1/32	1	2/Int.	N	0	NnP
					Dut-D
10.0.0.1/32	1	2/Int.	N	0	NnP
					Dut-E
10.20.1.2/32	1002	1/Int.	N	0	RnNP
					Dut-B
10.20.1.2/32	1002	2/Int.	N	0	NnP
					Dut-B
10.20.1.2/32	1002	2/Int.	N	0	NnP
					Dut-C
10.20.1.2/32	1002	2/Int.	N	0	RnNP
					Dut-D
10.20.1.2/32	1002	2/Int.	N	0	RnNP
					Dut-E
10.20.1.3/32	1003	2/Int.	N	0	RnNP
					Dut-B
10.20.1.3/32	1003	1/Int.	N	0	RnNP
					Dut-C
10.20.1.3/32	1003	2/Int.	N	0	NnP
					Dut-C
10.20.1.3/32	1003	2/Int.	N	0	NnP
					Dut-D
10.20.1.3/32	1003	2/Int.	N	0	RnNP
					Dut-E
10.20.1.4/32	1004	2/Int.	N	0	RnNP
					Dut-B
10.20.1.4/32	1004	2/Int.	N	0	RnNP
					Dut-C
10.20.1.4/32	1004	1/Int.	N	0	RnNP
					Dut-D
10.20.1.4/32	1004	2/Int.	N	0	NnP
					Dut-D
10.20.1.4/32	1004	2/Int.	N	0	NnP
					Dut-E
10.20.1.5/32	1005	2/Int.	N	0	RnNP
					Dut-B
10.20.1.5/32	1005	2/Int.	N	0	RnNP
					Dut-C
10.20.1.5/32	1005	2/Int.	N	0	RnNP
					Dut-D
10.20.1.5/32	1005	1/Int.	N	0	RnNP
					Dut-E
10.20.1.5/32	1005	2/Int.	N	0	NnP
					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 72: 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	Indicates 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-id-number

Displays multi-topology parameters.

Values 0, 2**ip-prefix/prefix-length**

Specifies the IP prefix and mask length.

alternative

Displays the level of protection per prefix.

Output

The following output is an example of IS-IS routes information, and [Table 73: 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 [Flags] Metric      Lvl/Typ Ver.   SysID/Hostname
NextHop      MT
-----
10.10.1.0/24          10          1/Int.  5      Dut-A
0.0.0.0              0
10.10.3.0/24[L] 20      1/Int. 137    Dut-B
10.10.1.2              0
10.10.4.0/24          20          1/Int. 137    Dut-B
10.10.1.2              0
10.10.5.0/24          30          1/Int. 137    Dut-B
10.10.1.2              0
10.10.9.0/24          60          1/Int.  52    Dut-F
10.10.21.6            0
10.10.10.0/24         70          1/Int.  52    Dut-F
10.10.21.6            0
10.10.12.0/24         20          1/Int. 137    Dut-B
10.10.1.2              0
10.10.13.0/24         10          1/Int.  7      Dut-A
0.0.0.0              0
10.10.14.0/24 [L] 20      1/Int.  52    Dut-F
10.10.21.6            0
10.10.15.0/24         30          1/Int. 137    Dut-B
10.10.1.2              0
10.10.16.0/24         30          1/Int. 137    Dut-B
10.10.1.2              0
10.10.21.0/24         10          1/Int.  48    Dut-A
0.0.0.0              0
10.10.22.0/24         30          1/Int. 137    Dut-B
10.10.1.2              0
10.20.1.1/32          0           1/Int.  10    Dut-A
0.0.0.0              0
10.20.1.2/32          10          1/Int. 137    Dut-B
10.10.1.2              0
10.20.1.3/32          20          1/Int. 137    Dut-B
10.10.1.2              0
10.20.1.4/32          20          1/Int. 137    Dut-B
10.10.1.2              0
10.20.1.5/32          30          1/Int. 137    Dut-B
10.10.1.2              0
```

```

10.20.1.6/32          10          1/Int.  52      Dut-F
  10.10.21.6          0
3FFE::A0A:100/120    10          1/Int.   5      Dut-A
  ::                  0
10.10.1.0/24         10          1/Int.  65      Dut-A
  0.0.0.0             2
10.10.13.0/24        10          1/Int.  65      Dut-A
  0.0.0.0             2
10.10.21.0/24        10          1/Int.  65      Dut-A
  0.0.0.0             2
10.20.1.1/32         0          1/Int.  65      Dut-A
  0.0.0.0             2
3FFE::A0A:100/120    10          1/Int.  65      Dut-A
  ::                  2
3FFE::A0A:300/120    20          1/Int. 116      Dut-B
  FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A0A:400/120    20          1/Int. 116      Dut-B
  FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A0A:500/120    30          1/Int. 130      Dut-B
  FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A0A:900/120    60          1/Int.  71      Dut-F
  FE80::2285:FFFF:FE00:0-"ies-1-3FFE::A0A:1501" 2
3FFE::A0A:A00/120    70          1/Int.  71      Dut-F
  FE80::2285:FFFF:FE00:0-"ies-1-3FFE::A0A:1501" 2
3FFE::A0A:C00/120    20          1/Int. 116      Dut-B
  FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A0A:D00/120    10          1/Int.  65      Dut-A
  ::                  2
3FFE::A0A:E00/120    20          1/Int.  71      Dut-F
  FE80::2285:FFFF:FE00:0-"ies-1-3FFE::A0A:1501" 2
3FFE::A0A:F00/120    30          1/Int. 130      Dut-B
  FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A0A:1000/120   30          1/Int. 130      Dut-B
  FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A0A:1500/120   10          1/Int.  65      Dut-A
  ::                  2
3FFE::A0A:1600/120   30          1/Int. 127      Dut-B
  FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A14:101/128    0          1/Int.  65      Dut-A
  ::                  2
3FFE::A14:102/128    10          1/Int. 116      Dut-B
  FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A14:103/128    20          1/Int. 130      Dut-B
  FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A14:104/128    20          1/Int. 127      Dut-B
  FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A14:105/128    30          1/Int. 130      Dut-B
  FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A14:106/128    10          1/Int.  71      Dut-F
  FE80::2285:FFFF:FE00:0-"ies-1-3FFE::A0A:1501" 2
-----
No. of Routes: 20
Flags: L = LFA nexthop available
=====
*A:Dut-A#

*A:Dut-B# show router isis routes alternative

=====
Route Table
=====
Prefix [Flags]          Metric   Lvl/Typ   Ver.   SysID/Hostname
NextHop                MT        AdminTag
Alt-Nexthop            Alt-Metric

```

```

-----
10.20.1.2/32          0          1/Int.    3      Dut-B
 0.0.0.0              0
10.20.1.3/32         10         2/Int.    2      Dut-C
 10.20.3.3            0
 10.20.3.3 (lfa)      15
10.20.1.4/32         10         2/Int.    3      Dut-D
 10.20.4.4            0
10.20.1.5/32         20         2/Int.    3      Dut-C
 10.20.3.3            0
10.20.1.6/32         20         2/Int.    3      Dut-D
 10.20.4.4            0
10.20.3.0/24         10         1/Int.    3      Dut-B
 0.0.0.0              0
10.20.4.0/24         10         1/Int.    3      Dut-B
 0.0.0.0              0
10.20.5.0/24         20         2/Int.    2      Dut-C
 10.20.3.3            0
10.20.6.0/24         20         2/Int.    4      Dut-D
 10.20.4.4            0
10.20.9.0/24         20         2/Int.    3      Dut-D
 10.20.4.4            0
10.20.10.0/24        30         2/Int.    3      Dut-C
 10.20.3.3            0
-----
Routes : 11
Flags: LFA = Loop-Free Alternate nexthop
=====
*A:Dut-B#

```

Table 73: 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 the next hop (or the hostname, if possible)
Hostname	Displays the hostname for the specific system ID

spf

Syntax

spf [detail]

Context

show>router>isis

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information about SPF calculation.

Parameters

detail
Displays detailed information about SPF.

Output

The following output is an example of IS-IS SPF information, and [Table 74: Output fields: IS-IS SPF](#) describes the output fields.

Sample output

```
A:ALA-A# show router isis spf
=====
Router Base ISIS Instance 1 Path Table
=====
Node                               Interface    Nexthop
-----
abr_sjc.00                         if2/2       dist_oak
abr_sjc.00                         if2/3       dist_nj
dist_oak.00                        if2/2       dist_oak
dist_nj.00                         if2/3       dist_nj
acc_nj.00                          if2/3       dist_nj
acc_ri.00                          if2/3       dist_nj
core_west.00                       if2/8       core_west
core_east.00                       lag-1       core_east
asbr_west.00                       if2/8       core_west
asbr_east.00                       if2/5       asbr_east
abr_sjc.00                         lag-1       core_east
abr_sjc.00                         if2/8       core_west
abr_lax.00                         lag-1       core_east
abr_lax.00                         if2/8       core_west
abr_dfw.00                         if2/5       asbr_east
abr_dfw.00                         lag-1       core_east
abr_dfw.00                         if2/8       core_west
dist_arl.00                        if2/5       asbr_east
dist_arl.00                        lag-1       core_east
dist_arl.00                        if2/8       core_west
dist_msq.00                        if2/5       asbr_east
dist_msq.00                        lag-1       core_east
dist_msq.00                        if2/8       core_west
acc_arl.00                         if2/5       asbr_east
acc_arl.00                         lag-1       core_east
acc_arl.00                         if2/8       core_west
acc_msq.00                         if2/5       asbr_east
acc_msq.00                         lag-1       core_east
acc_msq.00                         if2/8       core_west
acc_msq.03                         if2/5       asbr_east
acc_msq.03                         lag-1       core_east
acc_msq.03                         if2/8       core_west
acc_msq.04                         if2/5       asbr_east
acc_msq.04                         lag-1       core_east
acc_msq.04                         if2/8       core_west
=====
```

```
A:ALA-A#

A:ALA-A# show router isis spf detail
=====
Router Base ISIS Instance 1 Path Table
=====
Node       : abr_sjc.00           Metric : 20
Interface  : if2/2               SNPA    : 00:00:00:00:00:00
Nexthop    : dist_oak

Node       : abr_sjc.00           Metric : 20
Interface  : if2/3               SNPA    : 00:00:00:00:00:00
Nexthop    : dist_nj

Node       : dist_oak.00          Metric : 10
Interface  : if2/2               SNPA    : 00:00:00:00:00:00
Nexthop    : dist_oak

Node       : dist_nj.00           Metric : 10
Interface  : if2/3               SNPA    : 00:00:00:00:00:00
Nexthop    : dist_nj

Node       : acc_nj.00            Metric : 20
Interface  : if2/3               SNPA    : 00:00:00:00:00:00
Nexthop    : dist_nj

Node       : acc_ri.00            Metric : 20
Interface  : if2/3               SNPA    : 00:00:00:00:00:00
Nexthop    : dist_nj

Node       : core_west.00         Metric : 10
Interface  : if2/8               SNPA    : 00:00:00:00:00:00
Nexthop    : core_west

...
=====
A:ALA-A#
```

Table 74: Output fields: IS-IS SPF

Label	Description
Node	Displays the route node and mask
Interface	Displays the outgoing interface name for the route
Metric	Displays the route metric
Nexthop	Displays the system ID of the next hop or hostname
SNPA	Displays the Subnetwork Points of Attachment (SNPA) where a router is physically attached to a subnetwork

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.

Parameters

detail
Displays detailed information about SPF events.

Output

The following output is an example of IS-IS SPF event information, and [Table 75: 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 75: 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

Label	Description
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 76: 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 76: 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 SPF 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	The count of link state PDUs that had to be retransmitted by this instance of the protocol
LFA Runs	Displays the number of times 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 77: Output fields: IS-IS status](#) describes the output fields.

Sample output

```
*A:Dut-A>config>router>isis# show router isis status
=====
Router Base ISIS Instance 1 Status
=====
System Id           : 0100.2000.1001
Admin State         : Up
Ipv4 Routing        : Enabled
Last Enabled        : 02/13/2008 02:22:38
Level Capability     : L1L2
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     : Enabled
L2 Wide Metrics     : Enabled
L1 LSDB Overload    : Disabled
L2 LSDB Overload    : Disabled
L1 LSPs             : 6
L2 LSPs             : 6
Last SPF            : 02/13/2008 19:32:16
SPF Wait            : 10 sec (Max)   1000 ms (Initial)   1000 ms (Second)
Export Policies     : None
Multicast Import    : None
Multi-topology      : Disabled
Area Addresses      : 01
Ldp Sync Admin State : Up
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-A>config>router>isis#

*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 77: 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

Label	Description
	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 and 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
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	Indicates whether the IGP-LDP synchronization feature is enabled or disabled on all interfaces participating in the OSPF 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 addresses.

Parameters

ip-prefix/prefix-length

Specifies IP prefix and mask length.

Output

The following output is an example of IS-IS summary address information, and [Table 78: 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 78: Output fields: IS-IS summary address

Label	Description
Address	Displays the IP address
Level	Specifies 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 topology information.

Parameters

ipv4-unicast
Displays IPv4 unicast parameters.

ipv6-unicast
Displays IPv6 unicast parameters.

mt-id-number
Displays multi-topology parameters.

Values 0, 2

detail
Displays detailed topology information.

Output

The following output is an example of IS-IS topology information, and [Table 79: Output-fields: IS-IS topology](#) describes the output fields.

Sample output

```
*A:7210>show>router>isis# topology

=====
Router Base ISIS Instance 1 Topology Table
=====
Node                               Interface                     Nexthop
-----
IS-IS IPv6 paths (MT-ID 2),  Level 1
-----
P1_Router1.00                      B_to_ixia                     P1_Router1
-----
IS-IS IP paths (MT-ID 0),  Level 2
-----
Dut-A.00                           B_to_A                         Dut-A
Dut-A.30                           B_to_A                         Dut-A
Dut-A.31                           B_to_A                         Dut-A
Dut-A.32                           B_to_A                         Dut-A
Dut-A.33                           B_to_A                         Dut-A
Dut-A.34                           B_to_A                         Dut-A
Dut-A.35                           B_to_A                         Dut-A
Dut-A.36                           B_to_A                         Dut-A
Dut-A.37                           B_to_A                         Dut-A
Dut-A.38                           B_to_A                         Dut-A
Dut-A.39                           B_to_A                         Dut-A
Dut-A.3A                           B_to_A                         Dut-A
Dut-C.00                           B_to_C                         Dut-C
-----
IS-IS IPv6 paths (MT-ID 2),  Level 2
-----
```

```
Dut-A.00      B_to_A      Dut-A
Dut-A.30      B_to_A      Dut-A
Dut-A.31      B_to_A      Dut-A
Dut-A.32      B_to_A      Dut-A
Dut-A.33      B_to_A      Dut-A
Dut-A.34      B_to_A      Dut-A
Dut-A.35      B_to_A      Dut-A
Dut-A.36      B_to_A      Dut-A
Dut-A.37      B_to_A      Dut-A
Dut-A.38      B_to_A      Dut-A
Dut-A.39      B_to_A      Dut-A
Dut-A.3A      B_to_A      Dut-A
Dut-C.00      B_to_C      Dut-C
=====
*A:7210>show>router>isis#
```

Table 79: Output-fields: IS-IS topology

Label	Description
Node	Displays the IP address
Interface	Displays the interface name
Nexthop	Displays the next-hop IP address

5.10.2.3 Clear commands

isis

Syntax

isis [isis-instance]

Context

clear>router>isis

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 that only the specified entries are removed from the IS-IS adjacency database.

Values 6 octets 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 that contains information about PDUs.

Parameters

system-id

Specifies that only the specified entries are removed from the IS-IS link-state database.

Values 6 octets 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 re-evaluates 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.

5.10.2.4 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 the 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

Debugs adjacencies with the specified interface.

ip-address

Debugs adjacencies with the specified IP address.

Values ipv4-address : a.b.c.d

nbr-system-id

Debugs adjacencies with the specified system ID.

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 constraint-based shortest path first (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 IS-IS interface.

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

Parameters

ip-int-name

Debugs the interface information associated with the specified interface.

ip-address

Debugs the interface information associated with the specified IP address.

Values ipv4-address: a.b.c.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

Debugs IS-IS leak information associated with the specified IP address.

Values ipv4-address: a.b.c.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 the Link State DataBase (LSDB).

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

Parameters

level-number

Debugs for level 1 or level 2.

system-id

Debugs the LSDB associated with the specified system ID.

lsp-id

Debugs the LSDB associated with the specified LSP ID.

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 miscellaneous IS-IS events.
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

- packet-type***

Specifies the packet type to debug.

Values	ptop-hello, l1-hello, l2-hello, l1-psnp, l2-psnp, l1-csnp, l2-csnp, l1-lsp, l2-lsp
- ip-int-name***

Specifies the interface name to debug.
- ip-address***

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

Values	ipv4-address: a.b.c.d
- detail***

Specifies to provide detailed debugging information.

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 IS-IS route table manager (RTM).
The **no** form of this command disables debugging.

Parameters

ip-address

Displays the interface information associated with the specified IP address.

Values

ipv4-address:	a.b.c.d
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

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.

Parameters

level-number

Specifies level 1 or level 2.

system-id

Specifies the system ID.

6 BGP

This chapter provides information about configuring BGP.



Note:

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

6.1 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 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 VPN-IPv4, VPN-IPv6, and L2-VPN (BGP-AD) BGP address families are supported on the 7210 SAS-Mxp (standalone), 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), 7210 SAS-Sx 10/100GE (standalone and standalone-VC), and 7210 SAS-T (network).
- The MVPN-IPv4 BGP address family is supported on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), and 7210 SAS-T (network).
- The EVPN BGP address family is supported only on the 7210 SAS-Mxp (standalone), 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-Sx/S 1/10GE (standalone).

6.2 BGP communication

There are two types of BGP peers: internal BGP (iBGP) and external BGP (eBGP) ([Figure 24: 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.

6.2.1 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**

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 24: BGP configuration



6.3 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 for neighbors to exchange routing information and updates. Neighbors exchange BGP open messages that includes information such as ASNs, 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 waits 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 is not established and updates are not 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.

6.4 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.

6.5 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:

7210 SAS-R6 and 7210 SAS-R12 devices can be configured as route reflector clients or servers. These devices support route reflector server functionality for VPN-IPv4, VPN-IPv6, and BGP LU routes. All other 7210 SAS devices can be configured only as route reflector clients.


A route reflector (RR) provides route reflection services to iBGP peers called clients. Other iBGP peers of the RR are called non-clients. An RR and its client peers form a cluster. A large AS can be subdivided into multiple clusters, each identified by a unique 32-bit cluster ID. 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 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, SR-E, SR-F). To prevent loops, iBGP learned routes are not re-advertised to other iBGP peers.

Figure 25: 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 26: 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 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 26: 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 (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 the operator needs 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.

A router becomes a route reflector whenever it has one or more client iBGP sessions. A client iBGP session is created with the **cluster** command, which also indicates the cluster ID of the client. The router ID is typically used as the cluster ID, but this is not necessary.

Basic route reflection operation (without **add-path** configured) can be summarized as follows:

- If the best and valid path for an NLRI is learned from a client and **disable-client-reflect** is not configured, then that route is advertised to all clients, non-clients, and eBGP peers (as allowed by policy). If the client that advertised the best and valid path is a neighbor to which the **split-horizon** command (at the **bgp**, **group**, or **neighbor** level) applies, then the route is not advertised back to the sending client. In the route that is reflected to clients and non-clients:
 - The route reflector adds an ORIGINATOR_ID attribute if it does not already exist; the ORIGINATOR_ID indicates the BGP identifier (router ID) of the client that originated the route.
 - The route reflector prepends the cluster ID of the client that advertised the route and then the cluster ID of the client receiving the route (if applicable) to the CLUSTER_LIST attribute, creating the attribute if it does not already exist.

- If the best and valid path for an NLRI is learned from a client and **disable-client-reflect** is configured, then that route is advertised to all clients in other clusters, non-clients, and eBGP peers (as allowed by policy). In the route that is reflected to clients in other clusters and non-clients:
 - The route reflector adds an **ORIGINATOR_ID** attribute if it does not already exist; the **ORIGINATOR_ID** indicates the BGP identifier (router ID) of the client that originated the route.
 - The route reflector prepends the cluster ID of the client that advertised the route and then the cluster ID of the client receiving the route (if applicable) to the **CLUSTER_LIST** attribute, creating the attribute if it does not already exist.
- If the best and valid path for an NLRI is learned from a non-client, then that route is advertised to all clients and eBGP peers (as allowed by policy). In the route that is reflected to clients:
 - The route reflector adds an **ORIGINATOR_ID** attribute if it does not already exist; the **ORIGINATOR_ID** indicates the BGP identifier (router ID) of the non-client that originated the route.
 - The route reflector prepends the cluster ID of the client receiving the route to the **CLUSTER_LIST** attribute, creating the attribute if it already exists.
- If the best and valid path for an NLRI is learned from an eBGP peer, then that route is advertised to all clients, non-clients, and other eBGP peers (as allowed by policy). The **ORIGINATOR_ID** and **CLUSTER_LIST** attributes are not added to the route.
- If the best and valid path for an NLRI is locally originated by the RR — that is, it was learned through means other than BGP — then that route is advertised to all clients, non-clients, and eBGP peers (as allowed by policy). The **ORIGINATOR_ID** and **CLUSTER_LIST** attributes are not added to the route.

The **ORIGINATOR_ID** and **CLUSTER_LIST** attributes allow BGP to detect the looping of a route within the AS. If any router receives a BGP route with an **ORIGINATOR_ID** attribute containing its own BGP identifier, the route is considered invalid. In addition, if a route reflector receives a BGP route with a **CLUSTER_LIST** attribute containing a locally configured cluster ID, the route is considered invalid. Invalid routes are not installed in the route table and are not advertised to other BGP peers.

If VPN-IPv4 and VPN-IPv6 routes are advertised by a 7210 SAS router, the BGP next-hop address is set as follows for a reflected route:

- When a route is reflected from one iBGP peer to another iBGP peer, the RR server does not modify the next hop by default; however, if the **next-hop-self** command is applied to the RR and **enable-rr-vpn-forwarding** is configured, then this combination of commands changes the next hop to the local address of the RR used with the peer.
- If BGP 3107 LU service optimization has been configured using the **config>router>bgp>group>neighbor>advertise-label ipv4 use-svc-routes** command along with the **enable-rr-vpn-forwarding** command, only those BGP 3107 LU routes for which the node installs a swap entry for a VPN label are installed in the FIB.



Note:

On the 7210 SAS, the **next-hop-self** command can be used only in conjunction with the **enable-rr-vpn-forwarding** command or for BGP 3107 LU routes.

6.6 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.

6.7 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.

6.8 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.

6.9 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
```

**Note:**

The 7210 SAS-Mxp does not support the use of RLFA with BGP 3107 labeled routes.

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. When using RSVP LSP, only FRR one-to-one is supported for services using BGP 3107 labelled routes.

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.

**Note:**

RLFA used with SR is not supported for BGP 3107 tunnels when services are resolved to use BGP 3107 tunnels.

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.

6.9.1 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.

6.10 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 criteria are used to determine the best path:

1. Routes are not considered if they are unreachable.
2. An RTM 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 instead of 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:

7210 SAS devices do not support BGP ECMP (multipath). An ECMP value is set to 1 and the following are assumed:

- 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 is 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.
- 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.

6.11 Enabling best external

Enabling the best-external feature is supported only at the **config>router>bgp** level. This feature can be enabled/disabled on a per address family basis, with IPv4 and IPv6 as the only options supported initially. Enabling best-external for IPv4 causes the new advertisement rules to apply to both regular IPv4 unicast routes as well as labeled-IPv4 (SAFI4) routes. Similarly, enabling best-external for IPv6 causes the new advertisement rules to apply to both regular IPv6 unicast routes as well as labeled-IPv6 (SAFI4) routes.

The **advertise-external** command cannot be applied to a route reflector unless client-to-client reflection is disabled (disable-client-reflect in the CLI).

6.11.1 BGP decision process with best external

When best-external is enabled for an address family, all routes belonging to that address family must be classified internally as either "internal" or "external". A route is "internal" if:

- It was received from an iBGP peer in the same AS.
- It was originated by a router in the same or a different RR cluster of the same AS.

A route is "external" if it was received from an eBGP peer in a different AS.

The tie-breaking steps of the decision process are run as usual on all of the routes (both "internal" and "external") for a particular destination until only one path, the best path, is left. If this is an external route then the decision process must be rerun on only the "internal" routes to find the single best path in that subset. This "best internal" route is advertised to confed-eBGP peers, as described in [Advertisement rules with best external](#).

If the overall best path found by the first run of the decision process is an internal route with NEXT_HOP *n* the decision process must be rerun on only the "external" routes with NEXT_HOP not equal to *n* to find the single best path in that subset. This "best external" route is advertised to iBGP peers, as described in [Advertisement rules with best external](#).

6.11.2 Advertisement rules with best external

The advertisement rules when advertise-external is enabled can be summarized as follows:

- If a router has advertise-external enabled and its best overall route is an internal route then this best route should be advertised to:

- all iBGP RR clients (if the route came from a non-client peer) or all iBGP non-clients (if the route came from a client peer).
- all eBGP peers
- If there is a best external route, it should be sent to iBGP client and non-client peers instead of the best overall route.
- If a router has advertise-external enabled and its best overall route is an external route then this best route should be advertised to all iBGP peers and all eBGP peers.

6.11.3 Displaying best-external routes

BGP **show** commands display the following information for this feature:

- For each RIB-IN entry in the output of the **show router bgp routes prefix hunt** command there is a Flags field that indicates the origin of the route and whether it is valid, best, used, and so on. This feature reflects an "Advertised" value in the Flags field. This indicates that the route was advertised to one or more peers. If the "Advertised" flag is present but the "Best" flag is not the operator can determine that the route was probably a best-external.
- The **show router bgp neighbor advertised-routes** command displays all advertised routes to that peer, including routes that were overall best, best-external and best-internal.
- The advertise-external configuration (specifically the address families for which it is enabled) is displayed as part of the **show router bgp** output.

Note that the overall best, best-external, and best-internal routes for a prefix can be determined from the output of the **show router bgp routes prefix** command. The first external route to be displayed in the output is always be the best-external route and the first internal route to be displayed in the output is always be the best-internal route. Only one of these routes will have the "Best" flag set, and this will be the overall best route.

6.12 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 with 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 26.x.Rx* for more information about NG-MVPN with BGP signaling)

6.12.1 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.

6.12.1.1 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.

The following families support next-hop indirection on the 7210 SAS:

- label-IPv4
- MVPN-IPv4 (not supported on the 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE)
- VPN-IPv4
- label-IPv6
- VPN-IPv6
- L2-VPN
- PW route

6.13 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.

When a VPN route is rejected by an import policy or not imported by any services, it is deleted from the RIB-IN. For VPN-IPv4 and VPN-IPv6 routes, this behavior can be changed by configuring the **mp-bgp-keep** command. Configuring this option maintains rejected VPN-IP routes in the RIB-IN and issuing a Route Refresh message is not required for an import policy change.

6.13.1 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)

The BGP fast reroute feature is described in [BGP fast reroute](#) and [BGP fast reroute in a VPRN](#).

6.13.2 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, VPN-IPv4, and VPN-IPv6 routes. For information about BGP fast reroute in a VPRN, see [BGP fast reroute in a VPRN](#).

The following table describes the scenarios supported in the base router BGP context.

Table 80: 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

6.13.2.1 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.

A prefix supports ECMP paths or a backup path, but not both. 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.

6.13.2.2 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 because of an unresolved next hop

6.13.3 BGP fast reroute in a VPRN

In a VPRN context, BGP fast reroute is supported for VPN-IPv4 and VPN-IPv6 routes. On the 7210 SAS, BGP fast reroute is not supported for unlabeled IPv4 and unlabeled IPv6 routes.

The following table describes the supported VPRN scenarios.

Table 81: BGP fast reroute scenarios (VPRN context)

Ingress packet	Primary route	Backup route	PIC
IPv4 (ingress PE)	IPv4 route with next-hop A resolved by an IPv4 route	IPv4 route with next-hop B resolved by an IPv4 route	No
IPv4 (ingress PE)	VPN-IPv4 route with next-hop A resolved by an LDP, RSVP, or BGP tunnel	VPN-IPv4 route with next-hop A resolved by an LDP, RSVP, or BGP tunnel	Yes
MPLS (egress PE)	IPv4 route with next-hop A resolved by an IPv4 route	IPv4 route with next-hop B resolved by an IPv4 route	No
MPLS (egress PE)	IPv4 route with next-hop A resolved by an IPv4 route	VPN-IPv4 route with next-hop B resolved by an LDP, RSVP, or BGP tunnel	No

6.13.3.1 BGP fast reroute in a VPRN configuration

In a VPRN context, BGP fast reroute is optional and must be enabled. Fast reroute can be applied to all IPv4 prefixes, all IPv6 prefixes, all IPv4 and IPv6 prefixes, or to a specific set of IPv4 and IPv6 prefixes.

If all IP prefixes require backup path protection, use a combination of the BGP instance-level **backup-path** and VPRN-level **enable-bgp-vpn-backup** commands. The VPRN BGP **backup-path** command enables BGP fast reroute for all IPv4 prefixes and/or all IPv6 prefixes that have a best path through a VPRN BGP peer. The VPRN-level **enable-bgp-vpn-backup** command enables BGP fast reroute for all IPv4 prefixes and/or all IPv6 prefixes that have a best path through a remote PE peer.

6.13.4 RIB-OUT features

This section describes features related to RIB-OUT processing.

6.13.4.1 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 specified BGP routes toward 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.

6.13.4.2 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.

6.13.4.3 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
- MVPN-IPv4
- L2-VPN (BGP-AD)
- EVPN



Note:

BGP address family support varies per 7210 SAS platform. RTC is supported only for the BGP families that the specific 7210 SAS platform supports. See [BGP overview](#) for more information.

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.

6.13.4.4 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
- MVPN-IPv4

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.

6.13.4.5 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-paths](#).

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.

6.13.4.6 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.

6.14 Add-paths

6.14.1 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 able 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 a VPN-IPv4 NLRI with a path identifier.

Figure 27: BGP Update message with path identifier for VPN-IPv4 NLRI



The following figure shows an Update message carrying an IPv4 NLRI with a path identifier.

Figure 28: BGP Update message with path identifier for IPv4 NLRI



Currently, add-path is only supported by 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
- VPN-IPv4
- VPN-IPv6

6.14.2 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-paths 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-paths 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: a route learned from an iBGP neighbor cannot be readvertised to another iBGP neighbor unless the router is configured as a route reflector.

6.14.3 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.

6.14.4 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 able 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. Route type support is as follows:

- label-IPv4
- label-IPv6
- VPN-IPv4
- VPN-IPv6

6.14.5 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 , ..., 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)

6.15 AIGP metric



Note:

Accumulated Interior Gateway Protocol (AIGP) is only supported on 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

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 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.

6.16 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.

6.16.1 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.

6.16.2 BGP advertisement

BGP advertisement allows a BGP router to indicate to a peer, using the Optional Parameter, the features that it supports so that the router and peer can coordinate and use only the features that both support. Each capability in the Optional Parameter is TLV-encoded with a unique type code. The 7210 SAS supports the following capability codes:

- Multiprotocol BGP (code 1)
- Route refresh (code 2)
- Outbound route filtering (code 3)
- Graceful restart (code 64)
- 4-octet ASN (code 65)
- Add-path (code 69)

6.16.3 Changing the local ASN

Changing the local AS 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 reestablish the peer relationships with all peers in the group with the new local ASN
- at the neighbor level, causes BGP to reestablish the peer relationship with the new local ASN

6.16.4 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.

6.16.5 Hold time and keep alive timer dependencies

The BGP hold time specifies the maximum time BGP waits 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
This applies to all peers.
- group level
This applies to all peers in group.
- neighbor level
This applies only to a 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 is 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.

6.16.6 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.

6.16.7 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.

6.16.8 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 because 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.

6.17 Configuration guidelines for BGP


The following are the configuration guidelines for BGP:

- The 7210 SAS platforms can act only as a route reflector clients, except for the 7210 SAS-R6 and 7210 SAS-R12, which can act as route reflector servers.
- The 7210 SAS platforms support IPv4 family for PE-CE eBGP instance and for RFC 3107 labeled IPv4 routes. It does not support the IPv4 family in the base routing instance to exchange IPv4 routes.
- The 7210 SAS platforms support IPv6 family for PE-CE eBGP instance and for RFC 3107 labeled IPv6 routes. It does not support the IPv6 family in the base routing instance to exchange IPv6 routes.

6.18 BGP configuration process overview

The following figure shows the process to provision basic BGP parameters.

Figure 29: BGP configuration and implementation flow



6.19 Configuration notes

This section describes BGP configuration restrictions.

6.19.1 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.

6.19.1.1 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, then all BGP routes are accepted.
- If no export route policy statements specified, then all best and used BGP routes are advertised and non-BGP routes are not advertised.

6.19.1.2 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 82: 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 ¹⁶

If SNMP is used to set a value of X to the MIB variable in the following table, there are three possible results:

Table 83: 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 to 255 and X = 65535, MIB variable is 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:

¹⁶ A value of 0 is supported when the rapid-update command is applied to an address family that supports it.

Example: Sample 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: Sample 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: Sample 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: Sample 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
```

6.20 Configuring BGP with CLI

This section provides information to configure BGP using the command line interface.

6.20.1 BGP configuration overview

6.20.1.1 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. 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 origin. The router ID must be a valid host address.

6.20.1.2 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.

6.20.1.3 Internal and external BGP configurations

A BGP system consists of ASs that 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. Because 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.

6.21 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, 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 other parameters are optional.

```
info
```

```
#-----
echo "IP Configuration"
#-----
...
    autonomous-system 200
    router-id 10.10.10.103
#-----
...
#-----
echo "BGP Configuration"
#-----
    bgp
        exit
        cluster 0.0.0.100
        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
                no bfd-enable
                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
    exit
#-----
....
A:ALA-48>config>router#
```

6.22 Common configuration tasks

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 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.

6.22.1 Configuring a basic autonomous system

About this task

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:

1. Prepare a plan detailing the autonomous systems, the 7210 SAS node belonging to each group, group names, and peering connections.
2. Associate each 7210 SAS node with an autonomous system number.
3. Configure each 7210 SAS node with a router ID.
4. Associate each 7210 SAS node with a peer group name.
5. Specify the local IP address that will be used by the group or neighbor when communicating with BGP peers.
6. Specify neighbors.
7. Specify the autonomous system number associated with each neighbor.

6.22.2 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 ASNs by default. This means autonomous-system can have any value from 1 to 4294967295.

Example: Autonomous system configuration command usage

```
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#
```

6.22.3 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 a 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#
```

6.22.4 BGP components

The following section describes the syntax used to configure the BGP attributes.

6.22.5 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

Example

```
config>router# bgp          (global level)
                  group      (group level)
                  neighbor    (neighbor level)
```

**Note:**

Careful planning is essential when implementing 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#
```

6.22.6 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

The following is a sample 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#
```

6.22.7 Configuring neighbor attributes

After you create a group name and assign options, add neighbors within the same autonomous system to create IBGP connections and 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 10.5.0.2
        hold-time 90
        keepalive 30
        min-as-origination 15
        local-preference 170
        peer-as 10701
      exit
      neighbor 10.5.1.2
        hold-time 90
        keepalive 30
        min-as-origination 15
        local-preference 100
        min-route-advertisement 30
        preference 170
        peer-as 10702
      exit
    exit
...
-----
ALA-B>config>router>bgp#

```

6.22.8 Configuring AIGP



Note:

AIGP is only supported on 7210 SAS-Mxp.

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

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

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
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
  -----
```

6.22.9 BGP configuration management tasks

This section describes the BGP configuration management tasks.

6.22.9.1 Modifying an ASN

You can modify an ASN on a 7210 SAS but 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.

Because 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 systems, 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 reestablishes 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

```
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
```

6.22.9.2 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

```
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

This following is a sample BGP configuration 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#
```

6.22.9.3 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

```
config>router# router-id 10.10.10.104
config>router# no shutdown
config>router>bgp# shutdown
config>router>bgp# no shutdown
```

Example

The following is a sample 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#
```

6.22.9.4 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

```
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

The following is a sample "headquarters1" configuration output with the neighbor 10.0.0.103 removed.

```
ALA-B>config>router>bgp# info  
-----  
    group "headquarters1"  
        description "HQ execs"  
        local-address 10.0.0.104  
        neighbor 10.0.0.5  
            passive  
            peer-as 300  
        exit  
    exit  
-----  
ALA-B>config>router>bgp#
```

6.22.9.5 Deleting groups

To delete a group, the neighbor configurations must first be shut down. After each neighbor shuts 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

```
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
```

Trying to delete the group without shutting it down results in the following message:

```
ALA-B>config>router>bgp# no group headquarters1
MINOR: CLI BGP Peer Group should be shutdown before deleted. BGP Peer Group not
deleted.
```

6.22.9.6 Editing BGP parameters

Use the following syntax to change existing BGP parameters. 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.

6.23 BGP command reference

- [Command hierarchies](#)
- [Command descriptions](#)

6.23.1 Command hierarchies

- [Configuration commands](#)
- [Global BGP commands](#)
- [Group BGP commands](#)
- [Neighbor BGP commands](#)
- [Other BGP-related commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

6.23.1.1 Configuration commands

6.23.1.2 Global BGP commands

```
config
- router [router-name]
- [no] bgp
- [no] add-paths
- ipv4 send send-limit
- ipv4 send send-limit receive [none]
- no ipv4
- ipv6 send send-limit
- ipv6 send send-limit receive [none]
- no ipv6
- vpn-ipv4 send send-limit
- vpn-ipv4 send send-limit receive [none]
- no vpn-ipv4
- vpn-ipv6 send send-limit
- vpn-ipv6 send send-limit receive [none]
- no vpn-ipv6
- [no] advertise-inactive
- [no] aggregator-id-zero
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- auth-keychain name
- [no] backup-path [ipv4]
- 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
- cluster cluster-id
- no cluster
- connect-retry seconds
- no connect-retry
- [no] damping
```

```

- description description-string
- no description
- [no] disable-4byte-asn
- [no] disable-client-reflect
- disable-communities [standard] [extended]
- no disable-communities
- [no] disable-fast-external-failover
- [no] enable-peer-tracking
- [no] enable-rr-vpn-forwarding
- export policy-name [policy-name...(up to 15 max)]
- no export
- family [ipv4] [vpn-ipv4] [ipv6] [vpn-ipv6] [l2-vpn] [ms-pw] [mvpn-ipv4] [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
- [no] mp-bgp-keep
- multihop tll-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] [mvpn-ipv4] [evpn]
- no rapid-update
- [no] rapid-withdrawal
- [no] remove-private [limited]
- router-id ip-address
- no router-id
- [no] shutdown
- [no] split-horizon
- [no] vpn-apply-export
- [no] vpn-apply-import

```

6.23.1.3 Group BGP commands

```

config
- router [router-name]
- [no] bgp
- [no] group name
- [no] add-paths
- ipv4 send send-limit
- ipv4 send send-limit receive [none]
- no ipv4
- ipv6 send send-limit
- ipv6 send send-limit receive [none]
- no ipv6
- vpn-ipv4 send send-limit
- vpn-ipv4 send send-limit receive [none]
- no vpn-ipv4
- vpn-ipv6 send send-limit
- vpn-ipv6 send send-limit receive [none]
- no vpn-ipv6
- [no] advertise-inactive
- [no] aggregator-id-zero
- [no] aigp
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- auth-keychain name
- cluster cluster-id
- no cluster
- 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
- [no] disable-client-reflect
- 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] [ms-pw] [mvpn-ipv4]
[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 next-hop-self
- [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
```

6.23.1.4 Neighbor BGP commands

```
config
- router [router-name]
  - [no] bgp
    - [no] group name
      - [no] neighbor ip-address
        - [no] add-paths
          - ipv4 send send-limit
          - ipv4 send send-limit receive [none]
          - no ipv4
          - ipv6 send send-limit
          - ipv6 send send-limit receive [none]
          - no ipv6
          - no vpn-ipv4
          - vpn-ipv4 send send-limit
          - vpn-ipv4 send send-limit receive [none]
          - no vpn-ipv6
          - vpn-ipv6 send send-limit
          - vpn-ipv6 send send-limit receive [none]
        - [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
        - cluster cluster-id
        - no cluster
        - 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
- [no] disable-client-reflect
- 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] [ms-pw] [mvpn-ipv4]

[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 next-hop-self
- [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

```

6.23.1.5 Other BGP-related commands

```

config
- router [router-name]
  - autonomous-system as-number
  - no autonomous-system

```

6.23.1.6 Show commands

```
show
- router [router-instance]
  - bgp
    - auth-keychain keychain-name
    - damping [ip-prefix[/prefix-length] [damp type | detail] [ipv4]
    - damping [ip-prefix | prefix-length] [detail] vpn-ipv4
    - group [name] [detail]
    - neighbor [ip-address [detail]]
    - neighbor [as-number [detail]]
    - neighbor ip-address [family [type mvpn-type]] filter1 [brief]
    - neighbor ip-address [family] filter2
    - neighbor as-number [family] filter2
    - neighbor ip-address orf [filter3]
    - neighbor ip-address graceful-restart
    - paths
    - routes [family] [brief]
    - routes [family] prefix [detail | longer | hunt [brief]]
    - routes [family] community comm-id
    - routes l2-vpn l2vpn-type {[rd rd] | [siteid site-id] | [veid veid] [offset vpls-
base-offset]}
    - summary [all]
    - summary [family family] [neighbor ip-address]
```

6.23.1.7 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
```

6.23.1.8 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
```

6.23.2 Command descriptions

- [Configuration commands](#)
- [Other BGP-related commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

6.23.2.1 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

This command enables the add-paths configuration context and enables add-paths to be configured for one or more BGP route types. The BGP add-paths capability allows the router to send and receive multiple paths per prefix to and from a peer.

The **no** form of this command (**no add-paths**) removes add-paths from the configuration of BGP, the group, or the neighbor, causing sessions established using add-paths to go down and come back up without the add-paths capability.

Default

no add-paths

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-paths capability for IPv4 routes. By default, the add-paths capability is disabled for IPv4 routes.



Note:

Add-paths are supported only for the label-IPv4 family.

The maximum number of paths to send per IPv4 NLRI is the configured *send-limit*, which is a mandatory parameter. The capability to receive multiple paths per prefix from a peer is configurable using the optional

receive keyword. If the **receive** keyword is not included in the command, the receive capability is enabled by default.

The **no** form of this command disables add-paths support for 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 are allowed to 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

Keyword to specify that the router negotiates the add-paths receive capability for IPv4 routes with its peers.

none

Keyword to specify 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-paths capability for IPv6 routes. By default, the add-paths capability is disabled for IPv6 routes.

**Note:**

Add-paths are supported only for the label-IPv4 family.

The maximum number of paths to send per IPv6 NLRI is the configured *send-limit*, which is a mandatory parameter. The capability to receive multiple paths per prefix from a peer is configurable using the optional **receive** keyword. If the **receive** keyword is not included in the command, the receive capability is enabled by default.

The **no** form of this command disables add-paths support for IPv6 routes, causing sessions established using add-paths for IPv6 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 IPv6 NLRI that are allowed to 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

Keyword to specify that the router negotiates the add-paths receive capability for IPv6 routes with its peers.

none

Keyword to specify that the router does not negotiate the add-paths receive capability for IPv6 routes with its peers.

vpn-ipv4**Syntax**

vpn-ipv4 send *send-limit* **receive** [**none**]

vpn-ipv4 send *send-limit*

no vpn-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-paths capability for VPN-IPv4 routes. By default, add-paths is not enabled for VPN-IPv4 routes.

The maximum number of paths per VPN-IPv4 NLRI to send is the configured *send-limit*, which is a mandatory parameter. The capability to receive multiple paths per prefix from a peer is configurable using the optional **receive** keyword. If the **receive** keyword is not included in the command, the receive capability is enabled by default.

The **no** form of this command disables add-paths support for VPN-IPv4 routes, causing sessions established using add-paths for VPN-IPv4 to go down and come back up without the add-paths capability.

Default

no vpn-ipv4

Parameters

send-limit

Specifies the maximum number of paths per VPN-IPv4 NLRI that are allowed to be advertised to add-paths peers (the actual number of advertised routes may be less depending on the next-hop diversity requirement, other configuration options, route policies, or route advertisement rules).

Values 1 to 16, none

Default max

receive

Keyword to specify the router negotiates the add-paths receive capability for VPN-IPv4 routes with its peers.

none

Keyword to specify the router does not negotiate the add-paths receive capability for VPN-IPv4 routes with its peers.

vpn-ipv6

Syntax

vpn-ipv6 send *send-limit* **receive** [**none**]

vpn-ipv6 send *send-limit*

no vpn-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-paths capability for VPN-IPv6 routes. By default, add-paths is not enabled for VPN-IPv6 routes.

The maximum number of paths per VPN-IPv6 NLRI to send is the configured *send-limit*, which is a mandatory parameter. The capability to receive multiple paths per prefix from a peer is configurable using the optional **receive** keyword. If the **receive** keyword is not included in the command the receive capability is enabled by default.

The **no** form of this command disables add-paths support for VPN-IPv6 routes, causing sessions established using add-paths for VPN-IPv6 to go down and come back up without the add-paths capability.

Default

no vpn-ipv6

Parameters

send-limit

Specifies the maximum number of paths per VPN-IPv6 NLRI that are allowed to be advertised to add-paths peers (the actual number of advertised routes may be less depending on the next-hop diversity requirement, other configuration options, route policies and/or route advertisement rules).

Values 1 to 16, none

Default max

receive

Keyword to specify the router negotiates the add-paths receive capability for VPN-IPv6 routes with its peers.

none

Keyword to specify the router does not negotiate the Add-Paths receive capability for VPN-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 **use-svc-routes** parameter limits the number of BGP 3107 IPv4 labeled routes that are installed in the MPLS FIB. If this parameter is enabled, only those 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 conditions trigger the installation of the MPLS label into the MPLS FIB for the received BGP 3107 IPv4 labeled route:

- configuration of SDP to use a 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 a PE that is reachable using the labeled route
- configuration of a BGP session to a BGP peer, using the **bgp>neighbor** CLI command, and the BGP peer is reachable using the labeled route
- on the 7210 SAS-R6 and 7210 SAS-R12, if this command is enabled along with the **config router bgp enable-rr-vpn-forwarding** command, only BGP 3107 routes required to resolve the VPN-IPv4 or VPN-IPv6 routes, and for which the 7210 SAS node is swapping the VPN label, are added to the FIB (the rest are only held in the RIB)
- other IP applications, such as FTP and SSH, do not trigger the installation of the IPv4 labeled routes into the MPLS FIB

The **no** form of this command disables all configured options.

Default

no advertise-label

Parameters

ipv4

Keyword to specify the advertisement label address family for core IPv4 routes. This keyword can be specified only for an IPv4 peer.

use-svc-routes

Keyword to enable the user to limit the number of BGP 3107 labeled routes that 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. This parameter is only supported on the 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE.

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, while this command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of this command used at the global level reverts to default where BGP adds the 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

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

This command enables accumulated Interior Gateway Protocol (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.

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

Default

MD5 Authentication disabled

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

Keyword to specify that the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Keyword to specify that the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

backup-path

Syntax

[no] backup-path [ipv4]

Context

config>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the computation and use of a backup path for labeled IPv4 BGP-learned prefixes belonging to the base router. Multiple paths must be received for a prefix to take advantage of this feature. When 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. When 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 prefixes do not have a backup path installed in the IOM.

The **no** form of this command disables the use of a backup path.

Default

no backup-path

Parameters

ipv4

Keyword that enables BGP fast reroute for labeled IPv4 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 Multi-Exit Discriminator (MED) path attribute.

The default behavior of 7210 SAS (equivalent to the **no** form of this command) is to only compare two routes on the basis of MED if they have the same neighbor AS (the first non-confed AS in the received AS_PATH attribute). Also by default, a route without a MED attribute is handled the same as though it had a MED attribute with the value 0.

This 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 the **zero** nor **infinity** keyword is specified, the **zero** option is inferred, meaning a route without a MED is handled as though it had a MED attribute with the value 0. When the **strict-as** keyword is present, MED is only compared between paths from the same neighbor AS; 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 only compares two routes on the basis of MED if they have the same neighbor AS.

Default

no always-compare-med

Parameters

zero

Keyword to specify 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

Keyword to specify 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

Keyword to specify that BGP paths are to be compared even with different neighbor AS.

as-path-ignore

Syntax

as-path-ignore [ipv4] [vpn-ipv4] [l2-vpn] [mvpn-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 specifies whether the AS path is used to determine the best BGP route.

If this command is enabled, the AS paths of incoming routes are not used in the route selection process. The **no** form of this command removes the parameter from the configuration.

Default

no as-path-ignore

Parameters

ipv4

Keyword to specify that the AS-path length is ignored for all IPv4 routes.

vpn-ipv4

Keyword to specify that the length AS-path is ignored for all IPv4 VPRN routes.

l2-vpn

Keyword to specify that the AS-path length is ignored for all L2-VPN NLRIs.

mvpn-ipv4

Keyword to specify that the AS-path length is ignored for all mVPN IPv4 multicast routes. Supported on 7210 SAS-T, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp devices only.

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 configures 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 reverts to the default behavior whereby 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 the 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 do not change the best path even if y is less than x during BGP identifier comparison.

The **no** form of this command reverts to the default behavior of selecting the route with the lowest BGP identifier (y) as best.

Default

no ignore-router-id

cluster

Syntax

cluster *cluster-id*

no cluster

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

7210 SAS-R6 and 7210 SAS-R12

Description

This command configures the cluster ID for a route reflector server.

Route reflectors reduce the number of iBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in an AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients.

When a route reflector receives a route, it must first select the best path from all the paths received. If the route was received from a non-client peer, the route reflector sends the route to all clients in the cluster. If the route came from a client peer, the route reflector sends the route to all non-client peers and to all client peers except the originator.

For redundancy, a cluster can have multiple route reflectors.

The **no** form of this command deletes the cluster ID and disables the route reflection at the global BGP level or for the specified group or neighbor.

Default

no cluster

Parameters

cluster-id

Specifies the route reflector cluster ID, expressed in dotted-decimal notation.

Values	Any 32-bit number in dotted-decimal notation. (0.0.0.1 to 255.255.255.255)

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.

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 that 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.

The **no** form of this command used at the global level reverts route damping.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters:

Half-life:	15 minutes
Max-suppress:	60 minutes
Suppress-threshold:	3000
Reuse-threshold:	750

Default

no damping

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 enables 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 three level of the hierarchy so it can be specified down to the per peer basis.

If this command is enabled, 4-byte ASN support should not be negotiated with the associated remote 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 reverts to the default behavior.

Default

no disable-capability-negotiation

disable-client-reflect

Syntax

[no] **disable-client-reflect**

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

7210 SAS-R6 and 7210 SAS-R12

Description

This command disables the reflection of routes by the route reflector at the BGP, group, or neighbor level.

This command only disables the reflection of routes from other client peers. Routes learned from non-client peers are still reflected to all clients.

The **no** form of this command re-enables client reflection of routes.

Default

no disable-client-reflect

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

Keyword to specify standard communities that existed before VPRNs or 2547.

extended

Keyword to specify 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 the use of the IGP next hop to the BGP next hop as the next hop of last resort.

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 does not wait for the hold timer to expire; therefore, the BGP reconvergence process is accelerated.

The **no** form of this command disables peer tracking.

Default

no enable-peer-tracking

enable-rr-vpn-forwarding

Syntax

[no] enable-rr-vpn-forwarding

Context

config>router>bgp

Platforms

7210 SAS-R6 and 7210 SAS-R12

Description

This command enables a route reflector of VPN-IP routes to be deployed in the datapath between two BGP peers (a peer X and a peer Y) in a next-hop resolution.



Note:

Scaling and convergence should be considered before enabling this command.

When this command is configured, all received VPN-IP routes, regardless of route target, are imported into the dummy VRF, where the BGP next hops are resolved. The **label-route-transport-tunnel** command in the **config>router>bgp>next-hop-resolution** context determines what types of tunnels are eligible to resolve the next hops.

If a received VPN-IP route from iBGP peer X is resolved and selected as best so that it can be readvertised to an iBGP peer Y, and the BGP next hop is modified toward peer Y (by using the **next-hop-self** command in the Y group or neighbor context, or by using the **next-hop action** command in an export policy applied to Y), BGP allocates a new VPN service label value for the route, signals that new label value to Y, and programs the IOM to do the corresponding label swap operation. The supported combinations of X and Y are the following:

- from X (client) to Y (client)
- from X (client) to Y (non-client)
- from X (non-client) to Y (client)

The **no** form of this command causes the route to be readvertised without a new service label, or a new service label to not be advertised between the two peers.

Default

no enable-rr-vpn-forwarding

export

Syntax

export *policy-name* [*policy-name*...up to 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. 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 15 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.

family

Syntax

family [ipv4] [vpn-ipv4] [ipv6] [vpn-ipv6] [l2-vpn] [ms-pw] [mvpn-ipv4] [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 family or families to support over BGP peerings in the base router. This command is additive; entering 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 supported address family is reset back to the default.

Default

family ipv4

Parameters

ipv4

Keyword to provision support for IPv4 routing information.

vpn-ipv4

Keyword word to exchange IPv4 VPN routing information.

ms-pw

Keyword to exchange dynamic MS-PW related information.

ipv6

Keyword to exchange IPv6 routing information.

vpn-ipv6

Keyword to exchange IPv6 VPN routing information.

l2-vpn

Keyword to exchange Layer 2 VPN information.

mvpn-ipv4

Keyword to exchange multicast VPN related information. This keyword is supported only on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), and 7210 SAS-T (network). This family is not supported for 7210 SAS-Sx 10/100GE devices.

route-target

Keyword to exchange RT constrained route information.

split-horizon

Syntax

[no] split-horizon

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

7210 SAS-R6 and 7210 SAS-R12

Description

This command enables the use of split horizon. Split horizon prevents routes from being reflected back to a peer that sends the best route. It applies to routes of all address families and to eBGP or iBGP sending peers.

The **no** form of this command means that no effort is taken to prevent a best route from being reflected back to the sending peer.

Default

no split-horizon

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 causes the base instance BGP import route policies to be applied 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 time BGP waits between successive messages (either keepalive or update) from its peer, before closing the connection. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in group), or neighbor level (only applies to specified peer). The most specific value is used.

Even though the 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

Keyword to specify that 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 use 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 15 policy names can be specified. The first policy that matches is applied.

When multiple **import** commands are issued, the last command entered overrides the previous command.

When an import policy is not specified, BGP routes are accepted by default.

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

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. The 7210 SAS implementation allows the **keepalive** value and the **hold-time** interval to be independently set; however, 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 **keepalive** is set to one-third of the current **hold-time** value.
- If the specified **hold-time** interval is less than the configured **keepalive** value, 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 the **keepalive** value is ignored. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

keepalive 30

Parameters

seconds

Specifies the keepalive timer, in seconds, expressed as a decimal integer.

Values 0 to 21845

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 address expressed in dotted decimal notation. Allowed value is a valid routable IP address on the router, either an interface or system IP address.

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

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 command can be configured 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). By specifying this command at each neighbor level, it is possible to have a separate AS number for each eBGP session.

When a command is entered multiple times for the same AS, the last command entered is used in the configuration. Add or remove the **private** keyword 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 reestablish 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 reestablish the peer relationship with the new local ASN.

This is an optional command and can be used in a circumstance like the following 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 removes 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 AS number, expressed as a decimal integer.

Values 1 to 65535

private

Keyword to specify that 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 command is used if the BGP route arrives from a BGP peer without the **local-preference** command configured.

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 use 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 command 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 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

loop-detect ignore-loop

Parameters

drop-peer

Keyword that sends a notification to the remote peer and drops the session.

discard-route

Keyword that discards 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 is not active for an established peer until the connection is reestablished for the peer.

ignore-loop

Keyword that ignores routes with loops in the AS path but maintains peering.

off

Keyword that 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 MED and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set.

The specified value can be overridden by any value set via a route policy.

This command 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

Keyword to specify 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 at which a path attribute, originated by the local router, can be advertised to a peer.

This command 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 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

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 at which a prefix can be advertised to a peer.

This command 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 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

min-route-advertisement 30

Parameters

seconds

Specifies the minimum route advertising interval, in seconds, expressed as a decimal integer.

Values 1 to 255

mp-bgp-keep

Syntax

[no] mp-bgp-keep

Context

config>router>bgp

Description

When this command is enabled, route refresh messages are not required or issued when VPN route policy changes are made; RIB-IN retains all MP-BGP routes.

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

Default

no mp-bgp-keep

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 conveys to the BGP instance that the eBGP peers are directly connected.

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

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-self

Syntax

[no] next-hop-self

Context

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

7210 SAS-R6 and 7210 SAS-R12

Description

This command configures BGP to advertise routes to members of a group or to a specific neighbor using a local address of the BGP instance as the BGP next-hop address. This command is set regardless of the route source (eBGP or iBGP) or its family. When used with VPN-IPv4 and VPN-IPv6 routes, the **enable-rr-vpn-forwarding** command should also be configured.

The **no** form of this command uses standard protocol behavior to decide whether to set **next-hop-self** in advertised routes.

Default

no next-hop-self

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 label route context following 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 label 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 used in the resolution of BGP label routes using tunnels to BGP peers.

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

7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

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.

Default

no sr-isis

sr-ospf

Syntax

[no] sr-ospf

Context

config>router>bgp>next-hop-res>lbl-rt-tunn>family>resolution-filter

Platforms

7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

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 OSPF instances with the same lowest preference, the OSPF instance with the lowest ID number.

Default

no sr-ospf

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

Commands in this context configure sending or accepting extended-community-based BGP filters.
For the **no** form of this 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 ORF negotiation with a peer and to accept filters that the peer wants to send.

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

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 outbound route filtering (ORF) negotiation with a peer.

This command also causes the router to send a community filter, prefix filter, or AS path filter configured as an inbound filter on the BGP session to its peer as an ORF Action ADD.

The **no** form of this command causes the router to remove the send capability in the BGP ORF negotiation with a peer.

The **no** form also causes the router to send an ORF remove action for a community filter, prefix filter, or AS path filter configured as an inbound filter on the BGP session to its peer.

If the *comm-id* parameters are not exclusively route target communities, the router extracts appropriate route targets and uses those. If the *comm-id* parameters specified contain no route targets, the router does not send an ORF.

Default

no send-orf

Parameters

comm-id

Specifies a community policy that consists exclusively of route target extended communities. If the policy is not specified, the ORF policy is automatically generated from configured route target lists, accepted client route target ORFs, and locally configured route targets.

Values target: {*ip-addr* : comm-val | *2byte-asnumber* : ext-comm-val | *4-byte-asnumber* : comm-val}
 ip-addr: a.b.c.d
 comm-val: 0 to 65535
 2byte-asnumber: 0 to 65535
 ext-comm-val: 0 to 4294967295
 4byte-asnumber: 0 to 4294967295

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 removes the specified neighbor and the entire configuration associated with the neighbor. The neighbor must be administratively shut down before attempting to delete it. If the neighbor is not shut down, the command does not result in any action except a warning message on the console indicating that the neighbor is still administratively up.

Parameters

ip-address

Specifies the IP address of the BGP peer router, in dotted decimal notation.

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

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 AS number for the remote peer. The peer AS number must be configured for each configured peer.

For eBGP peers, the peer AS number configured must be different from the AS number configured for this router under the global level, because the peer is in a different AS than this router.

For iBGP peers, the peer AS number must be the same as the AS number of this router configured under the global level.

This is a required command for each configured peer. This command can be configured under the group level for all neighbors in a particular group.

Parameters

as-number

Specifies the AS 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 is initially set to the egress interface MTU. The DF bit is also set so that if a router along the path of the TCP connection cannot handle a packet of a particular size without fragmenting, an ICMP message is sent back 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 command 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 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 minutes

rapid-update

Syntax

rapid-update [*I2-vpn*] [*mvpn-ipv4*] [*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.

mvpn-ipv4

Keyword to enable the BGP rapid update for the MVPN-IPv4 address family. The MVPN-IPv4 address is a variable size value consisting of the 1-byte route type, 1-byte length, and variable size that is route type specific. Route type defines encoding for the route type specific field. Length indicates the length in octets of the route type specific field.

This keyword is supported only on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), and 7210 SAS-T (network). This family is not supported for 7210 SAS-Sx 10/100GE devices.

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. This keyword is supported only on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

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 this command from the configuration and reverts withdrawal processing to the normal behavior.

Default

no rapid-withdrawal

prefix-limit

Syntax

prefix-limit *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 sent. When the limit is exceeded, the BGP peering is dropped and disabled.

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

Default

no prefix-limit

Parameters

limit

Specifies the number of routes that can be learned from a peer, expressed as a decimal integer.

Values 1 to 4294967295

log-only

Keyword that enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, the BGP peering is not dropped.

percentage

Specifies the threshold value (as a percentage) that triggers a warning message to be sent.

Values 1 to 100

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 ASNs to be removed from the AS path before advertising them to BGP peers.

When this command is set at the global level, it applies to all peers regardless of group or neighbor configuration. When the command 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 ASNs that are defined by IANA as private. These are ASNs in the range of 64512 through 65535, inclusive.

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

no remove-private

Parameters

limited

Optional keyword that removes private ASNs up to the first public ASN encountered, at which point it 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, no router ID is configured for BGP; the system interface IP address is used.

Parameters

ip-address

Specifies the router ID, expressed in dotted decimal notation. Allowed value is a valid routable IP address on the router, either an interface or system IP address. Nokia highly recommends that this address be the system IP address.

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 Protocol Handling

On all 7210 SAS platforms, BGP is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure router bgp** command instantiates the protocol in the **no shutdown** state, and resources are allocated to enable the node to process the protocol.

To deallocate resources, issue the **configure router bgp shutdown** and **configure router no bgp** commands to allow the node to boot up correctly after the reboot. It is not sufficient to only issue a **configure router bgp shutdown** command.

The resources for BGP are allocated when the BGP context is enabled either in the base routing instance or the VPRN service instance. Resources are deallocated when the configuration of the last BGP context under either base routing instances or VPRN service is removed or shut down.

BGP Global

The BGP protocol is created in the **no shutdown** state on all 7210 SAS platforms.

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.

Specifying the **internal** parameter type indicates the peer is an iBGP peer; specifying the **external** parameter type indicates 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 that configures the peer as internal.

external

Keyword that configures the peer as external.

6.23.2.2 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 autonomous system number (ASN) 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 AS number is not used until the BGP instance is restarted either by administratively disabling and enabling (**shutdown** or **no shutdown**) the BGP instance or rebooting the system with the new configuration.

Parameters

as-number

Specifies the AS number, 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 configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized, the new router ID is used. This can result in an interim period of time when different protocols use different router IDs.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for each protocol that uses the router ID, or restart the entire router.

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 to reverts to the default value.

Parameters

ip-address

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

6.23.2.3 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-name*: Base, management
 service-id: 1 to 2147483647

Default Base

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 particular authentication keychain.

Parameters

keychain

Specifies an existing keychain name, up to 32 characters.

Output

The following outputs are examples of BGP authentication keychain information, and [Table 84: Output fields: BGP PIC](#) describes the output fields.

- [Sample output](#)
- [Sample output: BGP PIC](#)

Sample output

```
*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
10.1.0.2              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
Cluster Id         : 0.0.0.100      Client Reflect  : Enabled
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#
```

Sample output: BGP PIC

```
*A:7210SAS>show>service>id# show service id 1 base
=====
Service Basic Information
=====
Service Id         : 1              Vpn Id          : 0
Service Type       : VPRN
Name               : (Not Specified)
Description        : Default Description For VPRN ID 1
Customer Id        : 1
Last Status Change: 01/08/2000 22:57:35
Last Mgmt Change   : 01/08/2000 22:57:35
Admin State        : Up             Oper State       : Up
Route Dist.        : 100:1          VPRN Type       : regular
```

```

AS Number      : 100          Router Id       : 10.1.1.1
ECMP           : Enabled      ECMP Max Routes : 1
Max IPv4 Routes : No Limit    Auto Bind      : MPLS
Max IPv6 Routes : No Limit
Ignore NH Metric : Disabled
Hash Label     : Disabled
Vrf Target     : target:200:1
Vrf Import     : None
Vrf Export     : None
MVPN Vrf Target : None
MVPN Vrf Import : None
MVPN Vrf Export : None
Label mode     : vrf
BGP VPN Backup  : ipv4 ipv6

```

```

SAP Count      : 1          SDP Bind Count   : 3

```

Service Access & Destination Points

Identifier	Type	AdmMTU	OprMTU	Adm	Opr
sap:1/1/2:1	q-tag	9212	9212	Up	Up
sdp:1002:1 S(2.2.2.2)	Spok	0	9186	Up	Up
sdp:1003:1 S(3.3.3.3)	Spok	0	9186	Up	Up
sdp:1004:1 S(4.4.4.4)	Spok	0	9186	Up	Up

=====

```

*A:7210SAS>show>service>id#

```

Table 84: Output fields: BGP PIC

Label	Description
Service Id	Displays the service ID
VPN Id	Displays the VPN ID
Service Type	Displays the service type
Name	Displays the service name
Description	Displays the description of the service
Customer Id	Displays the customer ID
Last Status Change	Displays the date and time of the most recent change in the administrative or operating status of the service
Last Mgmt Change	Displays the date and time of the most recent management-initiated change to this service
Admin State	Displays the required state of the service
Oper State	Displays the current operational state of the service
Route Dist.	Displays the route distribution number
VPRN Type	Only valid in services that accept mesh SDP bindings

Label	Description
	It validates the VC ID portion of each mesh SDP binding defined in the service.
AS Number	Displays the autonomous system number
Router Id	Displays the router ID for this service
ECMP Max Routes	Displays the maximum number of routes that can be received from the neighbors in the group or for the specific neighbor
ECMP	Displays equal cost multipath information
Max IPv4 Routes	Displays the maximum number of IPv4 routes that can be used for path sharing
Auto Bind	Displays the automatic binding type for the SDP assigned to this service
Max IPv6 Routes	Displays the maximum number of IPv6 routes that can be used for path sharing
Ignore NH Metric	Indicates whether ignore NH metric is enabled or disabled
Hash Label	Indicates whether the hash label is enabled or disabled
Vrf Target	Displays the route target in the VRF applied to this service
Vrf Import	Displays the VRF import policy applied to this service
Vrf Export	Displays the VRF export policy applied to this service
MVPN Vrf Target	Displays the route target in the MVPN VRF applied to this service
MVPN Vrf Import	Displays the MVPN VRF import policy applied to this service
MVPN Vrf Export	Displays the MVPN VRF export policy applied to this service
Label mode	Displays the label mode
BGP VPN Backup	Indicates whether the BGP VPN backup is enabled or disabled
SAP Count	Displays the number of SAPs specified for this service
SDP Bind Count	Displays the number of SDPs bound to this service
Service Access and Destination Points	
Identifier	Displays the SAP or SDP identifier
type	Indicates whether this service SDP binding is a spoke or a mesh

Label	Description
AdmMTU	Displays the required largest service frame size (in octets) that can be transmitted through this SAP or SDP to the far-end router, without requiring the packet to be fragmented
OprMTU	Displays the actual largest service frame size (in octets) that can be transmitted through this SAP or SDP to the far-end router, without requiring the packet to be fragmented
Adm	Displays the administrative state of this SAP or SDP
Opr	Displays the operational state of this SAP or SDP

damping

Syntax

damping [*ip-prefix* [*ip-prefix-length*]] [*damp-type*] [**detail**]

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 because of route flapping. This command can be entered with or without a route parameter.

When the **detail** keyword is included, more detailed information displays.

When only the command is entered (without any parameters included except **detail**), all dampened routes are listed.

When a parameter is specified, the matching route or routes are listed.

When a **decayed**, **history**, or **suppressed** keyword is specified, only those types of dampened routes are listed.

Parameters

ip-prefix/ip-prefix-length

Displays damping information for the specified IP prefix and length.

Values

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

damp-type

Specifies the type of damping to display.

- Values**
- decayed** — Displays damping entries that are decayed but are not suppressed.
 - history** — Displays damping entries that are withdrawn but have history.
 - suppressed** — Displays damping entries suppressed because of route damping.

detail

Displays detailed information.

Output

The following outputs are examples of BGP damping information, and [Table 85: Output fields: 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 : 10.32.27.203
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 : 10.32.27.203
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 : 10.32.27.203
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 : 15.203.192.0/18
-----
Network      : 15.203.192.0/18   Peer      : 10.0.28.1
NextHop      : 10.0.28.1         Reuse time : 00h00m00s
Peer AS      : 60203             Peer Router-Id : 10.32.27.203
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 15.203.192.0/18 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 10.203.192.0/18
=====
-----
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 : 10.32.27.203
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 : 10.32.27.203
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 : 10.32.27.203
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 : 10.32.27.203
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 : 10.32.27.203
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 85: Output fields: BGP damping

Label	Description
BGP Router ID	Displays the local BGP router ID
Local AS	Displays the configured autonomous system number
Network	Displays the route IP prefix and mask length for the route
Flags	Legend: Status codes: u- used, s-suppressed, h-history, d-decayed, *-valid. If a "*" is not present, 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
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 autonomous system number of the advertising router

Label	Description
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 hour/minute/second (HH:MM:SS) format
Last update	Displays the time when BGP was updated last in 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
Reuse time	Displays the time when the route can be reused
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 this command is issued with a specific group name, information pertaining only to 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**
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.

Parameters

name

Displays information for the BGP group specified, up to 32 characters.

detail

Displays detailed information.

Output

The following outputs are examples of BGP group information, and [Table 86: Output fields: BGP group](#) describes the output fields.

- [Sample detailed output](#)
- [Sample output](#)

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           State           : Up
Peer AS         : 40000             Local AS        : 65206
Local Address    : n/a              Loop Detect     : Ignore
Connect Retry   : 20                Authentication : None
Local Pref      : 100               MED Out        : 0
Multihop        : 0 (Default)
Min Route Advt. : 30                Min AS Originate : 15
Prefix Limit    : No Limit          Passive         : Disabled
Next Hop Self   : Disabled          Aggregator ID 0 : Disabled
Remove Private  : Disabled          Damping        : Disabled
Export Policy    : direct2bgp
Hold Time       : 90                Keep Alive      : 30
Cluster Id      : None              Client Reflect  : Enabled
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#

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      State      : Up
Peer AS        : n/a          Local AS   : 100
Local Address   : n/a          Loop Detect : Ignore
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           Keep Alive   : 30
Cluster Id     : None          Client Reflect : Disabled
NLRI           : Unicast        Preference   : 101
TTL Security   : Disabled       Min TTL Value : n/a
Graceful Restart : Disabled      Stale Routes Time: n/a
Auth key chain  : n/a           Bfd Enabled   : Yes

List of Peers
- 3.3.3.3 :
  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#

```

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
Cluster Id   : None          Client Reflect : Enabled
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#

```

Table 86: Output fields: 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 ASNs from the AS path attribute in updates sent to the neighbor in the peer group Enabled — BGP removes all private ASNs 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
Cluster Id	Displays the configured route reflector cluster ID None — no cluster ID has been configured
Client Reflect	Disabled — the BGP route reflector does 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

Label	Description
	Unicast — IPv4 unicast routing information can be carried
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 [*as-number* [**detail**] *filter2*]
neighbor *ip-address* [*family* [**type** *mvpn-type*]] *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 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 about 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 is listed (see second output example).



Note:

This information is not available when using SNMP.

When either **history** or **suppressed** is specified, the routes learned from those peers that either have a history or are suppressed (respectively) 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 the peer is operationally disabled and must be restarted.

Parameters

ip-address

Displays information for the specified IP address.

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

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 and not those capable of exchanging IP-VPN routes.

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 (ORF) for the BGP instance. ORF 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 or 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 example of BGP neighbor information. The associated tables describe the output fields.

- [Sample output](#), [Sample output — detailed](#), [Table 87: Output fields: BGP neighbor](#)
- [Sample output — received routes](#), [Table 88: Output fields: BGP neighbor received routes](#)
- [Sample output — add-path](#), [Table 89: Output fields: show neighbor add-path](#)

Sample output

```
*A:7210-SAS>show>router>bgp# neighbor

=====
BGP Neighbor
=====
-----
Peer   : 10.1.1.1
Group  : sample
-----
Peer AS      : 12345      Peer Port      : 0
Peer Address : 10.1.1.1
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   : 10.1.3.4
Group  : test
-----

```

```

Peer AS      : 0          Peer Port      : 0
Peer Address : 10.1.3.4   Local Port     : 0
Local AS     : 12345
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
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  : 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: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 Port       : 179
Peer Address  : 10.20.1.3
Local AS      : 300          Local Port    : 49635
Local Address : 10.20.1.5
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     : 90           Keep Alive    : 30
Min Hold Time : 0
Active Hold Time : 90       Active Keep Alive : 30
Cluster Id    : None
Preference    : 170         Num of Update Flaps : 20
Recd. Paths   : 5
IPv4 Recd. Prefixes : 10    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    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     : 30        o/p Messages      : 26
i/p Octets       : 1321      o/p Octets        : 470
i/p Updates      : 8         o/p Updates       : 0
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     : 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              IPv6 Rej. Pfxs         : 0
VPN-IPv4 Rej. Pfxs    : 0              VPN-IPv6 Rej. Pfxs    : 0
Mc IPv4 Rej. Pfxs     : 0              Mc IPv6 Rej. Pfxs     : 0
MVPN-IPv4 Rej. Pfxs   : 0              MVPN-IPv6 Rej. Pfxs   : 0
Flow-IPv4 Rej. Pfxs   : 0              Flow-IPv6 Rej. Pfxs   : 0
L2-VPN Rej. Pfxs      : 0              MDT-SAFI Rej. Pfxs    : 0
Rte-Tgt Rej. Pfxs     : 0              MS-PW Rej. Pfxs      : 0
Mc Vpn Ipv4 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 State          : Established
Last Event        : rcvKeepAlive
Last Error        : Cease
Local Family      : IPv4
Remote Family     : IPv4
Hold Time         : 3                Keep Alive          : 1
Active Hold Time  : 3                Active Keep Alive   : 1
Cluster Id        : None
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 output — detailed

```

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 AS     : 65206      Local Port   : 0
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
Cluster Id    : None     Client Reflect : Enabled

```

```

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#

*A:SetupCLI>show>router>bgp# neighbor
=====
BGP Neighbor
=====
Peer   : 10.3.3.3
Group  : bgp_group_1 34567890123456789012
-----
Peer AS      : 20                      Peer Port       : 0
Peer Address : 10.3.3.3                 Local Port      : 0
Local AS     : 100                      Local Address   : 0.0.0.0
Peer Type    : Internal                 Last State      : Idle
State        : Active
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
Cluster Id   : 2.2.3.4
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
=====

```

Table 87: Output fields: 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 Established — BGP has successfully established a peering and is exchanging routing information
Last 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
Last Event	start — BGP has initialized the BGP neighbor stop — BGP has disabled the BGP neighbor open — BGP transport connection opened close — BGP transport connection closed openFail — BGP transport connection failed to open error — BGP transport connection error connectRetry — connect retry timer expired holdTime — hold time timer expired keepAlive — keepalive timer expired recvOpen — receive an OPEN message revKeepalive — receive a KEEPALIVE message recvUpdate — receive an UPDATE message recvNotify — receive a NOTIFICATION message None — no events have occurred
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.
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

Label	Description
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 next hop in route updates to the specified neighbor Enabled — BGP sends only its own IP address as the BGP next hop 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 does not remove all private ASNs from the AS path attribute, in updates sent to the specified neighbor Enabled — BGP removes all private ASNs from the AS path attribute, in updates sent to the specified neighbor
Passive	Disabled — BGP actively attempts to establish a BGP connection with the specified neighbor Enabled — BGP does 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 to 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
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 has been configured
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

Label	Description
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 as a result of 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 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 — 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 88: Output fields: 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:7210SAS# show router bgp neighbor 10.2.2.2

=====
BGP Neighbor
=====
-----
Peer   : 10.2.2.2
Group  : toPE
-----
Peer AS      : 100          Peer Port      : 50854
Peer Address : 10.2.2.2
Local AS     : 100          Local Port    : 179
Local Address: 10.1.1.1
Peer Type    : Internal
State       : Established   Last State    : Established
Last Event   : recvKeepAlive
Last Error   : Cease (Connection Collision Resolution)
Local Family : IPv4 VPN-IPv4 IPv6 VPN-IPv6
```

```

Remote Family      : IPv4 VPN-IPv4 IPv6 VPN-IPv6
Hold Time          : 90                      Keep Alive          : 30
Min Hold Time      : 0
Active Hold Time    : 90                      Active Keep Alive       : 30
Cluster Id         : None
Preference          : 170                      Num of Update 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                    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-SAFI Suppr* : 0                    FLOW-IPv4-SAFI Recd.* : 0
FLOW-IPv4-SAFI Activ* : 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
Backup Vpn IPv4 Pfxs : 0                      Backup Vpn IPv6 Pfxs   : 0
Input Queue           : 0                      Output Queue           : 0
i/p Messages          : 9042                    o/p Messages           : 65
i/p Octets             : 111                     o/p Octets             : 278
i/p Updates            : 0                      o/p Updates            : 0
TTL Security          : Disabled                  Min TTL Value          : n/a
Graceful Restart       : Disabled                  Stale Routes Time      : n/a
Advertise Inactive     : Disabled                  Peer Tracking          : Disabled
Advertise Label        : ipv4 ipv6
Auth key chain         : n/a
Disable Cap Nego       : Disabled                  Bfd Enabled            : Enabled
Flowspec Validate      : Disabled                  Default Route Tgt      : Disabled
L2 VPN Cisco Interop   : Disabled
Local Capability       : RtRefresh MPBGP 4byte ASN
Remote Capability      : RtRefresh MPBGP 4byte ASN
Local AddPath Capabi* : Send - VPN-IPv4 (1) VPN-IPv6 (4)
                       : Receive - VPN-IPv6
Remote AddPath Capab* : Send - VPN-IPv6
                       : Receive - VPN-IPv4 VPN-IPv6
Import Policy          : None Specified / Inherited
Export Policy          : P1

```

Neighbors : 1

=====

* indicates that the corresponding row element may have been truncated.

*A:7210SAS#

*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 89: 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 (shut down) 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 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 Established — BGP has successfully established a peering session and is exchanging routing information
Last 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 with this peer Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer

Label	Description
	Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION
Last Event	start — BGP has initialized the BGP neighbor stop — BGP has disabled the BGP neighbor open — BGP transport connection is opened close — BGP transport connection is closed openFail — BGP transport connection failed to open error — BGP transport connection error connectRetry — the connect retry timer expired holdTime — the hold time timer expired keepAlive — the keepalive timer expired rcvOpen — BGP has received an OPEN message revKeepalive — BGP has received a KEEPALIVE message rcvUpdate — BGP has received an UPDATE message rcvNotify — BGP has received a NOTIFICATION message None — no events have occurred
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
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

Label	Description
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 because of route damping
VPN-IPv4 Suppr. Pfxs	Displays the number of unique sets of VPN-IPv4 path attributes received from the BGP neighbor and suppressed because of 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
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 as a result of route damping
Backup IPv4 Pfxs	Displays the number of BGP Fast Reroute backup path IPv4 prefixes
Backup IPv6 Pfxs	Displays the number of BGP Fast Reroute backup path IPv6 prefixes
Backup Vpn IPv4 Pfxs	Displays the number of BGP Fast Reroute backup path VPN IPv4 prefixes

Label	Description
Backup Vpn IPv6 Pfxs	Displays the number of BGP Fast Reroute 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)
Peer Tracking	Displays the state of tracking a neighbor IP address in the routing table for a BGP session
Advertise Label	Indicates 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

Label	Description
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

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 and not those capable of exchanging IP-VPN routes.
 - vpn-ipv4** — Displays the BGP peers that are IP-VPN capable.
 - mcast-ipv4** — Displays the BGP peers that are mcast-ipv4 capable.

ip-address

Displays the next-hop information for the specified IP address.

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

detail

Displays a more detailed version of the output.

Output

The following output is an example of BGP next-hop information, and [Table 90: Output fields: 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 90: Output fields: BGP next-hop

Label	Description
BGP 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, 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

Label	Description
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 attribute information, and [Table 91: Output fields: 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 91: Output fields: BGP paths

Label	Description
BGP Router ID	Displays the local BGP router ID
AS	Displays the configured AS number
Local AS	Displays the configured local AS setting. If not configured, 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 Multi-Exit Discriminator 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 using a route policy.
Refs	Displays the number of routes using a specified set of path attributes
ASes	Displays the number of AS 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
Cluster List	Displays the route reflector cluster list

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]}
```

Context
show>router>bgp

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

Description
This command displays BGP route information.
If this command is issued without any parameters, the entire BGP routing table is displayed.
If this command is issued with an IP prefix/mask or IP address, the best match for the parameter is displayed.

Parameters

family
Specifies the type of routing information distributed by the BGP instance.

Values

- ipv4** — Displays the BGP peers that have the IPv4 family enabled; peers capable of exchanging IP-VPN routes are not displayed.
- vpn-ipv4** — Displays the BGP peers that are IP-VPN capable.
- ipv6** — Displays the BGP peers that are IPv6 capable.
- mcast-ipv4** — Displays the BGP peers that are Mcast IPv4 capable.
- mvpn-ipv4** — Displays the BGP peers that are MVPN IPv4 capable.

brief
Keyword that provides a summarized display of the set of peers tow hich a BGP route is advertised.

prefix
Specifies the type of routing information to display.

Values	
rd:[ip-address[/mask]]	
rd	ip-address:number1 as-number1:number2 as-number2:number3
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-ex*

Displays all routes with an AS path matching the specified regular expression *reg-exp*.

community *comm-id*

Displays all routes with the specified BGP community.

Values	[<i>as-number1:comm-val1</i> <i>ext-comm</i> <i>well-known-comm</i>]
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

rd

Specifies the route distinguisher.

Values	ip-addr:comm-val, 2byte-asnumber:ext-comm-val, 4byte-asnumber:comm-val

veid

Specifies the VE ID.

Values 0 to 4294967295

vpls-base-offset

Specifies the VPLS base offset value.

Values 0 to 4294967295

site-id

Specifies the site ID.

Values 0 to 4294967295

l2vpn-type

Specifies the L2 VPN type.

Values bgp ad, bgp-vpls, multi-homing

ms-pw

Displays routes for the MS-PW family.

Output

The following output is an example of BGP route information, and [Table 92: Output fields: BGP routes](#) describes the output fields.

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
Cluster       : No Cluster Members
Originator Id : None
Flags         : Used Valid Best Incomplete
AS-Path       : No As-Path
Peer Router Id : 10.20.1.1
Interface Name : ip-10.10.2.3
Aggregator    : None
MED           : None
-----
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
Cluster :No Cluster Members
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
Cluster :No Cluster Members
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
Cluster      : No Cluster Members
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
Originator Id : None
IPv4 Label   : 131065
Flags        : Used Valid Best IGP
Peer Router Id : 10.20.1.2
Interface Name : INT_T0_C3_D_1
Aggregator     : None
MED            : None
```

```

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
Atomic Aggr.      : Not Atomic
AIGP Metric       : None
Connector         : None
Community         : No Community Members
Cluster           : No Cluster Members
Originator Id     : None
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
Atomic Aggr.      : Not Atomic
AIGP Metric       : None
Connector         : None
Community         : No Community Members
Cluster           : No Cluster Members
Originator Id     : None
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

Interface Name : INT_T0_C4_E_1
Aggregator     : None
MED            : None

Peer Router Id : 10.20.1.4

Interface Name : INT_T0_C1_A
Aggregator     : None
MED            : None

Peer Router Id : 10.20.1.1

```

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
Interface Name    : NotAvailable

```

Aggregator AS : None	Aggregator : None
Atomic Aggr. : Not Atomic	MED : 100
AIGP Metric : 555	
Connector : None	
Community : No Community Members	
Cluster : 10.20.1.3	
Originator Id : 10.20.1.2	Peer Router Id : 10.20.1.4
IPv4 Label : 131065	
Origin : IGP	
AS-Path : 400 500	
Route Tag : 0	
Neighbor-AS : 400	
Orig Validation: NotFound	
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	Interface Name : NotAvailable
Aggregator AS : None	Aggregator : None
Atomic Aggr. : Not Atomic	MED : 100
AIGP Metric : 555	
Connector : None	
Community : No Community Members	
Cluster : 10.20.1.3	
Originator Id : 10.20.1.2	Peer Router Id : 10.20.1.2
IPv4 Label : 131065	
Origin : IGP	
AS-Path : 400 500	
Route Tag : 0	
Neighbor-AS : 400	
Orig Validation: NotFound	
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	Interface Name : NotAvailable
Aggregator AS : None	Aggregator : None
Atomic Aggr. : Not Atomic	MED : None
AIGP Metric : 555	
Connector : None	
Community : No Community Members	
Cluster : 10.20.1.3	
Originator Id : 10.20.1.2	Peer Router Id : 10.20.1.5
IPv4 Label : 131065	
Origin : IGP	
AS-Path : 400 500	
Route Tag : 0	
Neighbor-AS : 400	
Orig Validation: NotFound	
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	Interface Name : NotAvailable
Aggregator AS : None	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:7210SAS# show router bgp routes mvpn-ipv4
=====
BGP Router ID:16.16.16.16      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

=====
BGP MVPN-IPv4 Routes
=====
Flag RouteType OriginatorIP LocalPref MED
RD SourceAS VPNLabel
Nexthop SourceIP
As-Path GroupIP
-----
u*>i Intra-Ad 10.17.17.17 100 0
1:2 - -
17.17.17.17 - -
No As-Path - -
*i Source-Join 100 0
1:2 100 -
1.1.1.1 10.1.1.2
No As-Path 239.1.1.1
*i Source-Join - 100 0
1:2 100 -
2.2.2.2 10.1.1.2
No As-Path 239.1.1.1
*i Source-Join - 100 0
1:2 100 -
1.1.1.1 10.1.1.2
No As-Path 239.1.1.2
*i Source-Join - 100 0
1:2 100 -
2.2.2.2 10.1.1.2
No As-Path 239.1.1.2

*A:7210SAS#

*A:praragon-sim1# show router bgp routes mvpn-ipv4 type source-join
source-as 200 source-ip 10.100.1.2 group-ip 239.0.0.0 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
Aggregator AS   : None
Atomic Aggr.    : Not Atomic
Community       : target:10.20.1.3:2
Cluster         : No Cluster Members
Originator Id   : None
Flags           : Used Valid Best IGP
AS-Path         : No As-Path
Interface Name  : NotAvailable
Aggregator      : None
MED             : 0
Peer Router Id  : 10.20.1.4
-----
Routes : 1
=====
*A:praragon-sim1#

```

Table 92: Output fields: BGP 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, 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 using a route policy.
Flag	u — used s — suppressed

Label	Description
	h — history d — decayed * — valid i — igp e — egp ? — incomplete > — best
Aggregator AS	The aggregator AS value none — aggregator AS attributes are not present
Aggregator	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
Cluster	Displays the route reflector cluster list
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

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 the BGP peers that are IP-VPN capable.
- neighbor ip-address**
Clears damping information for entries received from the BGP neighbor.
- Values** ipv4-address: a.b.c.d

Output

The following output is an example of summary BGP neighbor information, and [Table 93: Output fields: 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 Supressed 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 Active Rts.: 0
Total VPN Remote Rts. : 0      Total VPN Hist. Rts. : 0
Total VPN Supp. 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 Rts. : 0          Total VPN Remote Active Rts.: 0
Total VPN Supp. Rts.  : 0          Total VPN Hist. 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)
-----
10.3.3.3      20        0    0    01h55m56s Active
                0        0
=====
A:SetupCLI>show>router#

```

Table 93: Output fields: BGP summary

Label	Description
BGP Router ID	Displays the local BGP router ID

Label	Description
AS	Displays the configured autonomous system number
Local AS	Displays the configured local AS setting. If not configured, 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
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 as a result of route damping
Total History Routes	Displays the total number of routes with history as a result of route damping
Total Decayed Routes	Displays the total number of decayed routes as a result of 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

Label	Description
Total VPN Supp.Rts.	Displays the total number of suppressed VPN routes as a result of route damping
Total VPN Hist. Rts.	Displays the total number of VPN routes with history as a result of route damping
Total VPN Decay Rts.	Displays the total number of decayed routes as a result of 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
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 current state (if not established) or the number of received routes, active routes and sent routes (if established)

6.23.2.4 Clear commands

damping

Syntax

damping [{*ip-prefix/ip-prefix-length*] [*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/ip-prefix-length

Clears damping information for entries that match the IP prefix and prefix length.

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

neighbor ip-address

Clears damping information for entries received from the BGP neighbor.

Values	<i>ipv4-address:</i>	a.b.c.d

group name

Clears damping information for entries received from any BGP neighbors in the peer group.

Values	32 characters maximum

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 prefix and mask length.

Values	<i>ip-prefix:</i>	a.b.c.d (host bits must be 0)
	<i>mask:</i>	0 to 32

neighbor *ip-address*

Clears route flap statistics for entries received from the specified BGP neighbor.

Values ipv4-address: a.b.c.d

group *group-name*

Clears route flap statistics for entries received from any BGP neighbors in the specified peer group.

regex *reg-exp*

Clears route flap statistics for all entries that have the regular expression and the AS path that matches the regular expression.

policy *policy-name*

Clears route flap statistics for entries that match the specified route policy.

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 that resets all EBGp neighbors.

all

Keyword that resets all BGP neighbors.

soft

Keyword that specifies the BGP neighbors reevaluate all routes in the local-RIB against the configured export policies.

soft-inbound

Keyword that specifies BGP neighbors reevaluate all routes in the RIB-In against the configured import policies.

statistics

Keyword that clears the BGP neighbor statistics.

end-of-rib

Keyword that clears 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.

6.23.2.5 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.

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 events affecting only the specified BGP neighbor.

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

group *name*

Debugs events affecting only the specified peer group and associated neighbors.

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 events affecting only the specified BGP neighbor.

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

group *name*

Debugs events affecting only the specified peer group and associated neighbors.

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 events affecting only the specified BGP neighbor.

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

group *name*

Debugs events affecting only the specified peer group and associated neighbors.

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 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 events affecting only the specified BGP neighbor.

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

group *name*

Debugs events affecting only the specified peer group and associated neighbors.

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 events affecting only the specified BGP neighbor.

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

group name

Debugs events affecting only the specified peer group and associated neighbors.

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 events affecting only the specified BGP neighbor.

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

group name

Debugs events affecting only the specified peer group and associated neighbors.

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 events affecting only the specified BGP neighbor.

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

group name

Debugs events affecting only the specified peer group and associated neighbors.

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 events affecting only the specified BGP neighbor.

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

group *name*

Debugs events affecting only the specified peer group and associated neighbors.

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 events affecting only the specified BGP neighbor.

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

group *name*

Debugs events affecting only the specified peer group and associated neighbors.

7 Route policies

This chapter provides information about configuring route policies.

7.1 Configuring route policies

The Nokia 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 does not place routes associated with a specific origin in the routing table. Those routes are not 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.

7.1.1 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 source and destination.

7.1.1.1 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 default actions are applied in the following cases:

- If a route policy does not specify a matching condition, all the routes being compared with the route policy are considered to be matches.
- If 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.

7.1.1.2 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.

7.1.1.3 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.

The Nokia implementation of route damping is based on the following parameters:

- **Figure of Merit**

A route is assigned a Figure of Merit (FoM), which is proportional to the frequency of flaps. FoM should be able to characterize a route 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 to the route could experience unstable routing as a result 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.

The two 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 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 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.

7.2 Regular expressions

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

A regular expression is expressed in terms of terms and operators. A term for an AS path regular expression is:

- regular expressions should always be enclosed in quotes
- an elementary term; for example, an ASN "200"
- a range term composed of two elementary terms separated by the '-' character like "200-300"
- the '.' dot wild-card character which matches any elementary term
- a regular expression enclosed in parenthesis "(" ")"
- 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:

- an elementary term which for a community string is any single digit like "4"
- a range term composed of two elementary terms separated by the '-' character like "2-3"
- a colon ':' to delimit the ASN from the community value
- the '.' dot wild-card character which matches any elementary term or ':'
- a regular expression enclosed in parenthesis "(" ")"
- 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 94: 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.
{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.

Operator	Description
\	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 95: AS path and community regular expression examples

AS path to match criteria	Regular expression	Example matches
Null AS path	null ¹⁷	Null AS path
AS path is 11	11	11
AS path is 11 22 33	11 22 33	11 22 33
Zero or more occurrences of ASN 11	11*	Null AS path 11 11 11 11 11 11 11 ... 11
Path of any length that begins with ASNs 11, 22, 33	11 22 33 .*	11 22 33 11 22 33 400 500 600
Path of any length that ends with ASNs 44, 55, 66	.* 44 55 66	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 ASNs 100 and 200, followed by one or more occurrences of the number 33	100 200 33+	100 200 33 100 200 33 33 100 200 33 33 33 100 200 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

¹⁷ The null keyword matches an empty AS path.


AS path to match criteria	Regular expression	Example matches
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 ASNs 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 11 11 11 22 11 11 11 11 22
Path of any length, except nonexistent, whose second ASN can be anything, including nonexistent	. .* or . .{0,}	100 100 200 11 22 33 44 55
ASN is 100. Community value is 200.	^100:200\$	100:200

AS path to match criteria	Regular expression	Example matches
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 ...

7.2.1 BGP and OSPF route policy support

OSPF and BGP requires route policy support. [Figure 30: BGP route policy diagram](#) and [Figure 31: OSPF route policy diagram](#) show where route policies are evaluated in the protocol. [Figure 30: BGP route policy diagram](#) shows BGP which applies a route policy as an internal part of the BGP route selection process. [Figure 31: 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 30: BGP route policy diagram




7.2.1.1 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.

Figure 31: OSPF route policy diagram



7.2.1.2 Re-advertised route policies

Occasionally, BGP routes may be readvertised from BGP into OSPF, IS-IS. OSPF export policies (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.

7.2.2 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 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, AS path, or community values to manipulate the control the route selection.
- When you want to control BGP route flapping (damping).

7.3 Route policy configuration process overview

The following figure shows the process to provision basic route policy parameters.

Figure 32: Route policy configuration and implementation flow



7.4 Configuration notes

This section describes route policy configuration restrictions.

7.4.1 General

When configuring policy statements, the policy statement name must be unique.

7.5 Configuring route policies with CLI

This section provides information to configure route policies using the command line interface.

7.6 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.

7.6.1 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 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:

- 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)
- damping

Before a route policy is applied, analyze the policy 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.

7.6.2 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 96: 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.	• 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 cannot be controlled using route policies	• 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.	• Internal routes: By default, all active BGP routes are advertised to BGP peers

Protocol	Import	Export
		External routes: By default all non-BGP learned routes are not advertised to BGP peers.
RIP	By default, all RIP-learned routes are accepted.	- Internal routes: By default, directly connected interfaces are not advertised using RIP unless a routing policy is attached and exported to the peers. By default, only RIP learned routes are advertised to RIP neighbors. - External routes: By default, all non-RIP learned routes are not advertised to RIP peers.

7.6.3 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.

[Figure 33: Route policy process example](#) 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 **action** command. Policies can be constructed to support multiple states to the evaluation and setting of various route attributes.


[Figure 34: Next policy logic example](#) shows the next-policy and next-entry route processes.

Figure 33: Route policy process example



sw0489

Figure 34: Next policy logic example



sw0496

7.6.4 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 as a result of network problems that 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 35: Damping example



7.7 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 the following.

A policy statement with the following parameters specified:

- At least one entry
- Entry action

Example

The following is a sample 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#
```

7.8 Configuring route policy components

The following section describes the syntax used to configure the route policy components.

7.8.1 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
```

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

The following shows the command usage to configure a policy statement. 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.

7.8.2 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.

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

```
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#
```

7.8.3 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. The default action is applied only to those routes that do not match any policy entries.

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.

Example

The following is a sample default action configuration output.

```
A:ALU-7210>config>router>policy-options# info
-----
        policy-statement "1"
        default-action accept
        metric set 10
        exit
        exit
-----
A:ALU-7210>config>router>policy-options#
```

7.8.4 Configuring an entry

An entry action must be specified. The other parameters in the **entry action** context are optional. See the [Route policy command reference](#) for the commands and syntax.

Example

The following are sample entry parameters and include the default action parameters, which were shown in the previous section.

```
A:ALA-B>config>router>policy-options# info
-----
    policy-statement "1"
      entry 1
        to
          neighbor 10.10.10.104
        exit
        action accept
        exit
      exit
      entry 2
        from
          protocol ospf 1
        exit
        to
          protocol ospf
          neighbor 10.10.0.91
        exit
        action accept
        exit
      exit
      default-action accept
      . . .
    exit
  exit
-----
A:ALA-B>config>router>policy-options#
```

7.8.5 Configuring damping

For each damping profile, all parameters must be configured.

The **suppress** value must be greater than the **reuse** value (see [Figure 35: 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 is 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

The following is a sample 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#
```

7.8.5.1 Configuring a prefix list

Example

The following is a sample 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#
```

7.9 Route policy configuration management tasks

This section describes the route policy configuration management tasks.

7.9.1 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

The following is a sample 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
-----
```

7.9.2 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

The following shows the 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
```

7.9.3 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

The following shows the command usage to delete a policy statement.

```
config>router>policy-options# begin
policy-options# no policy-statement 1
policy-options# commit
```

7.10 Use of route policies for IGMP filtering

Example

The following is a sample route policy configuration output that can be used 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#
```

7.11 Route policy command reference

7.11.1 Command hierarchies

- [Route policy configuration commands](#)
- [Show commands](#)

7.11.1.1 Route policy configuration commands

```

config
- [no] router
- [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
-     next-hop ip-address
-     no next-hop
-     [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
    - next-hop ip-address
    - no next-hop
    - [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
  - community name
  - no community
  - [no] external
  - family [ipv4] [ipv6] [vpn-ipv4] [vpn-ipv6] [l2-vpn] [ms-pw] [mvpn-
ipv4]
  - 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-id}]
  - [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-id}]

```

```
- 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]
```

7.11.1.2 Show commands

```
show
- router router-name
- policy [name | prefix-list name | admin]
```

7.11.2 Command descriptions

- [Generic commands](#)
- [Route policy options](#)
- [Route policy damping commands](#)
- [Route policy prefix commands](#)
- [Route policy entry match commands](#)
- [Route policy action commands](#)
- [Show commands](#)

7.11.2.1 Generic commands

abort

Syntax

abort

Context

config>router>policy-options

Platforms

Supported on all 7210 SAS platforms as described in this document, including platforms configured in the access-uplink operating mode

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, including platforms configured in the access-uplink operating mode

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, including platforms configured in the access-uplink operating mode

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, including platforms configured in the access-uplink operating mode

Description

This command creates a text description that 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.

7.11.2.2 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 configures 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.

regular-expression

Specifies the AS path regular expression (any string or **null**).

- Values**
- 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.
- 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 configures 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. Up to 15 community ID strings can be specified, to a maximum of 72 characters.
- Values**
- 72 chars max

2byte-asnumber:comm-val | reg-ex | ext-comm | well-known-comm

ext-comm

type:{ip-address:comm-val | reg-ex1&reg-ex2 | ip-address&reg-ex2 | 2byte-asnumber:ext-comm-val |4byte-asnumber:comm-val}

2byte-asnumber	0 to 65535
comm-val	0 to 65535
reg-ex	72 chars max
type	target, origin
ip-address	a.b.c.d
ext-comm-val	0 to 4294967295
4byte-asnumber	0 to 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* — Specifies the *as-num* is the Autonomous System (AS) number.

Values:

as-num: 1 to 65535
comm-value: 0 to 65535

- type **{target | origin}**:*as-num:comm.-value* — 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* allow the same values as described previously for regular community values.
- reg-ex1 reg-ex2* — Specifies 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, etc.), the entire string must be enclosed within double quotes.
- well-known-comm* — Keywords are the following: **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, including platforms configured in the access-uplink operating mode

Description

Commands in this context configure route policies.

In the access-uplink operating mode, route policies are 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 reevaluation.

By default, when a change is made to a policy in the **config>router>policy>options** context and then committed, the change is effective immediately. There may be circumstances when the changes should or must be delayed; for example, if a policy change is implemented that would effect 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.

When this 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. When a **triggered-policy** is enabled, any routine policy change or policy assignment change within the protocol does not take effect until the protocol is reset or a clear command is issued to reevaluate 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.

7.11.2.3 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 configures 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 value for the route damping profile.

The half life value is the time required for a route to remain stable in order for the Figure of Merit (FoM) value to be reduced by one half; for example, if the half life value is 6 (minutes) and the route remains stable for 6 minutes, 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.

The **no** form of this command removes the half life value 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 value for the route damping profile.

This value indicates the maximum time that a route can remain suppressed.

The **no** form of this command removes the maximum suppression value 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 value 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 value 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 value from the damping profile.

Parameters

integer

Specifies the suppress value, expressed as a decimal integer.

Values 1 to 20000

7.11.2.4 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 a prefix list entry, in dotted decimal notation.

Values	
ipv4-prefix:	a.b.c.d (host bits must be 0)
ipv4-prefix-length:	0 to 32
ipv6-prefix - x:x:x:x:x:x:x (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-le - [0 to 128]

<exact|longer|thro*> : keyword

<length> : [0 to 128] (prefix-length <= length)

<length1-length2> : length1/length - [0 to 128] (prefix-length <= length1 <=length2)

exact

Keyword to specify the prefix list entry only matches the route with the specified *ip-prefix* and prefix *mask* (length) values.

longer

Keyword to specify 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 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 specific range. The range is inclusive of start and end values.

Values 0 to 32, *length2* > *length1*

7.11.2.5 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, including platforms configured in the access-uplink operating mode

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 matching criteria defined (in which case, everything matches) but must at least define an action to be considered complete. Entries without an action are considered incomplete and are 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 needing to renumber 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, including platforms configured in the access-uplink operating mode

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] [ipv6] [vpn-ipv4] [vpn-ipv6] [l2-vpn] [mvpn-ipv4] [ms-pw]

no family

Context

config>router>policy-options>policy-statement>entry>from

Platforms

Supported on all 7210 SAS platforms as described in this document, including platforms configured in the access-uplink operating mode

Description

This command specifies address families as matching conditions.

Parameters

ipv4

Keyword to specify IPv4 routing information.

ipv6

Keyword to specify IPv6 routing information. This keyword is not supported on 7210 SAS platforms configured in the access-uplink operating mode.

vpn-ipv4

Keyword to specify IPv4 VPN routing information. This keyword is not supported on 7210 SAS platforms configured in the access-uplink operating mode.

vpn-ipv6

Keyword to specify IPv6 VPN routing information. This keyword is not supported on 7210 SAS platforms configured in the access-uplink operating mode.

l2-vpn

Keyword that exchanges Layer 2 VPN information. This keyword is not supported on 7210 SAS platforms configured in the access-uplink operating mode.

mvpn-ipv4

Keyword that exchanges multicast VPN related information. Supported on 7210 SAS-T, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE devices only. This keyword is not supported on 7210 SAS platforms configured in the access-uplink operating mode.

ms-pw

Keyword to specify ms-pw routing information. This keyword is not supported on 7210 SAS platforms configured in the access-uplink operating mode.

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 by 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 *ip-address:* a.b.c.d
 0 to 4294967295

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, including platforms configured in the access-uplink operating mode

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, including platforms configured in the access-uplink operating mode

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 that matches 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-addr

Specifies the neighbor IP address, in dotted decimal notation.

Values

ipv4-address: a.b.c.d

ipv6-address - x:x:x:x:x:x[-interface]

 x:x:x:x:x.d.d.d[-interface]

 x - [0 to FFFF]H

 d - [0 to 255]D

interface - 32 chars max, mandatory for link local addresses

name : [32 chars max]

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 that configures matching path information originating within the local AS.

egp

Keyword that configures matching path information originating in another AS.

incomplete

Keyword that configures matching path information learned by another method.

any

Keyword that specifies 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, including platforms configured in the access-uplink operating mode

Description

Commands in this context configure a route policy statement. The policy statement is a logical grouping of match and action criteria. The processing action taken is determined by the action associated with the entries configured in the policy statement.

In the access-uplink operating mode, route policy statements enable appropriate processing of IGMP group membership reports received from hosts.

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, including platforms configured in the access-uplink operating mode

Description



Note:

The **config>router>policy-options>policy-statement>entry>to** context is not supported on 7210 SAS platforms configured in the access-uplink operating mode.

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-id*}]

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 use as the match criterion.

Values direct, static, bgp, isis, ospf, rip, aggregate, bgp-vpn, igmp, ospf3, ldp, periodic

 The **rip** protocol value is supported only on 7210 SAS-Mxp.

instance-id
Specifies the OSPF or IS-IS instance.

Values 1 to 31

all
Keyword to specify OSPF- or ISIS-only.

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, including platforms configured in the access-uplink operating mode

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.

Values

 ipv4-address - a.b.c.d

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

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

 x - [0 to FFFF]H

d - [0 to 255]D

tag

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 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 criterion 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

type

Specifies to match OSPF routes with LSAs.

Values **1** — Matches OSPF routes with type 1 LSAs
 2 — Matches OSPF routes with type 2 LSAs

7.11.2.6 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, including platforms configured in the access-uplink operating mode

Description

Commands in this context configure actions to take for routes matching a route policy statement entry.

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

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

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

Default

no action

Parameters

accept

Keyword to specify routes matching the entry match criteria is accepted and propagated.

next-entry

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

next-policy

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

reject

Keyword to specify routes matching the entry match criteria would be rejected.

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

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12

Description

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 before 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 {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

Keyword to specify that the AS path list is to be prepended to an existing AS list.

replace

Keyword to specify that the 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

The command prepends a BGP AS number once or numerous times to the AS path attribute of routes matching the route policy statement entry.

If an AS number is not configured, the AS path is not changed.

If the optional *number* is specified, the AS number is prepended as many times as indicated by the number.

The **no** form of this command disables the AS path prepend action from the route policy entry.

Default

no as-path-prepend

Parameters

as-num

Specifies the AS number to prepend, expressed as a decimal integer.

Values 1 to 4294967295

repeat

Specifies the number of times to prepend the specified AS number, 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 that adds the specified community list to an existing list of communities.

remove

Keyword that removes the specified community from the existing list of communities.

replace

Keyword to specify that the community list replaces an 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 criterion 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 that disables 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, including platforms configured in the access-uplink operating mode

Description

This command configures actions for routes that do not match any route policy statement entries when the **accept** parameter is specified.

The **default-action** command can be set to all available action states. If the action states are accepted or rejected, the policy evaluation terminates and a result is returned.

If a default action is defined and no matches occur with the entries in the policy, the default action is used.

If a default action is defined and one or more matches occur with the entries of the policy, 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 routes matching the entry match criteria are accepted and propagated.

next-entry

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

reject

Keyword to specify routes matching the entry match criteria would 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 is advertised.

The value assigned to 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 specify *integer* is added to any existing metric. If the result of the addition results in a number greater than 4294967295, the value 4294967295 is used.

subtract

Keyword to specify *integer* is subtracted from any existing metric. If the result of the subtraction results in a number less than 0, the value of 0 is used.

set

Keyword to specify *integer* replaces any existing metric.

metric

Specifies the metric modifier, expressed as a decimal integer.

Values 0 to 4294967295

next-hop

Syntax

next-hop *ip-address*

no next-hop

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 the specified next-hop IP address to routes matching the policy statement entry.

If a next-hop IP address is not specified, the next-hop attribute is not changed.

The **no** form of this command disables assigning a next-hop address in the route policy entry.

Default

no next-hop

Parameters

ip-address

Specifies the next-hop IP address in dotted decimal notation.

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

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** command 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 that sets the path information as originating within the local AS.

egp

Keyword that sets the path information as originating in another AS.

incomplete

Keyword that sets 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 when more than two octets are configured.

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

Default

no tag

Parameters

tag

Specifies to assign an OSPF or ISIS tag to routes matching the entry.

Values

Accepts decimal or hex formats:

OSPF and ISIS:	[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 within 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

7.11.2.7 Show commands

```
policy
```

Syntax

```
policy [name | damping | prefix-list name | as-path name | community 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.
- damping**

Displays the damping profile for use in the route policy.
- prefix-list name**

Displays the prefix lists configured in the route policy.
- as-path name**

Displays AS path regular expression statements used in the route policy.
- community name**

Displays community lists used in the route policy.
- admin**

Keyword that displays 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 statement information, and [Table 97: Output fields: route policy](#) describes the output fields.

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#

```

Sample output: show router policy name

The **show router policy name** command displays information about a specific route policy.

```

A:ALA-1# show router policy "OSPF To OSPF"
    entry 10
        description "Entry For Policy Statement OSPF To OSPF"
        from
            protocol ospf
        exit
        to
            protocol ospf
        exit
        action accept
            tag 100
        exit
    exit
    default-action reject
ALA-1#

```

Sample output

```

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

```

```

protocol rip
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
protocol rip
no neighbor
no level
no prefix-list
exit

```

Table 97: Output fields: route 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

8 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.

8.1 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

8.2 Certificate Management

RFC 7468, Textual Encodings of PKIX, PKCS, and CMS Structures 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



Note:

Only PKIX and PKCS is supported.

8.3 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, Sx/S-1/10GE, and Sx/S-1/10GE-VC



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, Sx/S-1/10GE, and Sx/S-1/10GE-VC



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

8.4 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

8.5 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.

8.6 Internet Protocol (IP) — General

draft-grant-tacacs-02, The TACACS+ Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-vrrp-unified-spec-02, Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6 is supported on Mxp

RFC 768, User Datagram Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 793, Transmission Control Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 854, TELNET Protocol Specifications is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 951, Bootstrap Protocol (BOOTP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1034, Domain Names - Concepts and Facilities is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1035, Domain Names - Implementation and Specification is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1350, The TFTP Protocol (revision 2) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1534, Interoperation between DHCP and BOOTP is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1542, Clarifications and Extensions for the Bootstrap Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2131, Dynamic Host Configuration Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 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 5246, The Transport Layer Security (TLS) Protocol Version 1.2 – TLS client, RSA public key 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



Note:

Only TLS client.

RFC 5425, Transport Layer Security (TLS) Transport Mapping for Syslog – RFC 3164 with TLS 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



Note:

Only TLS client.

RFC 5656, Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer – ECDSA 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 7301, Transport Layer Security (TLS) Application Layer Protocol Negotiation Extension 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



Note:

Only TLS client.

RFC 8446, The Transport Layer Security (TLS) Protocol Version 1.3 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



Note:

Only TLS client.

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

8.7 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.

8.8 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.

8.9 IP — Version 6

RFC 2464, Transmission of IPv6 Packets over Ethernet 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



Note:

D, Dxp, K5, 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, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, K5, and T(A) for Management only.

RFC 4007, IPv6 Scoped Address 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



Note:

D, Dxp, K5, and T(A) for Management only.

RFC 4193, Unique Local IPv6 Unicast 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



Note:

D, Dxp, K5, and T(A) for Management only.

RFC 4291, Internet Protocol Version 6 (IPv6) Addressing 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



Note:

D, Dxp, K5, 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, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, K5, and T(A) for Management only.

RFC 4861, Neighbor Discovery for IP version 6 (IPv6) is supported on 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:

D, Dxp, K5, and T(A) for Management only.

RFC 4862, IPv6 Stateless Address Autoconfiguration (Router Only) is supported on 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:

D, Dxp, K5, and T(A) for Management only.

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, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

D, Dxp, K5, 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

8.10 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.

8.11 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

8.12 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

8.13 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

8.14 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

8.15 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.

8.16 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

8.17 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

8.18 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

8.19 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

8.20 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

8.21 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

8.22 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.

8.23 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, 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

8.24 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](#)